European Commission

Proceedings of the 2016 conference on

Big Data from Space

(BiDS'16)

15th-17th March 2016
Santa Cruz de Tenerife (Spain)

Edited by P. Soille and P.G. Marchetti

SatCen

European Commission

esa
European Space Agency

Joint Research Centre

EUR 27775 EN

Contact information
Name: Pierre Soille
Address: Joint Research Centre, Via Enrico Fermi 2749, TP 267, I-21027 Ispra (VA), Italy
E-mail: Pierre.Soille@jrc.ec.europa.eu
Tel.: +39 0332 78 9111

JRC Science Hub
https://ec.europa.eu/jrc

JRC100655

EUR 27775 EN

ISBN 978-92-79-56980-7 (PDF)
ISSN 1831-9424 (online)

doi:10.2788/854791 (online)

How to cite: P. Soille and P.G. Marchetti (Eds.) Proc. of the 2016 conference on Big Data from Space (BiDS'16), EUR 27775 EN, doi:10.2788/854791, 2016

Preface

Big Data from Space refers to Earth and Space observation data collected by space-borne and ground-based sensors. Whether for Earth or Space observation, they qualify being called 'big data' given the sheer volume of sensed data (archived data reaching the exabyte scale), their high velocity (new data is acquired almost on a continuous basis and with an increasing rate), their variety (data is delivered by sensors acting over various frequencies of the electromagnetic spectrum in passive and active modes), as well as their veracity (sensed data is associated with uncertainty and accuracy measurements). Last but not least, the value of Big Data from Space depends on our capacity to extract information and meaning from them.

Big Data from Space is further gaining momentum given the sharp increase in volume, velocity, and variety of data acquired from space-borne and ground-based sensors. Fortunately, this increase is paralleled by a tremendous amount of new developments related to big data in other fields and enabled by technological breakthroughs and new challenges in hardware and software developments, high performance/throughput computing, cloud computing, data science, and visualisation. In addition, the recent multiplication of free and open access initiatives to Big Data from Space is giving impulse to the field by widening substantially the spectrum of users as well as awareness among the public while offering new opportunities for scientists and value-added companies.

The goal of the Big Data from Space conference is to bring together researchers, engineers, developers, and users in the area of Big Data from Space. Following the huge success of the 2014 conference on Big Data from Space held at the ESA Centre for Earth Observation (ESRIN), the 2016 edition (BiDS'16) is again co-organised by ESA, the Joint Research Centre (JRC) of the European Commission, and the European Union Satellite Centre (SatCen). It was held at the Auditorio de Tenerife (Santa Cruz de Tenerife, Spain) from the 15th to the 17th of March 2016. The Teide Observatory on Tenerife together with the Roque de los Muchachos Observatory on the nearby La Palma highland rank among the best places on Earth for ground-based space observations. The ESA Optical Ground Station (OGS) telescope at the Teide observatory has been extensively used for experiments in laser communication now routinely used on the first European Data Relay Satellite (EDRS-A) launched on the 29th of January 2016. Dubbed the 'Space Data Highway', EDRS is revolutionising satellite communications as Europe's first optical space communication network, capable of relaying user data in near-real time at an unprecedented 1.8 Gbit/s. This Space Data Highway is a central component of the Big Data from Space given the sheer amount of data that is and will be generated by the Sentinel satellites operated by ESA in the framework of the Copernicus programme funded and managed by the European Commission. Indeed, since the last edition of the BiDS conference in November 2014, two Sentinel satellites have been launched: Sentinel-2A and Sentinel-3A launched on the 29th of June 2015 and the 4th of February 2016 respectively. They will deliver 0.8 TB (resp. 0.3 TB) of image data per day in full operational capacity and are key enablers of the Copernicus services and in particular those related to land and marine monitoring. They complement the Sentinel-1A synthetic aperture radar satellite launched in April 2014. The Sentinel 1, 2 and 3 are each designed as a two-satellite constellation. The launch of the first "B" unit, Sentinel-1B, is scheduled for spring 2016. Regarding space observations, virtual observatories are gaining increased attention as well as the GAIA astrometry mission launched in December 2013 and steadily populating its catalogue towards 1–2 billion sources by the end of the mission in 2022.

The focus of BiDS'16 is on the whole data life cycle, ranging from data acquisition by space borne and ground-based sensors to data management, analysis and exploitation in the domains of Earth Observation, Space Science, Solar System Objects, Space Situational Awareness, Secure Societies, etc. Special emphasis is put on highlighting synergies and cross-fertilisation opportunities. The main objectives of BiDS'16 are:

- Identify priorities for research, technology development and innovation;
- Widen competences and expertise of universities, research institutes, labs, SMEs and industrial actors;
- Foster networking of experts and users towards better access and sharing of data, tools and resources;
- Leverage innovation, spin-in and spin-off of technologies, and business development arising from research and industry progress;
- Increase and promote the value stemming from the huge quantity of data made available nowadays (and in the future);
- Contribute to the EO innovation for Europe, as one of the main pillars for the Ground Segment evolution strategy.

The presentations, discussions, and contacts established during the conference as well as the materials presented in these proceedings are contributing to these goals. A total of 111 papers were submitted for presentation at the conference. Following the peer-review process by members of the conference programme committee, 62 papers were selected for oral and 32 for poster presentations (for a total of 354 distinct co-authors with affiliations in 21 different countries). A demonstration/industry session (with 10 demos) has been organised along with the poster session, to give the possibility to present live demos on big data applications. Among the papers selected for oral presentation, 18 originate from 5 invited sessions devoted to a series of horizontal themes, namely, Organisational Challenges for Big Data Space Missions, Advances in Big Data using OGC Standards, Image Information and Data Mining for Big Data, Big Data Impact on Industry, and Data Science Skills.

The conference was inaugurated by the local authorities with an opening ceremony with a warm welcoming from the Tenerife Island Government, the University of La Laguna, the Canary Islands Institute of Astrophysics, and the City of Santa Cruz de Tenerife; followed by a keynote address on *SmartIsland and Big Data* by C. Alonso Rodríguez, Tenerife Island Government. The conference opening session was devoted to a series of enlightening talks from ESA, SatCen, and the European Commission:

- *Big Data from Space: the Copernicus contribution*
  by Pier Bargellini, European Space Agency, ESRIN
- *Big Data solutions to meet the CSDP operational needs in the Geoint field*
  by Pascal Legai, European Union Satellite Centre
- *INSPIRE Big Metadata: understanding implementation patterns*
  by Alessandro Annoni, European Commission, Joint Research Centre
- *Space data in data economy*
  by Marta Nagy-Rothengass, European Commission, DG CONNECT
- *Copernicus and Big Data: maximising the socio-economic benefits of Earth Observation in Europe*
  by Andreas Veispak, European Commission, DG GROW

Given the large number of high quality submissions, the 3 day conference programme was structured around 4 plenary sessions and 5 pairs of parallel sessions complemented by a poster and demonstration/industry session. These proceedings consist of a collection of 108 short papers corresponding to the oral and poster presentations presented at the conference (except for 2 oral-only presentations, see conference website). They are organised in sections matching the order of the conference sessions followed by the contributions that were presented during the poster session, also organised by topics. They

provide a snapshot of the current research activities, developments, and initiatives in Big Data from Space.

As tradition, the BiDS conference hosts contributions addressing the whole life-cycle of Big Data from Space, from university education to on board processing, down to image and data analysis, until downstream exploitation. In this latest BiDS conference edition, numerous contributions are addressing the infrastructures and platforms enabling to exploit the value behind the volume, velocity, and variety of Big Data from Space. This tendency is revealed by 2 oral sessions on Exploitation Platforms and a poster section on Architectures, Platforms, and Cloud Solutions. The growing interest in private and public cloud solutions reveals the increasing interest in bringing the processing closer to the data. The topics of Data Science, Data Analytics, Machine Learning, Toolboxes, and Visualisation are naturally gaining attention. Indeed, while a platform or infrastructure is an essential requirement, it needs to be complemented by all these ingredients to produce valuable information from Big Data from Space for the benefit of its users. Other generic key aspects of big data are mirrored by their respective sections in these proceedings. Additional conference materials such as electronic version of the slides presented at the conference, including those regarding the opening session talks and keynote lectures, are available on the conference website, see http://congrexprojects.com/bids16 for updates.

A great thanks goes to all authors and presenters of BiDS'16 as well as the numerous participants (nearly 300 coming from almost 40 different countries). Together, they have ensured the success of the 2016 conference on Big Data from Space. A special thank goes to the Programme Committee members and the additional reviewers for their thorough reviews and detailed comments that were taken into account by the authors when preparing the final version of their paper included in these proceedings.

This edition of the conference is deeply grateful to the thorough local organising support it received from the Parque Científico y Tecnológico de Tenerife (PCTT) and Council of Tenerife. They have both been very supportive in having BiDS'16 hosted on the Island.

Pierre Soille and Pier Giorgio Marchetti

Organising Committee

Sergio Albani	European Union Satellite Centre (SatCen), Spain
Jorge Amador Monteverde	European Space Agency (ESA), ESTEC, The Netherlands
Christophe Arviset	European Space Agency (ESA), ESAC, Spain
Simon Baillarin	Centre National d'Études Spatiales (CNES), France
Peter Baumann	Jacobs University, Germany
Lorenzo Bruzzone	University of Trento, Italy
Mihai Datcu	Deutsches Zentrum für Luft- und Raumfahrt (DLR), Germany
Michele Iapaolo	Randstad at ESA-ESRIN, Italy
Johan Knapen	Instituto de Astrofísica de Canarias, Spain
Pier Giorgio Marchetti	European Space Agency (ESA), ESRIN, Italy
Pierre Soille	European Commission, Joint Research Centre (JRC), Italy
Juan Luis Valero	European Union Satellite Centre (SatCen), Spain

Local Organising Committee

Madelon van Oostrom	Parque Científico y Tecnológico de Tenerife (PCTT), Spain
Iñaqui Dominguez	Meeting in, Santa Cruz de Tenerife, Spain

Programme Chairs

Pier Giorgio Marchetti	European Space Agency (ESA), ESRIN, Italy
Pierre Soille	European Commission, Joint Research Centre (JRC), Italy

Programme Committee

Selim Aksoy	Bilkent University, Turkey
Sergio Albani	European Union Satellite Centre (SatCen), Spain
Christophe Arviset	European Space Agency (ESA), ESAC, Spain
Simon Baillarin	Centre National d'Études Spatiales (CNES), France
Peter Baumann	Jacobs University Bremen, Germany
Victoria Bennett	STFC - RAL Space, U.K.
Francesca Bovolo	Fondazione Bruno Kessler, Italy
Lorenzo Bruzzone	University of Trento, Italy
Mihai Datcu	Deutsches Zentrum für Luft- und Raumfahrt (DLR), Germany
Liang Feng	University of Edinburgh, U.K.
Jean-Pierre Gleyzes	Centre National d'Études Spatiales (CNES), France
Jose Gomez-Dans	NCEO/ University College London, U.K.
Jutta Graf	Deutsches Zentrum für Luft- und Raumfahrt (DLR), Germany
Jacopo Grazzini	European Commission, EUROSTAT, Luxembourg
Michele Iapaolo	European Space Agency (ESA), ESRIN, Italy
Johan Knapen	Instituto de Astrofísica de Canarias, Spain

Henri Laur	European Space Agency (ESA), ESRIN, Italy
Guido Lemoine	European Commission, Joint Research Centre (JRC), Italy
Thierry Levoir	Centre National d'Études Spatiales (CNES), France
Pier Giorgio Marchetti	European Space Agency (ESA), ESRIN, Italy
Pierre-Philippe Mathieu	European Space Agency (ESA), ESRIN, Italy
Osamu Ochiai	GEO Secretariat, Switzerland
Peter Reinartz	Deutsches Zentrum für Luft- und Raumfahrt (DLR), Germany
Sven Schade	European Commission, Joint Research Centre (JRC), Italy
John Schnase	NASA Goddard Space Flight Center, U.S.A.
Pierre Soille	European Commission, Joint Research Centre (JRC), Italy
Vasileios Syrris	European Commission, Joint Research Centre (JRC), Italy
Raffaele Vitulli	European Space Agency (ESA), ESTEC, The Netherlands
Wolfgang Wagner	TU Wien, Austria
Michael Wise	ASTRON, The Netherlands

Additional Reviewer

Armin Burger	European Commission, Joint Research Centre (JRC), Italy

Table of Contents

Exploitation Platforms 1

Exploitation Platforms 2

Data Science Skills

High Performance Computing and Optimisation

Big Data Analysis and Analytics

Big Data Applications

Time Series and Change Detection

Services and Apps

Posters: Architectures, Platforms, and Cloud Solutions

Posters: Data Catalogues and Management

Posters: Data Science, Data Analytics, and Visualisation

Posters: Toolboxes, Machine Learning, and Grid Processing

Posters: Change Detection and Multitemporal Analysis

Posters: Content Based Image Retrieval and Mining

Posters: Big Data Apps and Value

BIG DATA AND THE SYNERGY BETWEEN MISSIONS AT THE INFRARED PROCESSING AND ANALYSIS CENTER (IPAC)

H. I. Teplitz, G.B. Berriman, V. Desai, S. Groom, G. Helou, D. Imel, D. Kirkpatrick, J. Mazzarella, P. Ogle, X. Wu

IPAC / Caltech

ABSTRACT

The Infrared Processing and Analysis Center (IPAC) at Caltech provides science and data management for astronomy missions, planning for future projects, and vital science archives. IPAC supports NASA, NSF and privately funded projects, including: Spitzer, WISE, NEOWISE, the Zwicky Transient Facility, and the US centers for Planck, Herschel and Euclid. IPAC is developing science user tools for the LSST project and is planning for WFIRST science center activities. IPAC operates the NASA/IPAC Infrared Science Archive (IRSA), the NASA/IPAC Extragalactic Database (NED), the NASA Exoplanet Archive, and the Keck Observatory Archive.

By supporting many astronomy projects at one center, IPAC has the advantage of the synergy between them. We provide powerful services to researchers that combine data sets across the infrared sky. The breadth of current infrared data presents many of the challenges of Big Data. In total, IPAC manages a data center of 10 PB (with 4 PB of unique astronomical data), including databases containing almost 200 billion rows. The data volume is expected to double in the next five years.

1. INTRODUCTION

Since its inception in 1984 for the purpose of creating and supporting the IRAS archive, IPAC has evolved into a multi-dimensional science center, a true "Portal to the Universe" as described in the report of the NRC Committee on NASA Science Centers. IPAC today is an active hub of science and data management for astronomy missions, of planning for future projects, and of heavily used science archives. These tasks are carried out by a staff of scientists with active involvement in research, working closely with engineering and technical professionals strongly motivated by the science mission.

IPAC supports several missions in operation or under development. The Spitzer Science Center provides science operations to NASA's Infrared Great Observatory, now in its Warm Mission phase. The NASA Herschel Science Center (NHSC) supports the US community in using Herschel, an "ESA Corner-Stone Mission" with enabling instrument contributions by NASA. IPAC hosts the central node in the US Planck Data Center, and delivered in late 2010 the Planck Early Release Compact Source Catalog; Planck is also an ESA mission with critical NASA contributions. WISE (Wide-Field Infrared Survey Explorer), a NASA MIDEX launched in December 2009, mapped the sky 1.2 times at 3.4, 4.6, 12, and 22 μm, finished the second map in a warm mode, and was placed into hibernation on January 31, 2011; it was reactivated in August 2013 as NEOWISE, dedicated to surveying for asteroids. IPAC operates the WISE Science Data Center (WSDC).

IPAC also hosts four distinct science archives: IRSA (NASA/IPAC InfraRed Science Archive), NED (NASA/IPAC Extragalactic Database), the NASA Exoplanet Archive, and the Keck Observatory Archive (KOA) that serve overlapping research Communities. The archives work together to provide interoperable services. IPAC is home to the NASA Exoplanet Science Institute.

Looking to the future, we are implementing the Euclid NASA Science Center at IPAC (ENSCI); Euclid is an ESA-led Dark Energy mission expected to launch in 2020. IPAC is participating in the WFIRST formulation activities, planning for the WFIRST Science Center in collaboration with the Space Telescope Science Institute and Goddard Space Flight Center. In addition to the NASA projects, IPAC is implementing science user tools for the Large Synoptic Survey Telescope (LSST), and a science data system for the Zwicky Transient Facility (ZTF), a LSST precursor using the Palomar 48-inch Schmidt telescope.

By supporting many astronomy projects at one center, IPAC has the advantage of the synergy between them. We provide powerful services to researchers that combine data sets across the multi-wavelength sky.

2. BIG DATA AT IPAC

The science and data centers for individual missions at IPAC have generated large and all-sky data products across a broad range of wavelengths. In total, IPAC manages a data center of 10 PB (with 4 PB of unique astronomical data), including databases containing almost 200 billion rows. The data center holds more than 500 servers, with over 5000 cores, and computing power of more than 70,000 Gflops. IPAC has recently upgraded to 10 Gbps internal

network speed. IPAC's operations have grown by a factor of 8 in the last three years. In this section, we present specific examples of the multi-faceted big data at IPAC.

2.1. WISE and NEOWISE databases

One of the large data challenges has been processing and serving the data from the WISE and NEOWISE projects (for more details see Groom et al. [1]). The WSDC processed the full sky images and extracted from them measurements of source photometry in each single exposure frame. Catalog sources were cross-matched with 2MASS and merged photometry is presented in the table. From these millions of images and billions of measurements, the WISE Moving Object Processing System (WMOPS, [2]) identifies solar system objects based on their relative motion in the repeated observations. During the one-year WISE prime mission alone, the NEOWISE project measured over 158,000 asteroids including about 700 near earth objects.

The size of the WISE datasets and online nature of the WISE archive loading process posed new challenges. The pre-hibernation WISE data products include over 40 billion rows in a single table, and each release of the NEOWISE Reactivation phase includes 20 billion rows of single exposure measurements. The sheer size of the data presented a strain on our indexing systems. Complicating things further, with new additions arriving almost daily, these large data sets had to be ingested and indexing performed while the data already in the archive were kept fully online and being actively used. IRSA implemented innovative indexing techniques, optimized to meet the required use cases for database queries. These cases include single position spatial searches (using a recursively subdivided triangular mesh) and simultaneous matching of large user-supplied lists of positions (using a file-based index outside of the database).

The archive built for WISE needed to be extended from and compatible with IRSA's existing architecture. The challenges presented by WISE were used as opportunities to extend IRSA's capabilities, so that the capabilities developed for WISE could be leveraged for other projects. Notably, the ZTF will require databases that are at least an order of magnitude larger.

2.2. NASA astrophysics archives: diverse holdings

NASA's infrared data products are served through IRSA. Today, IRSA offers all-sky data in 20 photometric bands (Figure 1). IRSA currently curates about 800 TB of data, including about 100 million images and almost 100 billion rows of tabular data. The volume of IRSA holdings has doubled since 2012, and currently stands at more than 10 times the holdings in 2008. The holdings are expected to grow to about 1.2 petabytes by 2018. The diversity of

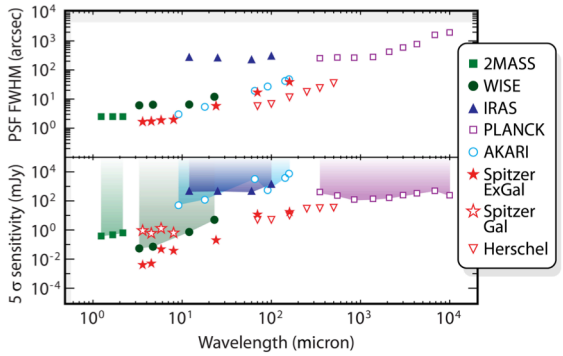

FIGURE 1. IRSA serves an unparalleled array of all-sky IR surveys providing a total of 20 bands from 1 micron to 10 mm. All-sky survey are shown as shaded regions.

IRSA's holdings across many missions contributes to the challenges of managing and presenting data effectively.

IRSA had over 13 million user queries in the 2014 calendar year. The growth in number of queries has followed a similar curve to the growth in holdings. In the past five years, IRSA has averaged almost 200 TB downloaded per year. The monthly rate is highly variable, with large spikes coinciding with major data releases.

NED provides efficient access to an extensive synthesis of extragalactic data combining and standardizing key measurements from NASA missions, sky surveys, and the astrophysics literature. It is the primary portal for systematically merged, value-added data and uniformly derived quantities across the electromagnetic spectrum for extragalactic objects. As the number, size, and complexity of data sets grow, the NED team keeps pace using increased automation guided by scientific expertise. A new rule-based, statistical tool for cross-matching sources enables efficient ingestion of large data sets from missions such as GALEX and WISE. Extragalactic papers have grown to 3,500 per year, with unique measurements for millions of objects.

3. VISUALIZATION

Visualization is an important component of the exploitation of big data in astronomy. To help users manage the complexity of diverse holdings (and keep development and maintenance costs under control), IPAC uses a common set of interface tools across many missions. These components (known as "Firefly by IPAC") offer enhanced visualization and data exploration options while still providing a common "look and feel." This technology has enabled IPAC to quickly and affordably create mission-specific interfaces for Spitzer, WISE, and Planck, and it has enabled a suite of multi-mission tools at IRSA. Unlike many traditional web

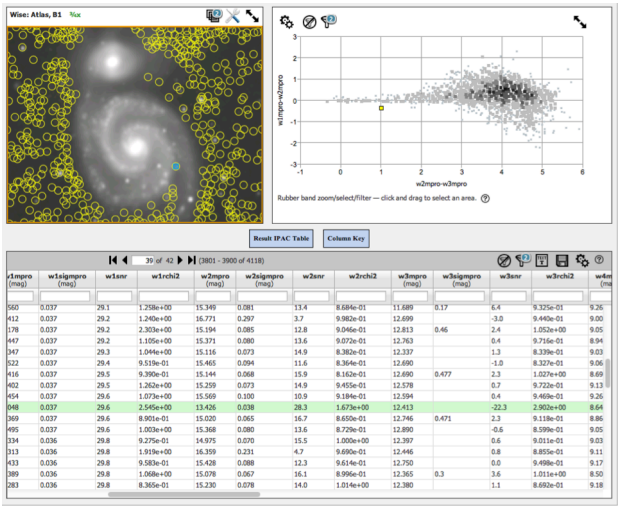

FIGURE 2: IRSA results visualization allows interactive exploration of tabular data, two dimensional plots, and images with source overlays.

applications where most of the processing occurs on the server, Firefly uses a "heavy client." This takes advantage of client desktop processing power, and allows for interactive features. The client-side is composed of dynamic HTML, CSS, and JavaScript to create a rich, interactive application. No plugins are required to use IRSA. Figure 2 presents an example of Firefly components used to provide visualization of a complex dataset [3].

Large archival data sets present a particular challenge. Even a thousand sources can make a catalog overlay on an image too confused to be useful. A scatter plot of a million sources may not be practical to display. IPAC is developing tools to enable users to explore large data sets and to refine their search criteria to identify a manageable subset of the data. Several features have been identified that are of general interest for future visualization tools:

- Plots and Catalog Overlays: The density of points in x-y plots should be visualized by user-selectable methods such as color, symbol size, or contours.
- Statistical methods: Data exploration tools must allow simple statistical views such as histograms.
- Tabular Data Exploration: Enhanced features for the exploration of large tables could include: user-selected reordering of columns; on-the-fly statistics on columns or combinations of columns; and more complex filtering ("ex post facto SQL").
- Multi-resolution Image Viewing: Exploring large surveys requires viewing sky maps at varying spatial resolution (to zoom in or out). There are two primary options to implement this service: (1) generating new sky maps at various resolutions, in steps of 2x; and (2) adapting existing sky-maps into Hierarchical Equal Area isoLatitude Pixelisation (HEALPix) format.

4. VIRTUAL OBSERVATORY

IPAC is committed to maintaining international interoperability between astrophysics archives. We are an active part of the NASA Astronomical Virtual Observatory (NAVO) and the International Virtual Observatory Alliance (IVOA). We offer millions of images and billions of rows of catalog data through VO-compliant program interfaces. Catalogs are available via the Simple Cone Search protocol and the Table Access Protocol (TAP). Images are available via the Simple Image Access (SIA) protocol. IPAC resources are discoverable in VO registries, allowing tools such as TOPCAT [4] access via the API.

5. DATA ANALYSIS TOOLS

As data sets grow in size relative to the local resources available to most researchers, it is the responsibility of the data centers to provide the tools and processing power required. Thus, it is widely recognized that analysis tools "near the data" are a key feature as we move into the era of large surveys and big data. IPAC has led the way in offering data analysis tools that run at the data center, allowing users to run the same software used by the mission. In this section, we give several specific examples designed for smaller data sets, which illustrate the approach.

5.1. IRAS analysis tools

In 1983, the Infrared Astronomical Satellite (IRAS) performed an unbiased, sensitive all sky survey at 12, 25, 60 and 100 microns, detecting about 350,000 infrared sources. The IRAS focal plane consisted of apertures leading to integrating cavities. As the satellite scanned, astronomical objects crossed the vertical array of sensors. The measurements were reported as Time Ordered Data.

IPAC produced the official data products for IRAS (image maps and source catalogs). To facilitate user-driven analysis, two data tools were developed. The first, "SCANPI", offers a factor of 2 to 5 gain in sensitivity over the IRAS Point Source Catalog by performing 1D scan averaging of raw survey data at specified arbitrary position. The second, "HIRES", returns IRAS survey images with higher resolution than the IRAS Sky Survey Atlas (ISSA).

HIRES uses the Maximum Correlation Method (MCM; [5]) to produce images with better than the nominal resolution. It is a powerful tool for studying morphology and separating confused sources. The MCM algorithm has also been applied to other applications. In particular, it became the basis of the coaddition pipeline for WISE, which is now available as a user tool through IRSA. HIRES itself has been adapted for use with Planck data. Planck is similar to IRAS, with a similar survey strategy and a focal plane filled with heterogeneous, not diffraction-limited detectors.

5.2. Herschel computer resources at IPAC

The NHSC provides large computer resources as well as long-term data storage. These resources have been used by scientists to reduce data that require more memory than available on local machines, to store data to be accessed by astronomers at different institutions, and to facilitate support by NSHC staff. The IPAC resources support the use of the Herschel Interactive Processing Environment (HIPE). The services are private, secure, virtual machines with a HIPE pre-installed for a requested period of time. Machine memory sizes range from 16 to 140 gigabytes, and can have multiple CPUs assigned to them. Data are secure and accessible by SFTP, rsync and webDAV. Rsync can be used to mirror data to other locations.

5.3. Montage astronomical image mosaic toolkit

The Montage Astronomical Image Mosaic toolkit [6] is a uniquely capable and scalable image processing toolkit that supports re-projection, transformation, subsetting, background matching, co-addition and mosaicking of FITS images, while preserving the calibration (photometric) and positional fidelity of the input images. It is available for download and as on-line service that creates custom image mosaics up to 1 degree on a side for 2MASS, SDSS, DSS and WISE. Montage has found wide applicability. Since 2005, there have been 120 citations to Montage in the peer-reviewed literature. These citations reference a wide range of data products and research areas, with particular applicability to Spitzer and WISE data. The components of Montage have been integrated into creation of data products that have had high impact in astronomy, such as several Spitzer Legacy and Exploration Science projects.

6. SCIENCE RESULTS

In this section, we highlight recent research projects enabled by IPAC's support for multi-mission big data investigations.

6.1. Proper motion searches for brown dwarfs

Combining all-sky data sets has proven invaluable in the hunt for cold and/or overlooked members of the immediate solar neighborhood. Kirkpatrick et al. ([7,8]) used WISE to identify 48,000 objects with significant motion over the six-month baseline provided by the first year of WISE data; Luhman [9] performed a similar search and found 762 new sources. Both surveys leveraged the longer time baseline provided by 2MASS, DSS1, and DSS2 to confirm the motion, or -- in the case of objects not seen in any of those earlier, shorter wavelength surveys -- to identify candidates that might be exceptionally cold. These surveys produced such discoveries as the third closest (sub)stellar system to the Sun (WISE 1049-5319AB at d=2.0 pc; [10]), the coldest known brown dwarf (WISE 0855-0714 at d=2.3 pc and Teff~250K; [8,9]), and many low-metallicity L dwarfs giving insight into the cooling rates of brown dwarfs at advanced ages [7,8]. Schneider et al. [11] used the longer time allowed by the reactivation to identify 20,568 motion objects; they likewise used 2MASS, DSS1, and DSS2 data as a crucial step in scrutinizing candidates. Gagne et al. [12] performed a 2MASS color cut to look for brown dwarfs that might be members of young, nearby moving groups; their candidate list of 98,970 motion candidates was the result of cross-matching to WISE, and in doing so they identified 228 possible new young, very low-mass brown dwarfs whose atmospheres can be studied as exoplanet proxies.

6.2. Superluminous spiral galaxies

Ogle et al. (2016, [13]) discovered a population of superluminous spiral galaxies, using data from NED joined from a wide range of surveys, including 2MASS, SDSS, GALEX, and WISE. The optical and near-IR surveys were essential to determine the galaxy masses, while WISE data gave star formation rates. Super spirals are the most optically luminous spiral galaxies found at redshift $z < 0.3$, selected to have luminosities more than 8 times the luminosity of the Milky Way galaxy and ultraviolet emission detected by GALEX. They are as bright as elliptical brightest cluster galaxies, but are found in a range of environments from isolation to cluster centers. Super spirals are giant, with diameters of 60-130 kpc, and also very massive, with stellar masses of 30-300 billion solar masses. They form stars at prodigious rates of 5-65 solar masses per year, primarily in their extended spiral disks. Local disk galaxies of this mass typically have lost their ability to form stars, through collisions and mergers with other galaxies or feedback by a supermassive black hole. Thus the discovery of these massive galaxies challenges theories of spiral galaxy formation and evolution.

7. REFERENCES

[1] Groom et al. 2014, ASPC, 485, 185
[2] Mainzer et al. 2011, ApJ, 731, 53
[3] Goldina et al. 2015, ASPC, 495, 137
[4] Taylor 2005, ASPC, 347, 29
[5] Aumann, et al., 1990, AJ, 99, 1674
[6] Jacob et al. 2009. Int J Comp and Engineering, 4, 1
[7] Kirkpatrick et al. 2014, ApJ, 783, 122
[8] Kirkpatrick et al. 2016, ApJS, submitted
[9] Luhman 2014, ApJ, 781, 4
[10] Luhman 2013, ApJL, 767, 1
[11] Schneider et al. 2016, ApJ, 817, 112
[12] Gagne et al. 2015, ApJ, 798, 73
[13] Ogle et al. 2016, ApJ, 817, 109

EUCLID: ORCHESTRATING THE SOFTWARE DEVELOPMENT AND THE SCIENTIFIC DATA PRODUCTION IN A MAP REDUCE PARADIGM

C. Dabin[1], M. Poncet[1], K. Noddle[2], M. Holliman[2], M. Melchior[3], O.R.Williams[7], A. Belikov[4], J. Koppenhoefer[5] , J. Hoar[6] on, behalf of Euclid Science Ground Segment System Team

[1] CNES - French Space Agency, Toulouse, France
[2] ROE - Royal Observatory, Edinburgh, United Kingdom
[3] FHNW - Fachhochschule Nordwestschweiz, Brugg, Switzerland
[4] Kapteyn Astronomical Institute University of Groningen, Groningen, The Netherlands
[5] MPE – Max-Planck-Institut für extraterrestrische Physik, Munch, Germany
[6] ESA – European Space Astronomy Centre, Madrid, Spain
[7]Donals Smits Center for Information Technology, University of Groningen, Groningen, The Netherlands

ABSTRACT

Euclid [1] is a high-precision survey mission developed in the frame of the Cosmic Vision Program of ESA in order to study the Dark Energy and the Dark Matter. Its Science Ground Segment (SGS) will have to deal with around 175 PetaBytes (PB) of data coming from Euclid satellite data, complex pipeline processing, external ground based observations or simulations, and with an output catalog containing the description of around 10 billion of objects with hundreds of attributes. Thus, the implementation of the SGS is a real challenge in terms of architecture and organization. This paper describes the Euclid Science Ground Segment Architecture, the organization of a collaborative software development and the data and processing distribution strategy in a distributed environment.

Index Terms— Euclid, Cosmic Vision, Science Ground Segment, Big Data, Architecture

1. EUCLID PROJECT

Euclid is a high-precision survey mission designed to answer fundamental questions on Dark Energy and Dark Matter. Euclid satellite will map the large-scale structure of the universe over the entire extragalactic sky out to a redshift of 2 (about 10 billion years ago). It will then cover the period over which dark energy accelerated the universe expansion. This will be accomplished thanks to a payload which consists of a 1.2m telescope, a visible imaging instrument (VIS) and a near-infrared instrument (NISP) with photometric and spectrometric capabilities. The satellite will be launched in 2020 on a Soyuz rocket from Kourou spaceport for a five years nominal mission. The spacecraft will be placed in a large second Sun-Earth Lagrange point (SEL2) halo orbit.

.

2. EUCLID GROUND SEGMENT

The ground segment is divided into the Operation Ground Segment (OGS) and the Science Ground Segment (SGS). The SGS [2] consists of a Science Operation Centre (SOC) and of several Science Data Centers (SDCs). The SDCs are in charge of the science data processing (Processing Functions) **software development**, and the **production** of science data products from Level1 (Raw Image) up to Level3 (Science ready data product : catalogs of galaxies, 2 point correlation functions, covariance matrices, mass maps, density maps, …), . There will be 8 SDCs over Europe and 1 in the US for which these responsibilities will be allocated or distributed.

During its lifetime, the SGS will have to handle an unprecedented volume of data for such a space mission: about 850 Gbit of compressed data will be downlinked every day, which means a total amount of at least 13 PetaBytes processed, transferred and stored scientific data for a given release. The SGS will have also to handle several external surveys for:

- the preparation of multiband ground based survey imaging data for combination with the Euclid data,
- the purpose of calibration and validation of Euclid measurements,
- the preparation of spectroscopic survey data sets for photometric redshift purposes.

It is expected to produce and store about a grand total of **175 PB** of data for the whole mission duration due to the three foreseen official data releases and multiple re processing and validation runs that will be required for the scientific validation of the data products. The Science Ground Segment will deliver and publish among other

science ready data products a catalog containing the description of more than **10 billions of sources with hundreds of columns.**

3. EUCLID SGS ARCHITECTURE

The Euclid SGS development is therefore a real challenge [3] in terms of software and infrastructure architecture design (software stack design, data storage, network, processing infrastructure and data/processing distribution strategy) and of organization. Thus, 9 Euclid SDCs and SOC will have to be federated, ensuring an optimized data storage and processing distribution and providing sufficient networking interconnection and bandwidth. In terms of organization, more than 14 countries will be involved in the project and hundreds of non-necessarily collocated people will have to work together either on scientific, software engineering or on IT aspects.
The reference architecture [4], currently proposed, for the SGS, will be based on:

- A **single mirrored metadata repository** which inventories, indexes and localizes the huge amount of distributed data: Euclid Archive System Data Processing System (EAS-DPS) along with interfaces and services for SGS subsystems [6]
- A **Monitoring & Control** (M&C) service allowing to monitor the status of the SGS as a whole or at SDC level,
- A **COmmon ORchestration System** (COORS) managing the distribution of the storage and the processing among the different SDCs (ensuring the best compromise between data availability and data transfers): "move the process not the data",
- A Euclid Archive System **Distributed Storage System** (EAS-DSS) providing a unified view on the SDCs distributed storage and managing the data transfers between SDCs,
- An **Infrastructure Abstraction Layer** (IAL) allowing the data processing software to run on any SDC independently of the underlying IT infrastructure, and simplifying the development of the processing software itself. It shows generic interfaces to the processing software and isolates it from the “plumbing” (e.g. it gathers input data, publish output data on behalf of the processing software).

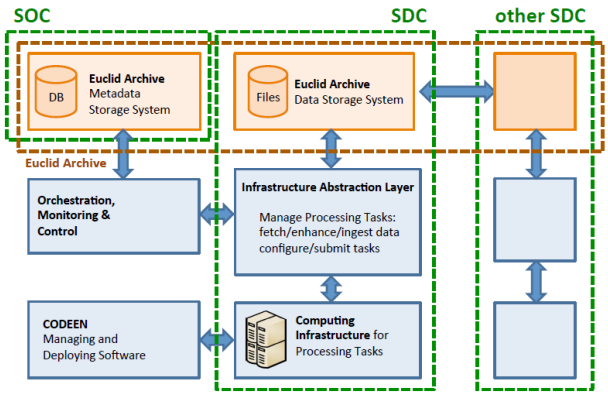

Figure 1 – SGS architecture components

This architecture concept has already been validated through "SGS Challenges", allowing namely to distribute and execute first simulation prototypes on any of the SDCs thanks to IAL and EAS prototypes. This development approach allows deploying working prototypes at early stages and is a great factor of motivation for the teams disseminated among different laboratories and Computing Centers around Europe.

4. EUCLID DATA PROCESSING SOFTWARE DEVELOPMENT APPROACH

The Pipeline Processing tasks are formally developed by many software development teams spread across Europe and US. To foster collaboration between these software development teams, to foster the homogeneity of the code: quality and coding standards, to foster the re use of common modules, to foster a clear non ambiguous definition of interfaces between these codes and to foster its portability and its deployment on various production infrastructures; a **common Euclid collaborative development platform (CODEEN)** has been set up. This platform is mainly based on **continuous integration principles** and associated tools. It automates unit tests, validation tests, pre integration, packaging and deployment on SDC’s. The Common Euclid development environment has been implemented into a Virtual Machine to make software developers autonomous.

A **Common Data Model** (CDM) shared by all software developers has been set up, This CDM relies on a common dictionary of data types from the most simple to the most complex. This CDM is the remit for the definition and implementation of the data products. The implementation of the EAS-DPS is based on this CDM.

5. EUCLID PROCESSING PIPELINE SOFTWARE STACK AND DESIGN PRINCIPLES

5.1. Software stack

The pipeline software stack is basically organized into a set of independent software layers that include applications and supporting libraries that are logically connected. The following figure depicts the different components of the architecture.

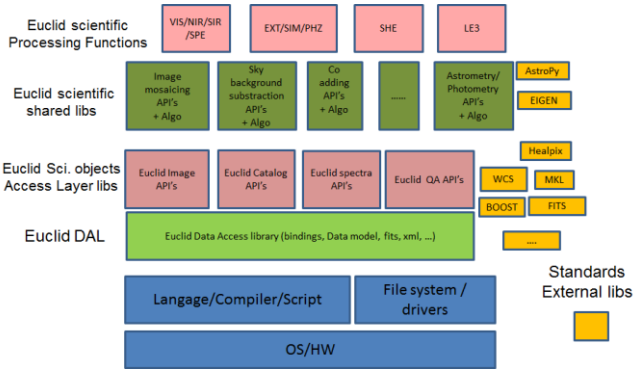

Figure 2 – SGS Software stack

5.2. Software Design principles

The following approaches are adopted:

• Co-locating the data and processing in the same place avoiding large data transfers

• Use of the infrastructure abstraction layer (IAL) that isolates and decouples the SDC infrastructure from the rest of the SGS.

This is practically implemented through the concept of dividing the work to be done by each SDC according to sky patches, not Processing Functions; that is it should be possible to run any Processing Function on any SDC computing facility and the data for each sky area should remain at the SDC responsible for that sky area.

Most of the processing functions software implementation inherits from existing codes that have been widely tested with various images and catalogs from ground or space based surveys or observatories. A specific attention and so a specific process has been set up to ensure the validation of these codes regarding the Euclid scientific requirements, the long term maintainability, the conformance with Euclid standards and the languages: C++ and Python that are supported by Euclid SGS.

6. DATA AND PROCESSING DISTRIBUTION PRINCIPLES: BIG DATA PARADIGM AT ARCHITECTURE LEVEL

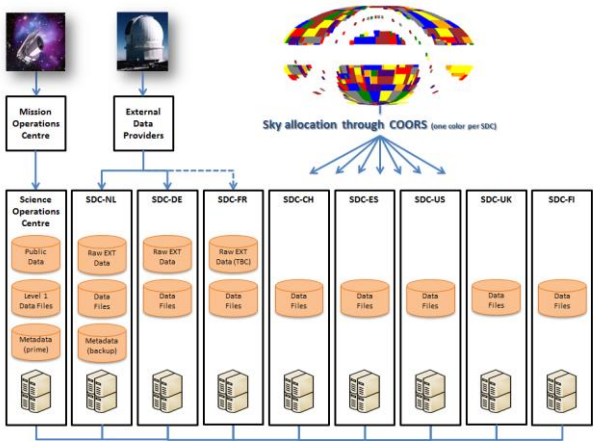

Figure 3 – SGS Observations Data distribution

A major part of the EUCLID data processing pipeline [5] (up to Level 3 processing) is parallelizable at the granularity of a given observation.

Hence, it has been possible to implement the "Big Data paradigm" at a multi-center scale:

6.1. The "Map" stage

First, each SDC is allocated areas of the survey (sky patches) for which it is responsible. The total area allocated to a SDC is defined as a function of the amount of resources available at that SDC since the size of the area of sky directly determines both the data volume output and the CPU resources necessary. A similar allocation can also applied for simulation data generation of large sky areas. There is some overlap between patches for processing reasons. In order to minimize data transfers, the system should maximize the amount of processing that can be done on that sky patch with the least additional data. Once an area of the sky is allocated to an SDC, the selected SDC will be responsible for running the entire processing tasks over that area of sky (except Level 3) and for maintaining the output data files. A second copy will be maintained at another SDC for robustness. To maintain some degree of flexibility a second SDC can be assigned a sky patch.

6.2. The "Reduce" stage

At Level-3 stage, the processing pipeline changes its shape. All the catalogued galaxies with the shear measurements, the spectra, the redshift and related attributes processed from single observations at level 2 are stored in the various SDC's. The relevant inputs needed for Level 3 are filtered out and retrieved to be processed by Level 3 processing functions in order to obtain high level scientific data products. This processing can be seen as a REDUCE step (as per Map/Reduce paradigm) at a multi-center level. The typology of LE3 processing algorithm is different from previous stages since they require to compute 'all the sky available' in a single run being able to produce some: 2 point correlation functions, 3 point correlation functions, E mode auto correlation, E and B mode cross correlation, covariance matrix,… for the various science domains that are: galaxy clustering, weak lensing, clusters of galaxies, time domain, milky way and nearby galaxies,... The challenge is now to cope with High Performance Computing codes on a reduced set of input data required by LE3 algorithms.

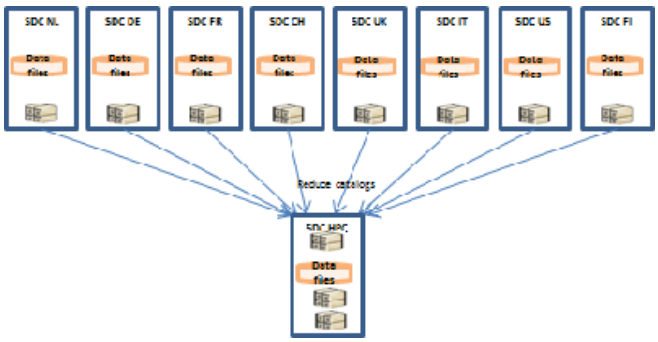

Figure 4 – SGS LE3 reduce stage

6.3. Technologies

The wide variety of technologies infrastructure supported by the various SDC's: protocol, storage, computing require considering virtualization technics and flexible deployment solutions. One of the solutions under test has been implemented by CERN and is CernVM ([7]). This is a baseline Virtual Software Appliance. The Appliance represents a complete, portable and easy to configure user environment for developing and running any scientific data analysis locally and on institutional and commercial clouds (OpenStack, Amazon EC2, Google Compute Engine), independently of Operating System software and hardware platform (Linux, Windows, MacOS).
The goal is to remove a need for the installation of the experiment software and to minimize the number of platforms (compiler-OS combinations) on which experiment software needs to be supported and tested.

7. CONCLUSION

The wide diversity of the algorithms to be implemented and deployed in this model of architecture forces a close cooperation between Infrastructure teams and software design architects to make sure that the model of architecture is correct. Each challenge, every 6 months, gives the opportunity to cross check and refine these assumptions.

8. REFERENCES

[1] Laureijs, R., et al., "Euclid: ESAs mission to map the geometry of the dark universe", Proc. SPIE 8442. doi:10.1117/12.926513

[2] F. Pasian & al, "The Science Ground Segment for the ESA Euclid mission", *SPIE*, 2012. doi: 10.1117/12.926026

[3] M. Poncet & al, "The challenge of Euclid Scientific Ground Segment Architecture", GSAW, Los Angeles, 2012. http://gsaw.org/wp-content/uploads/2013/07/2012s04poncet.pdf

[4] M. Poncet & al, "Euclid Scientific Ground Segment Organisation & Architecture", ESAW, Darmstadt, 2013.

[5] M. Sauvage "Data and Processing Flow Document" EUCL-0TS-ADD-8-002

[6] O. R Williams & al. "The Euclid Archive System : A data centric approach to big data", Big Data In Space, Tenerife, 2016

[7] CernVM software appliance. http://cernvm.cern.ch/portal/

BIG DATA, BIG CHALLENGES AND NEW PARADIGM FOR THE GAIA ARCHIVE

Christophe Arviset[(1)], Javier Durán[(1)], Juan González[(1)(4)], Raúl Gutiérrez[(1)], José Hernández[(2)], Uwe Lammers[(2)], Bruno Merin[(1)], Sara Nieto[(1)], William OMullane[(3)], Jesús Salgado[(1)], Juan Carlos Segovia[(1)]

[(1)]ESAC Science Data Centre, [(2)]Gaia Science Operations Centre, [(3)]Operations Development Division
ESA, European Space Astronomy Centre, 28692 Villanueva de la Cañada – Madrid, Spain
[(4)]ETSE Telecomunicación - Universidade de Vigo

ABSTRACT

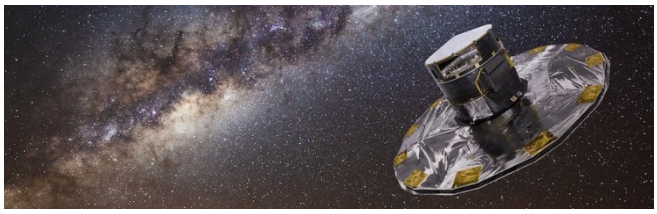

ESA's Gaia mission [1] will survey the sky for at least 5 years providing high accuracy astrometry, radial velocities and multi-colour photometry. The Data Analysis and Processing Consortium (DPAC) efforts will result in an astronomical catalogue with unprecedented accuracy and completeness of at least 1 billion (1E9) sources, and over 1PB of associated data products.

This brings big data challenges in storing, querying and distributing all the associated data and metadata, comparing them with other astronomical catalogues, enabling analysis, visualization, data mining and then sharing these results with other scientists. The amount of data involved forces a change of paradigm in dealing with astronomy archives. The usual usage of downloading the data to the users for her/him to work further on it needs towards evolve to a new way of working where the users' can send her/his code to the data, run it there on computing and storage services provided directly by the archive, where the data reside. The Gaia archive will provide an infrastructure to run added value interfaces and software on top of the Gaia data.

To answer to all these new technical challenges, the Gaia Archive has directly adopted VO standards [2] (TAP, ADQL, UWS, VOSpace, SAMP amongst others) to build the archive data management infrastructure, giving birth to one of the first "VO-built in" astronomy archive. Embracing VO technology from the start also ensures full interoperability of the Gaia archive to other VO compliant data archives and applications.

Index Terms— Gaia, Virtual Observatory (VO), astronomical archives, big data.

1. INTRODUCTION

The ESAC Science Data Centre (ESDC), located at ESAC, Madrid, Spain, is responsible for the design, implementation, operations and long term preservation of the science archives for all ESA space science astronomy and solar system missions.

ESA's Gaia cornerstone mission [1] was launched on 19th December 2013 and since then, has been doing a stereoscopic census of the galaxy. Gaia will produce a catalogue of 1-2 billion identified sources with unprecedented accuracy. The final delivery will also include all the single epoch CCD transit data that was used in the catalogue computation, reaching an estimate of around 1PB at the end of the mission in 2022 (Figure 1). The massive data processing challenge is being efficiently tackled by all DPAC partners and is better described in other paper [3].

Nominal Mission Data up to January 2016

Type of Data	Amount		
Science Telemetry	28 TB		
Astrometry transits	35.2 x 10^9	352.3 x 10^9	measurements
Photometry transits	35.2 x 10^9	73.6 x 10^9	low resolution spectra
Spectroscopy transits	2.3 x 10^9	6.9 x 10^9	high resolution spectra
Main database	90 TB	(700 TB expected after 5 years)	

with associated products, up to 1 PB of Gaia data by mission end

Figure 1: Gaia data : events and volume

The first Gaia catalogue will be publicly released to the scientific community around summer 2016, but the development of the Gaia archive has long started and an internal version is already available for the Gaia Consortium.

2. IS GAIA A "BIG DATA" PROJECT ?

The term Big Data can be used and understood differently by people, but today industry widely refers to the five "Vs"

(Volume, Velocity, Variety, Veracity and Value) when speaking about Big Data.

From the Volume perspective, with "only" 1 PB of data produced by the end of its lifetime, Gaia could hardly be considered as a big data mission, as many other astronomy projects (in particular ground based telescope) will produce way more data volume in shorter timescale.

Nonetheless, the billions of CCD transits, measurements and spectra (Figure 1) that will result into massive sources catalogues (Figure 2) for sure makes Gaia a big data challenge, from the data Velocity and Variety points of view.

Simulations Catalogues		
• GUMS		
• Milky Way:	$2x10^9$	rows
• Large Magellanic Cloud:	$7.5x10^6$	rows
• Small Magellanic Cloud:	$1.2x10^6$	rows
• Galaxies:	$38x10^6$	rows
• Quasars:	10^6	rows
• GOG	$1.8x10^9$	rows
External Catalogues		
• IGSL (Initial Gaia Source List)	$1.2x10^9$	rows
• 2MASS	$9.4x10^8$	rows
• Tycho2	$2.5x10^6$	rows
• UCAC4	$1.1x10^8$	rows
Gaia		
• TGAS (private validation team area)	$2.2x10^6$	rows
• **Foreseen, first Gaia catalogue**	**$>2x10^9$**	**rows**

Figure 2: Catalogues currently in the Gaia Archive

Ensuring the Veracity of the Gaia data represents one of the main big data challenge of the Gaia data processing [3]. And to finish, by performing astrometry, photometry and spectroscopy of about one billion objects in our Milky Way galaxy and beyond, the extent and content of the Gaia Catalogues will enable major progress to be made in many fields of galactic and stellar astronomy, hence its Value definitely places Gaia as a major Big Data project in astronomy.

3. STANDARD ARCHIVE ARCHITECTURE

Standard ESA space science archives architecture (Figure 3) is based on the OAIS architecture [3]. Users and scientists are used to interact with data or catalogues, either through a browser user interface or scriptable command line interface. They download the data and the full catalogue, usually via FTP, to their local disk and perform their science analysis on their local computer.

This working model works fine for small amount of data, but becomes difficult as data volume grows. To reduce the data transfer burden, the users try to select region of the sky and download only part of the catalogue, and then combine it with their own datasets or catalogues already stored on their disk. This can be describe as "move the data to the computing facility".

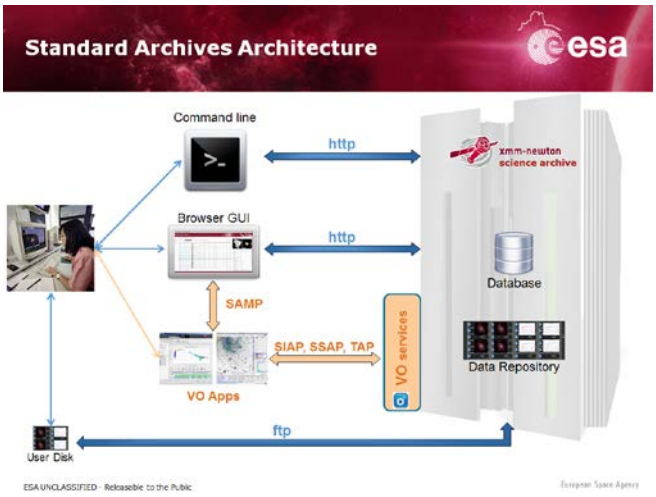

Figure 3: Standard ESA Archive Architecture

The Virtual Observatory (VO) provides an unified framework which enables transparent access to astronomical science data holdings coming from various different archives. The very same command can be sent to archives located in different locations (and using their own internal storage and database systems) and present results in a consistent way to the end user or applications. Early developed interoperability VO protocols (eg ConeSearch, Simple Image and Simple Spectra Access) greatly facilitates access to multiple datasets, but still assumes that data are finally being downloaded to the user's computer.

Usually, science archives implements a "VO layer" on top of the existing archive infrastructure, so all the data holdings can be accessible through these VO protocols. This ensures the interoperability of the archives with other VO compliant archives and applications.

4. NEW ARCHIVE PARADIGM FOR GAIA

With the avalanche of data in astronomy, the archive model previously described reaches its limit and a new paradigm needs to be established, the so called "move the code to the data".

For Gaia, the amount of data and meta data is so big that special computing infrastructure is required to efficiently handle Gaia data. For example, a query of a cone search on the ~1 billion Gaia catalogue might return 10 million sources. Another typical use case is to upload a table with sources and to cross match these with the Gaia catalogue. This operation is made possible with the VO Table Access Protocol (TAP), coupled with ADQL (Astronomical Data

Query Language, SQL with specialized astronomical searches). TAP also supports asynchronous query, as such a cross match can require too much time to be perform interactively. Universal Worker Service (UWS) enters in action to manage these asynchronous jobs. This back-end infrastructure serves all the front-end interfaces available at the Gaia archive.

Special emphasis has been put on the database design (based on PostgreSQL and pgSphere add-on which provides spherical data types, functions, and operators for PostgreSQL) and associated indexing. Furthermore, the Gaia archive is hosted on a powerful server and the most popular catalogues are stored on PCIe SSDs disks to ensure good performance of these crossmatch functions. Some examples of time required for some of the complex crossmatches are given in Figure 4.

Catalogue 1	Catalogue 2	Radius (arcsec)	# results	Time
Tycho2 2.5×10^6 rows	2MASS PSC 4.7×10^8 rows	1″	2,495,304	49s
Tycho2 2.5×10^6 rows	2MASS PSC 4.7×10^8 rows	5″	2,614,163	116s
Tycho2 2.5×10^6 rows	IGSL 1.2×10^9 rows	1″	2,600,542	46s
Tycho2 2.5×10^6 rows	IGSL 1.2×10^9 rows	5″	2,829,401	55s

Tycho2 vs. IGSL crossmatches are even faster than the ones with 2MASS as IGSL is located in the fastest local storage (PCIe), even when IGSL (similar to the final Gaia catalogue) is around 3 times bigger than 2MASS.

Figure 4: Crossmatch Performances

End user can then interact with the Gaia through a standard Graphical User Interface from any standard web browser. This offers a full ADQL query interface, with example to help the user familiarize with the Gaia catalogue content and structure.

In addition, it is expected that many of the users will interface with Gaia data directly through scriptable interface. All operations available from the Gaia archive GUI (ie TAP/ADQL queries) can also be included directly in user's scripts. Various examples of such scripts are provided in the most commonly used programming language in astronomy (Python, C, Java).

VO applications (eg Topcat, VOSpec) can access directly the Gaia data through the corresponding VO protocol, TAP for table, Simple Access Protocol for spectra, without the need to develop a "VO layer" as seen in the standard ESA archive architecture.

When the user retrieves big volume of Gaia data, she/he would be able to download it to her/his local disk via FTP, but it will probably be more efficient to leave it on the archive disk itself to avoid the burden of the network transfer. This can be done with VOSpace, virtual disk accessible by VO data access protocols. By keeping the data on her/his VOSpace, the user can continue to interact with it and as well share it with any other Gaia archive user, to facilitate scientific collaboration.

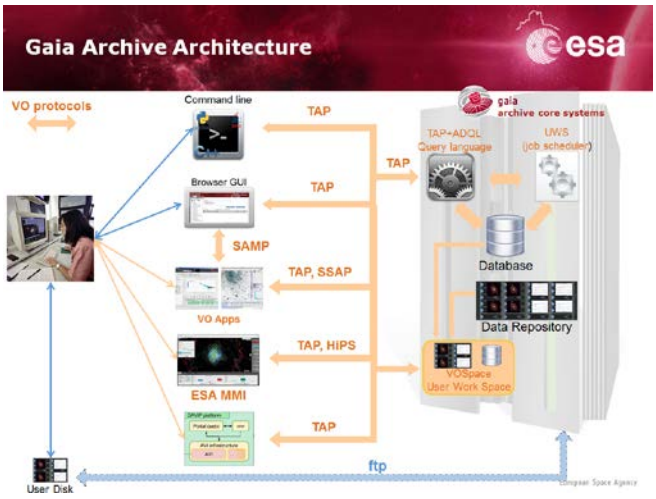

Figure 5: Gaia Archive VO built-in Architecture

Similar mechanism is provided for metadata storage. If the user would download the results of a complex search (ie returning millions of entries) on her/his computer, she/he will probably needs to ingest these into a local database to continue to work on them. It appears then more efficient to provide to the user some space directly within the Gaia archive database to store the results of her/his searches. As such, the user "database" becomes VO compliant and shareable as well with any other Gaia archive user.

Overall, the Gaia archive architecture depicted in Figure 5 has been built around these VO protocols and has become the first "VO built-in" archive. By design, the Gaia archive also becomes immediately interoperable with any other VO compliant archive and application.

5. GAIA ADDED VALUE INTERFACES

The user workspace can be brought one step further to enable the user to also run her/his own code directly on the Gaia archive through so called "Gaia Added Value Interfaces" (Gaia AVIs, Figure 6). Four AVI demonstrators are currently being developed for transient alerts, advanced visualization, spectral classification and temporal analysis. These AVIs will run using containers (Docker) and will make use of the Gaia Archive VO built-in protocols (TAP, VOSpace). A Gaia AVI Portal will be created and users will be able deposit their own code and can run it on their data located into their user workspace (VOSpace and database). AVI templates will be provided to help the user to develop their own AVIs.

The Gaia AVI project is currently being developed as a proof of concept project to be delivered in 2017 and could become the new framework for collaborative "Archive 2.0".

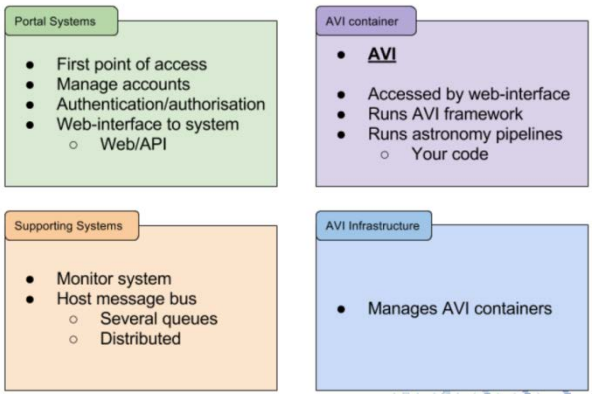

Figure 6: Gaia Archive AVI Building Blocks

6. BIG DATA VISUALIZATION

Before being able to search for Gaia data, it might be really helpful to provide visualization of the Gaia data in various ways, such as density maps, 1D histograms (Figure 7) or again interactive visualization through VO application (Aladin Lite), integrated into the archive GUI, through another VO protocol SAMP (Simple Application Messaging Protocol). The production of such graphs requires the use of big data reduction techniques, such as Map / Reduce. With 10 parallel threads on a powerful machine with big RAM (1TB), fast disks (PCIe SSDs with fast random IO) and efficient database (PostgreSQL), the production of the density maps for the GUMS simulated catalogue (2.14 billion rows) took less than 2 minutes and the ones for the IGSL (1.22 billion rows) as little as 65 seconds.

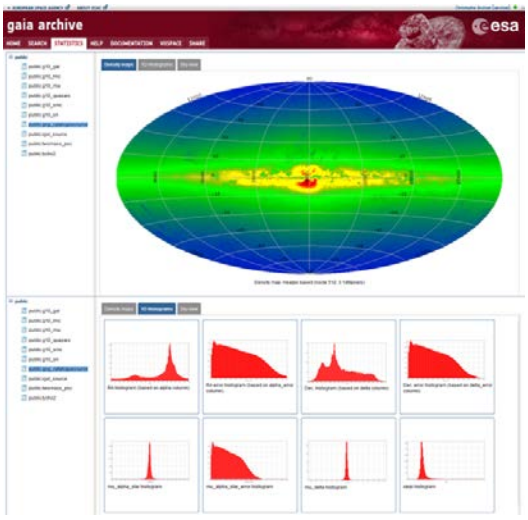

Figure 7: Gaia data density maps and 1D histograms

Another way to visualize Gaia data will be through the recently released science-driven discovery portal for all the ESA Astronomy Missions called ESA Sky that allow users to explore the multi-wavelength sky and to seamlessly retrieve science-ready data in all ESA Astronomy mission archives from a web application without prior-knowledge of any of the missions. Amongst other things, the system offers progressive multi-resolution all-sky projections of full mission datasets using a new generation of HiPS (Hierarchical Progressive Survey) files. HiPS is based on the HEALPix sky tessellation and is essentially a mapping of survey data at various spatial resolutions into a collection of HEALPix tiles. It is particular adapted to big data visualization as it allows a dedicated client/browser tool to access and display a survey progressively, based on the principle that "the more you zoom in on a particular area the more details show up".

7. CONCLUSION

Gaia, ESA's cornerstone mission currently in operations represents one of the major big data challenge in astronomy to date. A totally new archive architecture (both hardware and software) has been developed to tackle this challenge. It results into one of the first VO built-in science archive, paving the way towards flexible, open and interoperable archive services. The user will work directly with the data in the archive through dedicated user workspace, without the need to transfer it to her/his location, She/he will be able to become an actor of the archive with the possibility to bring her/his code to the data and share it with other archive users. This new "Archive 2.0" concept will be the mean to fully exploit the science legacy of the Gaia mission.

8. ACKNOWLEDGEMENTS

The Gaia Archive and its content are the result of the work of hundreds of people so the authors would like to particularly thank the Gaia Science Operations Centre at ESAC, the Data Processing and Analysis Consortium, the Gaia Added Value Interfaces teams and the ESAC Science Data Centre.

9. REFERENCES

[1] The Promises of Gaia, Prusti, T., 2012AN....333..453P http://ads.nao.ac.jp/abs/2012AN....333..453P
[2] IVOA standards, http://www.ivoa.net/documents/
[3] Gaia downlink data processing, Siddiqui, H. et al, 2014SPIE.9149E..2ES
[4] OAIS, Reference Model for an Open Archival Information System.

NEEDS FOR THE QUEST OF EARTH ANALOGS: THE PLATO SPACE-BASED MISSION

Deleuil, M., Rauer, H., Pollacco, D., Gizon, L., Walton, N.A., Piotto, G., and the PLATO Team

(1) Aix-Marseille University, LAM, Marseille, France; (2) DLR - Berlin, Germany; (3) Warwick University, UK; (4) MPS, Germany; (5) IoA, Cambridge, UK; (6) Universita di Padova, Italy

ABSTRACT

Space-based missions such as CoRoT and Kepler have demonstrated the power ultra-high precision photometry for planets hunting. However the detection of planetary transits requires the continuous monitoring of tens of thousands of stars over years in order to overcome the low transit probability and detect Earth analogs and thus the handling and the analysis of tens of thousands of light curves at high temporal cadence. While the detection is the first and unavoidable step, the next generation of instruments like PLATO, selected by ESA as the third M-class mission, intend to go one step further. The objective of PLATO is not only to ensure the discovery of large statistically meaningful samples of new planets, but also to fully characterise them. PLAT will enable an understanding of exo-planet diversity through a good estimate of their composition, and decipher how they formed and evolved.

Achieving such a goal requires to combine light curves analysis to information from others facilities. They are indeed needed for the careful characterization of the stellar population to be observed, for the identification of false positives, and for the final determination of planets parameters. Complementary observations, involving various technics such as high precision radial velocity measurements, high contrast imaging or spectropolarimetry, are thus full part of the PLATO mission, that will also benefit the GAIA catalog.

This paper presents the scientific objectives and the design of the mission. It reviews the needs of PLATO in terms of data analysis and ground-based resources, from the mere detection of transits to planetary systems fully characterized.

Index Terms— Planetary systems, transit survey

1. NEEDS FOR A NEW GENERATION OF SPACE-BASED MISSION

Space-based missions, CoRoT and Kepler, have provided us thousands of planet candidates and some hundreds of confirmed exoplanets through the analysis of their ultrahigh precision light curves. Among forefront results such as transiting multi planet systems [1] or circumbinary planets [2], unexpected populations in the small planet regime have been discovered: the Super-Earths (1.25 to 2 $R_\oplus$) [3], and the mini-Neptunes (2 to 4 $R_\oplus$) [4]. These small planets, which have no analog in the Solar System, are nearly ubiquitous and present a surprising diversity in their bulk properties that is still to be better understood. To that purpose, it is necessary to confront a large sample of them, with their parameters and properties accurately determined, to models that form planets and population synthesis.

Achieving such a goal requires to combine light curves analysis with information obtained from others methods for different reasons: 1) Number of stellar configurations: eclipsing binaries, triple systems, background eclipsing binaries.. can have a photometric signature similar to a planet one and should be thus carefully identified through others technics: spectroscopy, imaging... A few hundreds of planets have been secured this way, and a few thousands of Kepler candidates with a large range of orbital periods and sizes are still awaiting for further confirmation. 2) Transits being a geometric effect, they allow to measure the planet's radius. Radial velocity is a gravitational induced effect and is thus proportional to the mass of the planet. Measuring both the radius and the mass of exoplanets reveals their nature and give first constraints to their internal structure and formation mechanisms. 3) Any planet's parameter being measured relatively to the host star's one, the uncertainties on our knowledge of the host star and its fundamental parameters: radius, mass, and age, directly impact those on the planets' and prevent setting precise constraints to planet formation and evolution theories.

The issues we are facing today come from the faint magnitude of the host stars, as previous space-based missions targeted stars in the V-magnitude range 14 to 16 so that to survey a large number of stars and increase the chance of detection. On the opposite, radial velocities and asteroseismology require photons in order to achieve the high precision needed to measure small planet masses or the mode frequencies of the host star. The poor overlap between the domain of performances of these methods which give access to different planetary system parameters prevents their complete and deep characterization and the precise determination of the physical properties of the planets. It is even more critical in the small size domain where the amount of observing time that would be necessary for such confirmation/characterization process is by far too prohibitive with current facilities. In addition, the magnitude issue also prevents from any further atmosphere

characterization of these planets.

PLATO 2.0, selected by ESA as the third M-class mission [5], is designed to overcome these limitations. It is designed to optimize the number of stars to observe and their brightness so that to combine the best precision of all technics involved in the planet characterization: transit, radial velocity and asteroseismology of the host-star. The three technics are indeed fully complementary and their combination is the only way to determine accurate radii and masses especially for planets in the super-Earth to Neptune mass range and assess their properties.

To that purpose, the instrument will collect long, uninterrupted, ultrahigh precision photometric light-curves of thousands of relatively bright stars (mag V $\leq$ 11), observed for durations ranging from a few months up to 3 years. PLATO data will be used both for detecting exoplanets through the transit method, and for characterizing their host star via asteroseismology. The later will allow a seismic determination of stellar masses, radii and ages of stars with an accuracy of a few percents, by fitting stellar model to the frequencies of oscillation. The deep analysis of these light curves will be done in combination with various datasets from ground-based facilities so that to rule out false positives and get the complete set of planet's physical and orbital parameters. From this strategy, we expect PLATO 2.0 to allow us to distinguish between an almost coreless planet and a planet interior with a large iron core.

2. MISSION PROFILE

The instrument will consist of 34 telescopes mounted on a single plateform. This multi-telescope design enables to cover a very wide field of view, dictated by the paucity of bright stars, and the collecting area required to achieve a high sensitivity. The camera, based on a fully dioptric design with 6 lenses, will operate in white light. 32 will have a read-out cadence of the CCD of 25 sec. Two additional cameras will have a higher read-out cadence (2.5 s) and will be used for observing the brightest stars (mV $\simeq 4-8$), and as fine guide sensor for the spacecraft attitude control system.

PLATO is scheduled to be launched from Kourou by a Soyuz 2-1b rocket. It will perform its scientific observations in an operational orbit around the Earth-Sun Lagrange Point 2 (L2). The nominal mission lifetime is 6.5 years, with consumables for 2 additional years of operations. The observing duty cycle will be at least 95%.

While high-number of detections needs wide field, planets at large orbits need long duration pointing. The PLATO nominal science operation will be 6.5 years that will be divided up in two long pointings of 2 - 3 years and a number of shorter duration pointings of 2 to 5 months. PLATO shall be able to observe between 25 and 50% of the sky to allow for statistical studies of planetary formation under various conditions.

PLATO will survey about 85,000 stars in the range 4-11 mag suitable for super-Earths detection and full planet and host star characterization. Stars in the range 11-13 mag will be too faint for radial velocity measurements and precise asteroseismology analysis at the level of precision needed for the characterization of small size planet. They will however provide significant information for the statistical investigation of planet frequencies and of planetary system architectures. We expect in total about1 Million of stars to be observed in the complete range of magnitudes over the lifetime of the mission.

The final products of PLATO will be a catalogue of extrasolar planetary systems with their complete characteristics derived using follow-up observations and/or transit timing variations: orbital parameters, planet size, mass, density (average composition), insulation.. together with their host star parameters: radius, mass, age, effective temperature, rotation, and others physical or compositional properties.

In addition, a second catalog of planet candidates with their basic parameters and the associated planetary probability, still awaiting for a final confirmation by the community at large, will be issued as part of the PLATO legacy. The seismic analysis will not be limited to host stars only. It will be performed on a large sample of bright stars across the HR diagram so that to improve stellar models.

Generate these various products will require not only careful analyses of PLATO light curves but also to combine data from ground-based observations. This will be the task of the PLATO Data Center.

3. PLATO DATA PRODUCTS

The PLATO Data Center (PDC) is a key element of the mission. Under the responsibility of the PLATO Mission Consortium, it will support the Science Operations Centre. It will process the data stream of each star observed in a given field together with the associated target parameters from the Input catalog and ground-based observation products (see sect. 3). Before the launch it will be in charge of providing the PLATO Input Catalogue to the SOC for the scientific mission planning and later its regular updates necessary for the preparation of each new pointing. Once the instrument will be in operation, it will provide to the scientific community the PLATO data, that is validated and calibrated light curves and centroid curves, corrected for the main instrumental effects, and ready for the scientific analyses. In parallel, it will carry out the analysis of the light curves for the transit detection and the stellar analyses. PLATO will downlink some 430 Gbit/day using the ESA K-band Deep Space Network. The PLATO data processing is highly iterative, hence the end of mission data volumes will be in the mult-Petabyte regime.

Starting from the validated light curves, two data analysis pipelines will be run mostly in parallel, while cross-talking to each other and exchanging results. One pipeline will be dedicated to the transit search and the derivation of the plan-

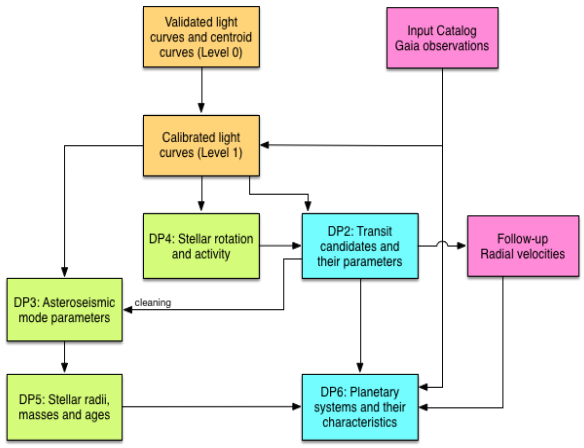

Fig. 1. Summary of the PLATO data flow and products

ets parameters. The second one will carried out the seismic analysis and will provide the inferred host star's fundamental parameters: mass, radius and age with some additional estimates such as the star's rotation or its photometric properties. All the information obtained on a given planetary system will be then gathered for a final homogeneous processing that will end up in the complete set of parameters that will be provided in the final catalogue of extra-solar planetary systems, issued at the end of the mission. This complexe processing is summarized on Fig. 1.

The Exoplanet pipeline (Exoplanet Analysis System) will perform the transit feature detection on all the light curves collected for a given field. Several tools have been already developed for the automated detection of planetary transits (see [6]). They are based on the main characteristics of the planet transit: a short, small, and periodic decrease of the stellar flux. Both detection completion and false alarm rate are method dependent. As a consequence, the use of complementary filtering and precise detrending of the light curves prior to the detection will be implemented in the Exoplanet Analysis System.

Once the light curves are filtered and normalized, transit parameters will be derived from a detailed modeling. Classical technics are based on the transit model by [7] and applies the Amoeba optimization procedure [8]. In this modeling special care should be brought on: 1) the identification and correction for possible nearby contaminant sources that may affect the transit depth, 2) The limb-darkening coefficients [9] which directly depends on the central star's properties. The two might directly impact the inferred planet's radius and will benefit complementary information from existing catalogs

The detection will trigger the vetting for the presence of false alarms.The planetary likelihood will be assessed taking into account different criteria: the shape of the transit, the possible existence of out-of-transit variations, the detection and depth of a secondary eclipse, the depth and the duration of the transit, the stars spectral type, the centroid measurements and the local neighborhood of the targets, where most of the false-positives are. In this process again all information available on the target itself but also on its neighborhood will be taken into account. It will end up with a prioritized list of planet candidates and their associated parameters and errors. The hierarchy in the list will depend on the planetary likelihood of the signal but also on its scientific merit: small size and long orbital period.

4. NEEDS FOR GROUND-BASED SUPPORT

In the preparation phase of the instrument, information from existing catalogs will be used to prepare the fields selection then the careful selection of targets that will be monitored. The optimal selection of both stellar fields and targets to be monitored relies on a good estimate of stars parameters: their magnitudes, parallaxes, effective temperature, but also their rotation or activity level. In addition, to guarantee the optimal photometry of the selected targets, their neighborhood should be also known within the PLATO's PSF down to about 5 magnitudes difference. This information will be obtained thanks to the analysis of accurate all-sky catalogs such as Hipparcos, Tycho-2, 2MASS, and the forthcoming GAIA catalog.

Once the instrument will be in operation the detection of planet candidates and their rank in the priority list will trigger ground-based observations, at least for the most promising one. The goals are twofold: (i) screen out the candidates list from the remaining false positives, (ii) measure accurate planetary masses through radial velocity measurements. Follow-up observations will be carried out by the PLATO community on a limited number and well chosen planet candidates so that to achieve the complete characterization of a significant number of small size planets in the habitable zone of their host star in a reasonable amount of telescope time. This optimal use of the telescope time relies on both the ranking of the planet candidates and a well define sequence of tests to be carried out using different facilities:

– Low-precision spectroscopy to identify possible spectroscopic binaries;

– High-resolution imaging to identify within the large PLATO PSF much fainter companion (physically bounded or not to the target);

– On and off transit photometry to secure the target as being the source of the observed transits;

– High-resolution spectroscopy to measure accurately the planet's mass;

– Rossiter-McLaughlin observations for the smallest objects in the Habitable Zone of heir host star.

The characterization of the most difficult case of an Earth-size planet orbiting a solar-type star at 1 AU requiring nearly two hundred of radial velocity measurements, the sequence resumed in Fig. 2 ensures to concentrate the follow-up effort on the best targets with the highest scientific merit.

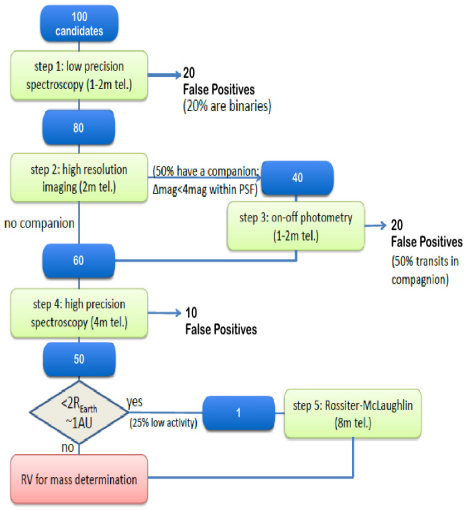

Fig. 2. Sequence of filtering observations. Dark blue boxes give the number of candidates entering the next step

The results of the follow-up observations will be gathered in the PDC. They will be used in the Exoplanet Analysis System for the final estimate of the complete planet's set of parameters. This processing will be done by fitting together, in a self-consistent manner, all the available datasets and not only the light curve. This combined modeling will allow to derive accurate planet parameters with rigorous confidence intervals.

Besides the large sample of candidates that will be securely identified as true planets and will constitute the PLATO catalog, an even larger sample of planet candidates will lack the filtering by complementary observations and will thus keep their status of candidate. They will however go through a complex process of planet validation [10]. Taking into account all the information available on the target and its neighborhood, the probability of each possible scenarios, planet and false positives, will be computed. The candidate is considered as validated if the posterior probability of the planetary hypothesis is significantly larger than the sum of probability of all false positive scenarios. The objective will be to allocate a planetary posterior probability associated to each candidate that could be taken into account when assessing their statistical properties.

5. CONCLUSION

On the roadmap toward Earth analogs and the understanding of the origins and evolutions of planets, characterization of small size exoplanets within the solar neighbourhood is now the next step. Measuring their masses, radii and densities with the highest accuracy is required so that to identify those suitable for studying their atmospheric composition and properties. This requirement lead to the design of PLATO. Continuous and ultrahigh precision photometry on bright stars, achievable from space only, will allow to measure planet transits and stars oscillations. The analysis in a self-consistent way of this unique data combined to datasets from well-identified ground-based instruments, will give us the complete and accurate set of planet's physical and orbital parameters, revaeling the properties of small planets over a large range of orbital period.

6. REFERENCES

[1] J. J. Lissauer, G. W. Marcy, et al., "Almost All of Kepler's Multiple-planet Candidates Are Planets," *Astrophys. Journ.*, vol. 750, pp. 112, May 2012.

[2] L. R. Doyle, J. A. Carter, et al., "Kepler-16: A Transiting Circumbinary Planet," *Science*, vol. 333, pp. 1602, Sept. 2011.

[3] A. Léger, D. Rouan, et al., "CoRoT-7b: the first super-Earth with measured radius," *Astron. & Astrophys.*, vol. 506, pp. 287–302, Oct. 2009.

[4] J. J. Lissauer, D. Ragozzine, et al., "Architecture and Dynamics of Kepler's Candidate Multiple Transiting Planet Systems," *Astrophys. Journ. Sup.*, vol. 197, pp. 8, Nov. 2011.

[5] H. Rauer, C. Catala, et al., "The PLATO 2.0 mission," *Experimental Astronomy*, vol. 38, pp. 249–330, Nov. 2014.

[6] J. Cabrera, S. Csizmadia, A. Erikson, H. Rauer, and S. Kirste, "A study of the performance of the transit detection tool DST in space-based surveys. Application of the CoRoT pipeline to Kepler data," *Astron. & Astrophys.*, vol. 548, pp. A44, Dec. 2012.

[7] K. Mandel and E. Agol, "Analytic Light Curves for Planetary Transit Searches," *Astrophys. Journ. Letters*, vol. 580, pp. L171–L175, Dec. 2002.

[8] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical recipes in C++ : the art of scientific computing*, 2002.

[9] S. Csizmadia, T. Pasternacki, C. Dreyer, J. Cabrera, A. Erikson, and H. Rauer, "The effect of stellar limb darkening values on the accuracy of the planet radii derived from photometric transit observations," *Astron. & Astrophys.*, vol. 549, pp. A9, Jan. 2013.

[10] R. F. Díaz, J. M. Almenara, et al., "PASTIS: Bayesian extrasolar planet validation - I. General framework, models, and performance," *Monthly Notices of the RAS*, vol. 441, pp. 983–1004, June 2014.

NOAA'S BIG DATA PARTNERSHIP: LEVERAGING THE VALUE OF DATA WITH INDUSTRY

Edward J. Kearns
Amy Gaskins
Brian Eiler
Maia Hansen
Alan Steremberg*
Chelsea Mastrapa

National Oceanic and Atmospheric Administration (NOAA)
United States Department of Commerce
Office of the Chief Information Officer
1315 East West Highway SSMC3
Silver Spring, MD 20910 USA
Ed.Kearns@noaa.gov, Amy.Gaskins@noaa.gov, Brian.Eiler@noaa.gov, Maia.Hansen@noaa.gov, Alan.Steremberg@gmail.com, Chelsea.Mastrapa@noaa.gov

ABSTRACT

NOAA's Big Data Partnership with US industry was announced in April 2015. Under the terms of a multi-year cooperative research agreement, cloud service providers Amazon, Google, IBM, Microsoft, and the Open Cloud Consortium have begun to store and disseminate NOAA data and information products. These data were selected based on prospective business opportunities with other industry and business entities. The Big Data Partnership is testing whether a self-sustaining business ecosystem can leverage the value inherent in NOAA's data in order to underwrite the costs of cloud storage and dissemination, as well as provide wider application of NOAA's data and new business opportunities for industry.

***Index Terms*—** Big Data, Cloud, Access, NOAA[1]

1. INTRODUCTION

At BiDS '14, Kearns et al. [1] introduced the emerging relationship between the National Oceanic and Atmospheric Administration (NOAA) and the cloud infrastructure business community. On April 21, 2015 the US Secretary of Commerce, Penny Pritzerk, announced that under the auspices of a cooperative research agreement, five major cloud infrastructure providers are signing agreements. These Big Data Partnership (BDP) agreements provided a mechanism by which NOAA could provide its data, the industry signees (hereafter "Collaborators") could provide their infrastructure, and industry, partners, and researchers could receive access to NOAA's data for little or no cost (Collaborators could only charge fees to recover costs associated with data acquisition, if any).

Following the signing of the Big Data Partnership agreements with NOAA's five collaborators, Amazon, Google, IBM, Microsoft, and the Open Commons Consortium, an engagement strategy was pursued by NOAA, which allowed the collaborators to select those NOAA datasets of greatest interest and value to them. Rather than have NOAA decide which datasets to contribute to the BDP, the Collaborators would decide which datasets made "business sense" to them and their "associated business partners" -typically companies and entities with specific technical markets and abilities.

Discussions took place among NOAA and the Collaborators to present candidate datasets of interest. NOAA data managers and subject matter experts discussed the proposed datasets and held follow-up meetings with each Collaborator and their associated business partners to review details. Typically, attempts were made to identify those datasets which had a demonstrable business use case, were currently unavailable or very difficult to access, had no restrictions on distribution, and had a usage pattern that enabled novel

*Alan Steremberg
660 4th Street #554
San Francisco, CA 94107

applications. Small individual datasets that were already easily accessed on-line and widely used, were not typically initially identified by the BDP collaborators, though in combination with other data were recognized as having significant "Big Data" value.

For those datasets of interest, the NOAA data experts, BDP Collaborators, and their associates engaged in technical interchange discussions about how to use the data, the meaning of the data, and how these data might be best copied from NOAA to their cloud infrastructure. NOAA managers needed to assess the NOAA resources required to deliver these data, while maintaining current operations and level of service to existing NOAA customers. Milestones, deliverables, and a timeline were developed around "best effort" service from NOAA, though the collaborators could, under the terms of the cooperative research agreement, provide additional resources to enable NOAA to expand and/or accelerate its efforts.

To better understand how these relationships will work, we will describe the first remotely-sensed dataset introduced to the BDP, even though it is not a space-based remote sensing dataset. We will follow with descriptions of approaches to be used for other datasets, including NOAA's satellite data, under discussion and development for the BDP.

2. EXAMPLE: NOAA WEATHER RADAR DATA

The NOAA Big Data Partnership copied one of NOAA's flagship data products to the industrial cloud which has the potential to enhance its use, promote new business opportunities, and increas its value for the US economy. The Next-Generation Doppler Weather Radar Level II, or NEXRAD L2, with over 270 terabytes (TB) and 180 mission files of historical archived data, have been transferred from the magnetic tape archive at the National Centers for Environmental Information in Asheville, NC, to the cloud infrastructures at Amazon Web Services and Microsoft Azure. Amazon took the data project a step further by partnering with Unidata, a member of the Open Cloud Consortium, to establish a realtime data feed. On October 27, 2015, Amazon announced the public release of these data for all users of their Amazon Web Services platform, free of charge, at http://aws.amazon.com/noaa-big-data/nexrad/.

Why was NEXRAD selected for the Big Data Partnership? The value of NEXRAD observations to the US economy is well established, with NEXRAD ranked #2 in National Observation Value for the recent US National Plan for Civil Earth Observations (2014). Early Big Data Partnership talks had involved NOAA's National Centers for Environmental Information (NCEI), which is the designated Archive for NEXRAD data. These services at NCEI have been financially supported by the NOAA National Weather Service (NWS) for archive and data access for over ten years. While all of the basic archived NEXRAD data are now publicly available upon request directly from NOAA, they are slow and difficult to access from the older magnetic tape library, and have an unwieldy size (270 TB for Level II alone) and particular format (compressed radial volume scans). Nevertheless, NEXRAD is a highly popular dataset for use in industry, with many industrial users for both realtime and archived NEXRAD data. Options include not only the "base" Level II data, but also Level III severe weather products, and a new NOAA integrated product called the Multi-Radar Multi-Sensor (MRMS) dataset. Progress on these other NEXRAD datasets is currently pending.

The Level II archive is of primary interest since multiple derivative uses are possible from this "base" data product - it is the feedstock for new innovations and information products. A comprehensive utilization of the entire NEXRAD archive outside of NOAA had never before been realized, though several industrial entities had tried and failed to do so. In addition, NCEI and the NOAA National Severe Storms Laboratory in Norman, OK, had recently extracted and reprocessed some of the NEXRAD L2 data in order to produce a retrospective (2001-2011) MRMS analysis. Half of that L2 dataset already resided on disk at the Cooperative Institute for Climate and Satellites in NC (CICS-NC), which made it possible to bypass the tape archive for the initial data deliveries and make rapid progress on this first BDP dataset. The MRMS dataset was also identified as a highly desired product, but was still undergoing archive and quality control by NCEI and was subject to some questions regarding intellectual property rights which could impact the "free and open" paradigm of the BDP.

NCEI successfully used its cooperative institute partner, CICS-NC, as a middleman to reduce impact on NCEI's federal labor, operational systems, and data networks, and has copied over 270 TB each to Amazon, Microsoft, and Google, with a smaller subset to the Open Commons Consortium, while maintaining its current level of service to NOAA/NCEI customers. Complete data

extraction from the NOAA archive was a time consuming effort, even under the optimal conditions. Given some system performance issues, extraction from the tape library took longer than initially expected. Some files required retransmission following transmission failures inventory, and verification was a tedious task The project required consistent communications among NOAA, CICS-NC, and the BDP partners to ensure efficient data delivery.

The success of the project required many discussions regarding data use and possible applications. Discussions concerned how to best format and arrange the data to optimize access, instead of optimizing the data for preservation and archive as NCEI had done previously. NOAA's NEXRAD technical experts and scientists played a key role in the success of the project. Although the BDP Collaborators are experts in the movement and storage of digital data within their cloud infrastructure, the data content experts were provided by NOAA. Moving the data was relatively straightforward to accomplish, while understanding the proper use and application of the data was more nuanced and required a sustained conversation among the stakeholders.

Involvement of an expert sub-partner(s) with a Collaborator proved to be very advantageous to both the Collaborator and NOAA. In the case of Amazon's NEXRAD L2 activities, the inclusion of representatives from both The Climate Corporation and Unidata allowed for the data to be re-organized, unpacked, and quality-controlled in a manner optimized to Amazon and NEXRAD data user community. NOAA directly benefitted from the partners' interactions since approximately 0.01% of the original archived files were found to be corrupted. With the partners' feedback and cooperation these files were able to be repaired and resubmitted to the official NOAA data archive.

3. BENEFITS OF THE BDP

Users of NOAA data, including commercial entities, will now have simpler, more complete access to those data. For the example with Amazon, as of October 27, 2015, all of NOAA's NEXRAD L2 data - from 1991 observations to data collected five minutes ago - are available for free (users need not possess an account on Amazon Web Services to discover or download these data), in one place, using a single interface for users. By utilizing the Collaborator's cloud resources, innovators can take advantage of all available data to develop new ideas without the trouble of moving the data across networks. In this case of NEXRAD data, users have access to the processing and visualization resources on Amazon that are integrated with that NEXRAD data storage. Since all data are present, users can develop new, comprehensive, accurate, statistical and probabilistic products that were not feasible before this project.

It is also important to note that industry, academia, other government agencies, and the public now have equal access to NOAA's data. Previously NOAA had a practical advantage on others, since it was possible to use all of this large-volume data only from inside the agency, and special arrangements were needed to receive access to realtime NEXRAD data. New business opportunities are now being created, facilitating the development of applications for custom severe weather products, bird migration studies, etc. Easier access creates possibilities for more users, ideas, and applications, and NOAA will learn more about these possibilities as the BDP project continues.

Perhaps the most significant benefit of the utilization of NOAA data on a commercial cloud platform may be the increased speed at which scientific progress and innovation can now be made. Since the data need not be pulled from a deep archive (all data are online and easily accessed), and need not transverse networks to reach processing (processors are essentially co-located with the data), data utilization is accelerated. Whereas the aforementioned reprocessing of NEXRAD L2 data to advanced integrated products took several years to accomplish ending in 2015, these same products could be now processed – it is estimated – in a matter of weeks. Will this open, free, and fast access to NOAA data speed the development of applications and allow for faster delivery of value-added information products to the market place? The BDP partners will track data usage and access patterns, which may help answer these questions.

4. SATELLITE AND OTHER DATA SETS

NOAA is currently engaged in discussions with the BDP Collaborators regarding the introduction of NOAA satellite data to their cloud platforms. Current access to realtime information from the Polar Orbiting Environmental Satellites (POES) and the Geostationary Orbiting Environmental Satellites (GOES) from NOAA's central processing facilities is limited by bandwidth and system capacities to a small number of non-government users. In addition, since the primary mission of NOAA's

satellites is in support of numerical weather prediction, which is a low latency and critical national application, there may be risks to that mission imposed by additional distribution requirements and system loads. To date, distribution of the realtime satellite data to the private sector and research have been necessarily limited and subject to strict review. Similar operational restrictions are anticipated for NOAA's new GOES-R and JPSS data streams that should begin in 2017. Would the distribution of NOAA's satellite data through one or more BDP Collaborators make them more easily accessible and useable by a wider community of users and for a wider array of applications? By distributing current POES and GOES data through the Collaborators' cloud resources, could the BDP experience help inform an advanced future distribution and utilization scheme for data from NOAA's upcoming GOES-R and Joint Polar Satellite System (JPSS) satellite series?

The current NOAA archive of POES and GOES time series have a significant number of data format and consistency issues that hinder their use. Access to these archived satellite data is also limited by the dependence on the same kinds of slow-performance tape libraries that underpin the aforementioned NEXRAD archive. The BDP is evaluating whether there is a value in providing reformatted and consistent archived data, adhering to the modern self-describing, machine independent formats used for JPSS and GOES-R data streams, to the Collaborators. This would follow the model being applied at this time by Amazon Web Services for NEXRAD L2, which results in one consistent presentation of all data, readily available through a single interface in a common format for both retrospective and realtime data.

Other NOAA data sets now being evaluated for inclusion in the BDP include operational numerical weather and climate forecasts, experimental numerical weather and climate forecasts, fisheries catch/bycatch data, and other ocean satellite and in situ products. The collection of many smaller datasets, such as the fisheries data, offer additional Big Data challenges for which the Collaborators' infrastructure may provide significant advantages.

5. SUMMARY

NOAA's Big Data Partnership (BDP) with Collaborators Amazon, Google, IBM, Microsoft, and the Open Commons Consortium began in April 2015 under five separate 3-year-renewable cooperative research agreements. The business proposition embodied by the BDP is whether the commercial value of NOAA's data is sufficient to allow the Collaborators to host and disseminate NOAA's data, while recovering their costs by providing commercial services and capabilities directly or via industrial partners. NOAA has been working with the Collaborators to identify datasets of mutual interest that also have a business case which may allow for such cost offsets. The first dataset transferred from NOAA as part of the BDP was over 270 TB of NOAA's NEXRAD weather radar Level II data, from 1991-present, to multiple BDP Collaborators. Significant improvements in access times to the NEXRAD data have been observed from Amazon's data store. New applications are being developed on Amazon Web Services and improvements to the NOAA archived data have been realized through cooperative data management activities between NOAA and Amazon's partners The Climate Corporation and Unidata (the other Collaborators have not yet made their NEXRAD holdings publicly available). NOAA and the BDP Collaborators are continuing to evaluate other datasets of interest, including NOAA's POES and GOES satellite data, for inclusion in the BDP. Experience with these current satellite data sets may provide valuable insights for future dissemination methods and utilization schemes for JPSS and GOES-R data that are expected to become available in 2017.

6. REFERENCES

[1] Kearns, E.J., 2014. (Invited) On the Preservation and Application of Climate Data Records in a Big Data World. Big Data From Space 2014. ESA, Frascati, Italy.

ENABLING HIGH-PERFORMANCE ACCESS TO BIG DATA FROM SPACE

Dr Matt Pritchard[1]
Jonathan Churchill[2]
1. Centre for Environmental Data Analysis, RAL Space
2. Scientific Computing Department
1,2: STFC Rutherford Appleton Laboratory, Harwell, Didcot, Oxfordshire OX11 0QX

ABSTRACT

The JASMIN high-performance data analysis infrastructure [1] provides a unique Petabyte-scale environment for environmental scientists to manipulate and analyse huge data sets including high-resolution climate models and multi-mission series of satellite observations. Whole data sets are systematically ingested, on behalf of the climate science and Earth observation communities, into the long-term curated data archive, while members of collaborative projects are able to bring together data sets, either from external sources or generated within the JASMIN environment, for use within their own analysis workflows and potential dissemination to other users. In such cases, it is important to consider the end-to-end path of network data transfers to ensure optimum performance while maintaining security in a multi-user environment.

We describe the JASMIN "Data Transfer Zone": an area of network dedicated to high-performance data transfers, constructed using industry best practices and report on activities underway to serve a large and data-hungry community with an increasing appetite for big data from space and other sources.

Index Terms— e-Infrastructure, Earth Observation, climate modelling, data processing, analysis, archive, transfer, network, high-performance computing

1. INTRODUCTION

This paper describes how the JASMIN data analysis environment has been augmented with capabilities for efficient high-performance data transfer. Through collaboration with international partners in the climate modelling data management community, sharing problems and adopting best practices, the benefits of efficient data transfer techniques are of benefit to CEDA's community of Earth Observation scientists. For example, CEDA is currently embarking on the task of constructing a UK archive of data from the ESA Sentinel series of satellites alongside its existing multi-petabyte data holdings.

We report how datasets required by users in the Atmospheric Science, Climate Modelling and Earth Observation communities are brought together onto the same interdisciplinary analysis platform for collaborative environmental science at the Petabyte scale, and the challenges faced along the way.

2. INTERNATIONAL CLIMATE NETWORK WORKING GROUP

The Centre for Environmental Data Analysis (CEDA) is a partner in the International Climate Network Working Group (ICNWG) [3], a project aimed at improving large-scale data transfers between institutions involved in sharing climate model data as nodes of the Earth System Grid Federation (ESGF) [4]. ICNWG is an international collaboration between institutions in the United States, UK, Netherlands, Germany, UK and the host National Research and Education Networks (NRENs) of those countries, for example JANET in the UK. Partners benefit from technical expertise provided by ESnet, the US NREN.

Through working together and attempting to solve real-world problems of meeting throughput targets for international data transfers, ICNWG partners have learned a great deal about best practice in this area and where possible adopted elements of network infrastructure design which separate science data transfers from standard "business" traffic.

3. OPTIMISATION

In addition to network design optimisation, much has been learned about the techniques required for efficient transfer of large data sets. For example, it is critical to understand and characterise the end-to-end path of a transfer so that any bottlenecks can be identified and rectified. Also, transfer techniques employed by end users do not necessarily make best use of available bandwidth, so it is important to understand the characteristics of the data being transferred, as well as the source and destination locations of the transfer in order to advise on the best technique(s) available for optimum transfer speeds.

Once efficient transfers have been established, it is also important to monitor performance over time so that performance does not degrade.

4. CURRENT USE CASES

4.1. CMIP6 data replication

As part of ESGF, CEDA acts as a data centre for replication of key parts of the CMIP6 data set activity ahead of the IPCC 6th Assessment Report. Huge institutional data flows both in to and out of CEDA via JASMIN are expected, for which efficient tools and practices will be essential.

4.2. UK Academic Mirror Archive for Sentinel data

CEDA is currently constructing a UK academic mirror archive of data from the ESA Sentinel missions and (as of Autumn 2015) is downloading 2 TB/day from ESA, likely to increase five-fold over the next few months.

4.3. End user transfers

Individual users of the JASMIN infrastructure can download data from the CEDA archives in the traditional manner. However many data sets, for example large multi-mission satellite data sets, are inconvenient to download and process elsewhere. In some cases, workflows require combination or comparision with other data, either already in the archive or to be brought in from their home institution or elsewhere. At other stages in the project lifecycle, project data outputs may need to be archived for long-term curation or disseminated to collaborators elsewhere. Users therefore need to be able to perform efficient bi-directional transfers themselves. This paper discusses the challenges of achieving this within a secure, multi-user environment.

5. REFERENCES

[1] Lawrence, B. N., Bennett, V. L., Churchill, J., Juckes, M., Kershaw, P., Pascoe, S., Pepler, S., Pritchard, M. and Stephens, A. (2013) Storing and manipulating environmental big data with JASMIN. Proceedings of IEEE Big Data 2013, p68-75 doi:10.1109/BigData.2013.6691556.

[2] V. Bennett et al., "EO science from big EO data on the JASMINCEMS infrastructure," BiDS14, ESA-ESRIN, Frascati, Italy, 2014.

[3] International Climate Network Working Group https://icnwg.llnl.gov/

[4] Earth System Grid Federation. www.esgf.org

ENABLING BIG-DATA SCIENCE WORKFLOWS IN EARTH OBSERVATION: A EUROPEAN RESEARCH NETWORKING PERSPECTIVE

Jakob Tendel[1], Maria Minaricova[2]

(1) Deutsches Forschungsnetz - DFN, Berlin, Germany
(2) GÉANT Ltd., Cambridge, UK

ABSTRACT

Most academic and publicly funded research institutions conducting cutting edge science with large datasets (Big-Data) are already connected to Research and Education (R&E) network infrastructure. The R&E networking community, woven together at the European level by GÉANT, offers great technical and organizational value to researchers and is actively reaching out to improve collaboration with other e-Infrastructures for enabling even better science in a rapidly evolving science data landscape. We present two case studies of such value added services provided to ESA and EUMETSAT for their part in the European Commission's flagship Earth Observation project Copernicus.

Index Terms— Big-Data, R&E, NREN, DFN, GÉANT, Copernicus

1. INTRODUCTION

Earth Observation is experiencing strong growth in data collection capacity on an unprecedented number of platforms with passive and active instruments of increasing precision and resolution; a trend which is constantly accelerating. Science using Big-Data requires equally state of the art technical foundations and data logistics to enable the most effective research. The core mission statement of the R&E networking community is to deliver these data logistics capabilities to the researcher.

2. RESEARCH NETWORKING

National R&E Networks (NRENs) are organizations in individual countries tasked with providing computer networking services specialized and customized to the needs of research and academic institutions. The primary service is high performance data networking, which should simultaneously provide high velocity throughput of large data transfers as well as low latency connectivity between remote sites for rapid querying of distributed data storage. Most

Figure 1 GÉANT global connectivity

European countries have a designated or de facto NREN such as DFN in Germany.

Deutsches Forschungsnetz e.V. (DFN) is the non-profit association that manages operation and development of Germany's R&E network. Founded in 1984, DFN today counts more than 330 member institutions representing the vast majority of German academia.

DFN's mission is to foster research and higher education by promoting innovation and development through building and providing powerful dedicated network resources (Tendel 2015). Together with its members, DFN is also developing a federated framework for organizing and managing cloud services. This extends the role of DFN from a network service organization to an enabler of e-Infrastructure processes for research and higher education communities.

Beyond their national scope, the European NRENs are participants in a continental collaboration within the GÉANT organization which is Europe's leading collaboration on network and related e-Infrastructure and services for the benefit of research and education, contributing to Europe's economic growth and competitiveness. The primary infrastructure is the GÉANT network itself, which provides high-speed international R&E networking, moving 1000 terabyte/day across Europe and to international partners on all continents (Figure 1) ; ("GÉANT Global Networking" 2016).

However, the organisation also develops, delivers and promotes advanced associated e-Infrastructure services, and supports innovation and knowledge-sharing amongst its members, partners and the wider research and education networking community. ("GÉANT Services" 2016)

3. CHALLENGES

Rapid growth in the amount of science data that is constantly being collected poses a challenge for data logistics such as satellite downlink capacity, transfer to processing facilities, dissemination and archiving, as well as increasing the complexity of performing meaningful data analysis. Without advanced data handling and analysis technologies, scientists will be unable to effectively manage and process the extremely large datasets involved and will be unable to conduct the best science possible ("GÉANT: Global Data Sharing for Earth Observation" 2016).

An important aspect with bearing on operating economics of R&E network operators is that connections (peerings) with general purpose internet networks are often not sized and budgeted for sustained science data traffic but are mostly intended for incidental internet use by R&E users.

Large data streams flowing out of the non-R&E internet into an NREN network can create significant direct or indirect cost by loading a peering point and creating the need for additional capacity, to be paid for by the NREN at considerable expense. Recouping this cost may take the form of directly charging the recipient of that data (research centre, university), or by increasing overall fees because of increased cost of business.

This reinforces the need for properly planning the data pathways of large international science projects to employ the available R&E networks as intended. The NREN/GÉANT infrastructure's core mission is the transfer of large science data and it has been designed and budgeted accordingly.

4. CASE STUDIES

We outline the capabilities of the European R&E networks, represented here by DFN and GÉANT, for addressing the needs of research organisations for international partnership management as well as advanced computer networking and services for Big-Data science projects.

In scenarios where an international research stakeholder requires contact with multiple NRENs, GÉANT could serve as a single point of contact to broker or completely take over separate bilateral contracts in multiple countries and to coordinate the subsequent collaboration.

We present two practical examples of such services in the context of our support for the Copernicus project and indicate some advantages that an infrastructure designed and optimized for research applications can have over commercial internet service.

In the context of Copernicus, DFN and GÉANT are collaborating with the European Space Agency (ESA) to implement a direct connection to GÉANT via the DFN network for high performance and cost-effective access to ESA data products not only to users within Germany and wider Europe but also those on international R&E networks.

At the same time we are also collaborating with the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT) to enable their dissemination of data products through GÉANT, including an innovative implementation of their ("EUMETSAT - EUMETCast" 2016) service using terrestrial transport with the IP multicast protocol for highly bandwidth-efficient near-real-time data streams (Henderson 2014) (Maurice and Minaricova 2014).

4.1 Copernicus: ESA

The European Space Agency ESA is responsible for providing the land surface data from the Sentinel satellites and contributing missions in near real time (NRT) and short-term rolling archive form. While several avenues of access to various collections of data products exist (e.g. Core Services, Collaborative Ground Segment, theme-specific Copernicus Services), this project focusses on providing R&E-connected users a data pathway to the open general purpose site ("ESA Scientific Data Hub" 2016) entirely over the R&E network infrastructure.

The source for ESA Sentinel data products accessed over the Scihub website is the internet-facing exit node of the private ESA Copernicus network which is located in Frankfurt, Germany and where it is already connected to the general-purpose internet. As a result of the collaboration described here, the Scihub is now about to be connected to GÉANT through Germany's national research network DFN. All three organizations worked together to develop a suitable technical and contractual solution which is now being implemented.

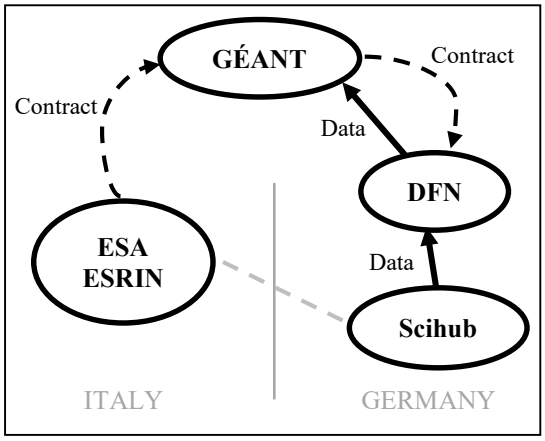

Figure 2 Contract and data relationships

The contractual and technical aspects of the project each presented specific challenges that had to be overcome along the way. Since the data access infrastructure is located in Germany, GÉANT access would be implemented by connecting to the local NREN, DFN, at a suitable service level that uses DFN's existing GÉANT connectivity. However, the ESA department responsible for procuring the Copernicus networking services is located at ESRIN in Italy, which creates cross-border contracting issues and uncertainty with DFN's mandate to supply only users in Germany.

The cross-border contract issue was greatly simplified by GÉANT's ability to act as an intermediary and contract manager for ESA, while leveraging the existing contractual relationship with DFN. GÉANT signed a high level agreement with ESRIN for R&E-based dissemination for Copernicus and sub-contracted the technical work in Germany to DFN (Figure 2).

The ESA Copernicus data node in Frankfurt a.M. is being connected with a 10Gbps access link to the nearest DFN router, also located in Frankfurt. As it happens, this router also hosts one of DFN's two 100Gbps GÉANT interconnects, so any transit traffic to GÉANT can be handled directly on the router fabric. In internet networking terms, the Scihub has been assigned a specially declared network address space (Autonomous System, AS) that can connect (peer) equally to the general purpose internet provider and now to the DFN network in Germany. Any user connected to an NREN with transit rights to GÉANT will have their traffic to and from the Scihub automatically routed though GÉANT and DFN, giving the expected performance, reliability, and cost level of using dedicated science data infrastructure.

One of the additional technical capabilities the R&E networks offer is support for increased data payload per network transmission unit (frame), compared to standard practice on commercial internet services. The broad use of these 'Jumbo Frames' on the R&E networks, which reduces by 6x the amount of protocol overhead, results in significantly higher effective data transfer rate at a given physical connection bandwidth.

The contractual framework for this collaboration was finalised in late 2015 and the technical implementation is on track to be completed in Q1 2016.

4.2 Copernicus: EUMETSAT

Beyond operating weather satellites, EUMETSAT is also responsible for providing data products to the national weather services of member states and international cooperating states with stringent requirements on timeliness and availability.

Aside from conventional download and archive retrieval, EUMETSAT's data delivery mechanism includes the EUMETCast service for dissemination of Near-Real-Time (NRT) satellite data products via digital satellite broadcast to authenticated subscribers. However, future earth observation satellite missions such as the Sentinels, METOP, and MTG will create several step changes in dissemination data rate, which will stretch the current NRT data infrastructure.

To augment the satellite-based service, EUMETSAT approached GÉANT and since 2014 a pilot cooperation has developed a ground-network-based service using a similar targeted broadcast paradigm, for dissemination from EUMETSAT's Darmstadt site in Germany through the DFN network and across GÉANT.

Based on the IP multicast protocol, this EUMETCast (terrestrial) pilot service is currently being tested for the transmission of an expanded set of NRT data products to a select group of EUMETSAT institutional partners on research networks in member countries and international partner countries. This cooperation is eventually intended to deliver NRT data products to R&E-connected users in 31 European countries and worldwide to Asia, Africa, and North America.

The main advantages of the network-based multicast approach versus satellite broadcast is the higher data transfer capacity, currently up to 100Mbps per data service but capable of more. Additional benefits are lower cost of hardware at the receiving end and the ability to maintain the structure and access levels for the data products and remain in the PUSH mode of data distribution.

Using the IP multicast protocol instead of separate per-recipient unicast data streams avoids bandwidth for several identical streams needlessly cumulating near the source by duplicating data flow to individual end-points only at the latest branch point along the network path (Figure 3). A key differentiator of the R&E networks' contribution here is the network-wide availability of multicast-capable infrastructure.

Aside from the networking services, GÉANT took on the role of single point of contact for EUMETSAT in coordinating the contributions of all the involved NRENs. This added service is particularly valued by EUMETSAT as

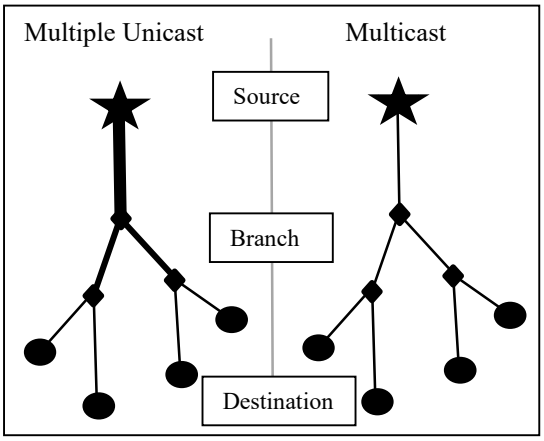

Figure 3 Data rate efficiency of multicast

it leverages GÉANT's existing relationships with NRENs and experience with coordinating projects in the heterogeneous R&E landscape.

Valuable learning and trust-building tools are the network-wide service quality monitoring, reporting, and escalation mechanisms deployed by GÉANT. Learning the day-to-day realities of running a near-operational multicast service on the R&E infrastructure is an important step toward demonstrating the technical characteristics and reliability required for full mission-critical operations. The pilot phase was begun in 2014 as a two year effort which will be evaluated toward the end of 2016 with a view toward further development of operational service characteristics and an expanded group of participants.

5. SUMMARY

Large international earth observation efforts are a natural fit for close cooperation with the R&E networking and e-Infrastructure community, even if they are not exclusively a research activity. The Copernicus project, with its strong focus on operational earth observation and support for decision makers from governments to SMEs, nonetheless generates data of great scientific interest to the research community.

For the ESA part of Copernicus the partners were able to agree and design a flexible and open network access mechanism that allows serving the different user groups according to their varying needs. Institutional and business users might require data access with contractual characteristics best served by a commercial provider, while Big-Data researchers benefit most from using the dedicated R&E networks with stable high-throughput performance and associated service and cost benefits.

For the EUMETSAT part of Copernicus, the R&E network community led by GÉANT took on the challenge of implementing a technically sophisticated multicast data distribution architecture across national and network domain boundaries while simultaneously developing and providing a service and contract management model to the data provider EUMETSAT. This model is unconventional in the R&E space, but serves to relieve EUMETSAT of the need to maintain active relationships with multiple NRENs. The pilot effort has been successful in demonstrating the required capabilities and the challenge now is to see how well the approach will scale with the number of partner sites and how to satisfy the requirements for contractually assured service levels that EUMETSAT sees for a fully operational mission-critical service in the context of their own international and institutional data dissemination obligations.

The R&E networking community offers many services to researchers, directly or indirectly through an institution's IT department. While most researchers on an R&E network will automatically benefit from the network capacity, aspects of the local network configuration can have great impact on long-distance transfer performance and consideration for campus best practices is strongly advised when establishing new data workflows. Other valuable R&E services can be requested, such as advanced lightpath configurations, eduGAIN for federated identity management and single-sign-on to multiple services and e-Infrastructures, or international project and stakeholder management at individual NRENs and GÉANT ("GÉANT Services" 2016).

GÉANT and Europe's NRENs are leading the collaboration on network and are partnering with related e-Infrastructure such as PRACE, EUDAT, or EGI, which may be relevant to earth observation researchers and are all reachable through the R&E networks. Resources are made available to assist research projects in extracting maximum value from the existing infrastructure and services, often in the form of support programs and awards ("Enlighten Your Research.net" 2016).

In the age of Big-Data science, thorough planning of the data logistics is just as important as compute resources or instrument characteristics. We invite all those involved in conducting or planning Big-Data science projects to engage with their national R&E network operator or with GÉANT for collaboration in preparing and implementing robust data management plans that will enable optimal science results.

6. REFERENCES

"Enlighten Your Research.net." 2016. Accessed February 12. https://www.enlightenyourresearch.net/about-2/.

"ESA Scientific Data Hub." 2016. Accessed February 12. https://scihub.copernicus.eu/.

"EUMETSAT - EUMETCast." 2016. Accessed February 12. http://www.eumetsat.int/website/home/Data/DataDelivery/EUMETCast/index.html.

"GÉANT: Global Data Sharing for Earth Observation." 2016. Accessed February 8. http://www.geant.org/Resources/Documents/GEANT-Global-data-sharing-for-Earth-Observation.pdf.

"GÉANT Global Networking." 2016. Accessed February 8. http://www.geant.org/Networks/Global_networking/Pages/Home.aspx.

"GÉANT Services." 2016. Accessed February 8. http://www.geant.org/Resources/#service_descriptions.

Henderson, Tamsin. 2014. "GÉANT Supporting World's Biggest Earth Observation Programme." *GÉANT Connect*, May, 15 edition.

Maurice, Paul, and Maria Minaricova. 2014. "EUMETSAT and GÉANT: Ensuring Delivery of Critical Data." *GÉANT Connect*, May, 15 edition.

Tendel, Jakob. 2015. "Globales VPN für deutsch- japanische Weltraum-Mission." *DFN Mitteilungen*, May, 88 edition.

SCIENTIFIC DATABASES, THE NEW SUBSTRATE FOR EARTH OBSERVATION

Kostis Kyzirakos, Stefan Manegold, Martin Kersten

Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands

ABSTRACT

Lots of Earth Observation data has become available at no charge in Europe and the US recently and there is a strong push for more open EO data. From the perspective of the user, EO data and auxiliary ground observations form "data lakes", which need to be combined and bridged to satisfy an information need. The goal of this paper is to present the distinctive features of the column-store MonetDB that enables users to manage, process and analyze sheer amounts of EO and geospatial data in a coherent way.

Index Terms— Scientific Databases, Array databases, Geospatial information, Linked Geospatial Data, Statistical analysis.

1. INTRODUCTION

With the launch of the European Copernicus program and its Sentinel satellites, more Earth Observation (EO) data than ever will not only be acquired but also be publicly available. Once all Sentinels have been launched, they are expected to transmit back to Earth several terabytes of data per day. However, availability and accessibility - both legally and practically - is only a first step. The EO data gains value once it is analyzed, correlated and enriched with other data sources, and turned into information and knowledge. The sheer volume of the data - both per time and aggregated over time - generated by the Sentinel missions of the Copernicus program poses sever data management and analysis challenges that exceed the capabilities of currently deployed data management and analysis solutions for EO data. The key challenges in the Big Data from Space domain are:

- Provide a temporal and an array-based computational paradigm.
- Handle the sovereignty of existing distributed EO data silos.
- Provide strong support for geospatial data manipulation.

In this paper we present the distinctive characteristics of the column-store MonetDB[1] that allows it to exploit the substrate of vast amounts of EO data. In the next sections we will present the components of MonetDB that allows the user to easily ingest and process EO data utilizing DBMS technologies. In addition we will present how MonetDB was used in the EU Project TELEIOS[2] for developing a Virtual Earth Observatory.

2. MonetDB

MonetDB is a column-store DBMS that provides a modern and scalable solution without calling for substantial hardware investments. The MonetDB system uses the column-at-a-time paradigm, just-in-time indexing, plan-for-reuse, where optimization opportunities are helped by its columnar result materialization. The architecture of MonetDB is presented in Figure 1. Let us now focus on the distinguishing features of MonetDB that address the key challenges in the Big Data from Space domain.

SciQL [1] is a novel SQL-based query language for scientific applications with arrays as first-class citizens. SciQL uses multi-dimensional arrays to represent EO data of various processing levels. This allows us to store EO data (e.g., satellite images) in the database, and query and manipulate their content transparently within the high-level declarative query language. This has three important advantages. First, it allows us to express low level image processing (e.g., cropping, resampling, georeferencing etc.) as well as image content analysis (e.g., feature extraction, pixel classification) in a user-friendly high level declarative language that provides efficient array manipulation primitives. Second, it opens up these algorithms to be optimized by the DBMS (extended) query optimizer. Third, using the seamless integration and symbiosis of relational tables and arrays, query processing and knowledge discovery can exploit both image metadata and image pixels at the same time.

The Data Vaults [2] is a mechanism that provides a true symbiosis between a DBMS and existing (remote) file-based repositories such as the ones used in EO applications right now. The data vault keeps the data in its original format and (remote) location, while at the same time enables transparent data and metadata access and analysis using the SciQL query language. The data vault framework of MonetDB define a virtual scientific data warehouse, that allows a scientist to simply attach an external file repository to the DBMS, e.g., using a URI, and let it conduct efficient and flexible query processing over data of interest. The main idea of the data vaults is to make the DBMS aware of external file formats and thus keep the knowledge and functionality to convert them into database objects inside the DBMS. Currently, the data vault framework of MonetDB supports array-based scientific file formats, such as GeoTIFF, FITS, mSEED, and netCDF, the geospatial file format ESRI Shapefile, and the LAS/LAZ file format that is used for storing LIDAR data sets.

MonetDB comes with an interface to the Simple Feature Specification (SFA) of the Open Geospatial Consortium (OGC) which opens the route to develop GIS applications. This standard defines relational schemata that support the storage, retrieval, query and update of sets of simple features using SQL. The OGC-SFA standard defines functions for requesting a specific representation of a geometry, functions for checking whether some condition holds for a geometry and functions for returning some properties of the geometry. In addition, the standard defines functions for testing named spatial relationships between two geometries and functions for constructing new geometries from existing geometries. Recently, MonetDB has been enhanced with support for handling massive point clouds [3]. By utilizing the LIDAR data vault and extending the GIS module of MonetDB, the user can efficiently manage and process vast collections of point cloud data.

By supporting the OGC-SFA standard, MonetDB serves as the back end for the semantic spatiotemporal RDF store Strabon [4].

This work has been partially funded by the NWO project COMMIT/.

[1] http://www.monetdb.org/

[2] http://www.earthobservatory.eu/

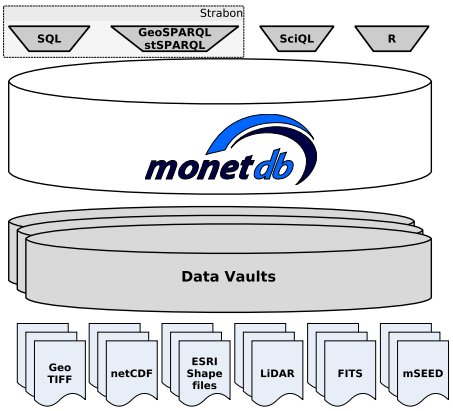

Fig. 1. MonetDB architecture

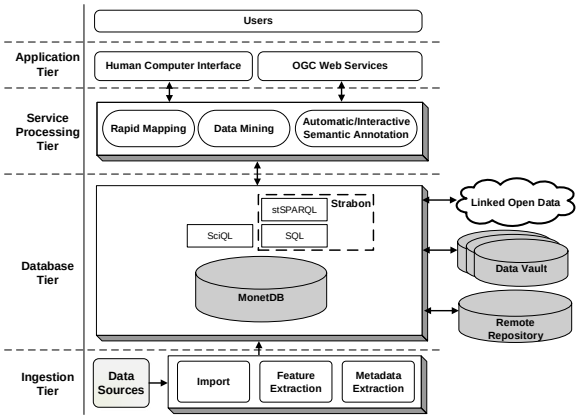

Fig. 2. Software architecture of TELEIOS Virtual Earth Observatory

Strabon allows the user to utilize a SPARQL front-end that supports the state-of-the-art query language stSPARQL and the recent OGC standard GeoSPARQL for storing and querying linked geospatial data that changes over time.

MonetDB comes with a fast connector to the statistical software suite R. While statistical packages are optimized for advanced algorithms, database systems provide fast access to large volumes of data. Taking the best of both worlds, this combination leads to a powerful symbiotic platform to speedup data discovery. For this purpose, MonetDB provides two-fold integration with the R statistical software. For SQL users, the MonetDB/R connector adds support for user-defined functions written in the R statistical programming language in the MonetDB SQL layer. For R users, the MonetDB.R package for R provides a transparent connection to MonetDB.

3. DEVELOPING A VIRTUAL EARTH OBSERVATORY

In this section we present how MonetDB was used in TELEIOS for developing a Virtual Earth Observatory. In TELEIOS, the problem of providing scalable access to vast amounts of EO data and discovering knowledge that can be used in applications was addressed by leveraging and extending data management technologies. A high level view of the software architecture is presented in Figure 2. We can distinguish four tiers in the TELEIOS architecture.

Ingestion Tier: It consists of components that perform data ingestion and content extraction. The data ingestion components transform the original satellite image into a table or array representation that is subsequently stored in the DBMS. The data vaults framework of MonetDB is used during this process for ingesting EO data from external files into database tables or arrays. As a result, the optimization and execution engines of the DBMS gain transparent access to the image content instead of treating it as a "black box" BLOB or a file external to the database. In addition, the ingestion components perform operations like cropping an image to keep only the area of interest and georeferencing an image to a specific coordinate reference system. Such operations are expressed as SciQL queries that are evaluated in MonetDB.

The content extraction components perform feature and metadata extraction from raw files and products. This module allows the user to perform classification operations over satellite acquisitions. Operations like the classification of image pixels as "fire", "potential fire" or "no fire" are expressed in SciQL and utilize its structural grouping capabilities.

Database Tier: It consists of components that provide access to data, metadata and semantic annotations. The main components in this tier are the systems MonetDB and Strabon. Data and metadata that were extracted from the input data during the ingestion phase are stored in MonetDB. Data of various processing levels are stored in MonetDB as multi-dimensional arrays.

Metadata and semantic annotations are stored in Strabon which utilizes MonetDB as a spatially-enabled relational back-end. Metadata and semantic annotations are expressed in RDF so that they can be easily combined with other publicly available linked data sources (e.g., GeoNames, LinkedGeoData, DBpedia) to allow for the expression of rich user queries using GeoSPARQL and stSPARQL.

Service Processing Tier: It consists of Rapid Mapping services, Data Mining services and services for Automatic or Interactive Semantic Annotation. These services enable an EO application to execute processing chains using SciQL and stSPARQL/GeoSPARQL. In the real-time fire monitoring service that was developed in TELEIOS, we performed low level image processing and analysis in a high-level declarative way by performing the classification of hotspots using SciQL and by improving the thematic accuracy of the generated products using stSPARQL/GeoSPARQL.

Application Tier: It consists of applications and services that provide domain specific support to the end user community.

4. THE FIRE MONITORING APPLICATION OF NOA

In this section we present how we instantiated the TELEIOS Virtual Earth Observatory for implementing the wildfire monitoring service provided by the National Observatory of Athens (NOA). NOA operates since 2007 an MSG/SEVIRI acquisition station and has been archiving and processing on a routine basis raw satellite images on a five and fifteen minutes basis from satellites MSG-1 and MSG-2.

Before TELEIOS, the wildfire monitoring service of NOA was organized as follows. The ground-based receiving antenna collects all spectral bands from MSG-1 and MSG-2 every five and fifteen minutes respectively. The raw data are decoded and stored as wavelet compressed images. Afterwards, an application written in Python process the raw data in real-time and performs the following operations: extract and store the raw file metadata in a relational database, filter the raw data files for keeping only data applicable for the fire monitoring scenario, trigger the processing chain and

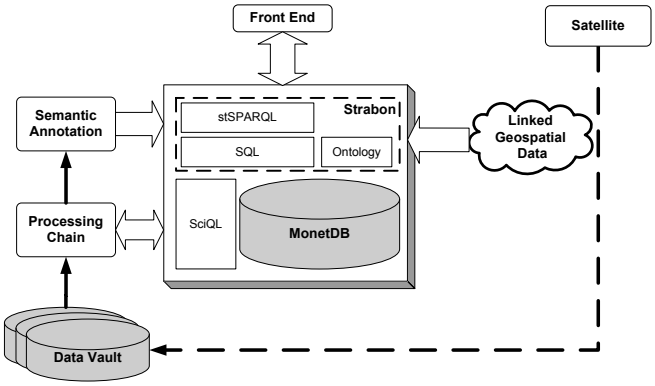

Fig. 3. The NOA fire monitoring service

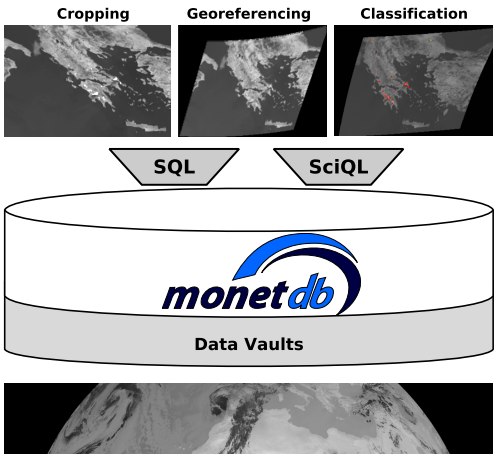

Fig. 4. Ingesting and processing SEVIRI data with MonetDB

dispatch the derived products for archiving. The processing chain comprises the following steps: cropping the raw image, georeferencing to the Hellenic Geodetic Reference System 1987, classifying the pixels as "fire" or "non-fire" using a thresholding algorithm [5] and exporting the final product to raster and vector formats for further processing and visualization.

The processing chain

Let us now present how we implemented the aforementioned processing chain of the fire monitoring service in MonetDB using the Data Vaults framework and SciQL.

Loading

The EO data consumed by the fire monitoring applications follow the High Rate Information Transmission (HRIT) or Low Rate Information Transmission (LRIT) formats. Loading such data in a DBMS requires an external library that converts them into a table or array representation that can be handled by the DBMS. In general, DBMS are not aware of external file formats and the knowledge of how such data can be converted to a tabular or array format is kept outside the DBMS.

In contrast, we exploited the extensibility of MonetDB and developed a module that allows the user to load directly an HRIT file into a relational table or a SciQL array. The module provides the SQL function "HRIT_load_image()" that takes as input a URI with the location of the HRIT file and returns a table/array.

Since the classification algorithm used for the fire monitoring application uses as input the IR bands 3.9, and 10.8, we invoke the "HRIT_load_image()" function for ingesting the input data into MonetDB. As a result, the contents of an acquisition are stored in a SciQL array that follows the schema that corresponds to the following SciQL statements:

```
CREATE ARRAY hrit_T039_image_array
      (x INTEGER DIMENSION, y INTEGER DIMENSION, v FLOAT);
CREATE ARRAY hrit_T108_image_array
      (x INTEGER DIMENSION, y INTEGER DIMENSION, v FLOAT);
```

Cropping and georeference

NOA is interested only in a specific part of the image that is received from the satellite. Cropping the area of interest from the input image is performed in a straightforward manner using a two dimensional range query. The following SciQL statements select part of an image and insert it into a new table:

```
-- array to (temporarily) hold cropped images
CREATE ARRAY cropped (
r INT DIMENSION [200],  -- [windowRows],
c INT DIMENSION [320],  -- [windowCols],
temp039  DOUBLE, temp108  DOUBLE);

CREATE FUNCTION XRIT_2_cropped ( datadir string,
satname string, imgdate string, imgtime string,
windowCols int, windowRows int)
RETURNS ARRAY (
r INT DIMENSION [windowRows], c INT DIMENSION [windowCols],
temp039  DOUBLE, temp108  DOUBLE, coordinate INT)
BEGIN
-- cropping window upper left corner columns offset
declare offsetCols int;
set offsetCols = 2100;
-- cropping window upper left corner row offset
declare offsetRows int;
set offsetRows =   50;
RETURN
SELECT [r - offsetRows], [c - offsetCols], temp039, temp108
FROM hrit_T039_image_array
WHERE
r >= offsetRows AND r < offsetRows + windowRows AND
c >= offsetCols AND c < offsetCols + windowCols;
END;

INSERT INTO cropped
SELECT [r], [c], temp039, temp108, coordinate
FROM XRIT_2_cropped(datadir, satname, imgdate, imgtime,
                    windowCols, windowRows );
```

After cropping the input image, the resulting image needs to be georeferenced. Since the MSG satellite is geostationary, a pre-calculated transformation needs to be applied to the cropped image. The image is re-sampled into a larger size and a two degree polynomial is applied in order to map pixels of the cropped image to the pixels of the georeferenced image. The following SciQL statements georeference an MSG1 image (the pre-calculated coefficients have been omitted for brevity).

```
-- array to (temporarily) hold georeferenced images (MSG1)
CREATE ARRAY georef1 (
rr INT DIMENSION [245], cc INT DIMENSION [319],
y DOUBLE,x DOUBLE, temp039 DOUBLE, temp108 DOUBLE);

CREATE FUNCTION XRIT_2_georef ( datadir string,
satname string, imgdate string, imgtime string,
```

```
ncols int, nrows int, xydim int)
returns array (
rr int dimension [nrows],
cc int dimension [ncols],
y double, x double, temp039 double, temp108 double)
BEGIN
-- simplified version for presentation
RETURN
SELECT [rr], [cc], y, x, temp039, temp108
FROM (
 SELECT newimage.rr AS rr, newimage.cc AS cc,
 newimage.y  AS y, newimage.x  AS x,
 coalesce( orgimage.temp039, 0 ) AS temp039,
 coalesce( orgimage.temp108, 0 ) AS temp108
 FROM (
  SELECT rr, cc, y, x,
  CAST( round( rowsMap, 0 ) AS int ) AS resampleRow,
  CAST( round( colsMap, 0 ) AS int ) AS resampleCol
  FROM (
   SELECT rr, cc, y, x,
          rxx * x * x + ryy * y * y + rx_ * x +
          ry_ * y + rxy * x * y + r__ AS rowsMap,
          cxx * x * x + cyy * y * y + cx_ * x +
          cy_ * y + cxy * x * y + c__ AS colsMap
   FROM (
    SELECT rr, cc,
           ulymap - rr * xydim AS y,
           ulxmap + cc * xydim AS x
    FROM (
      SELECT coordinate / ncols AS rr,
             coordinate % ncols AS cc
      FROM (
       SELECT coordinate
       FROM cropped
      UNION
       SELECT coordinate + windowRows * windowCols
              AS coordinate
       FROM cropped
       WHERE coordinate < arrayDiff
      ) AS newCoordinatesRC
     ) AS newCoordinatesYX
    ) AS newCoordinatesMap
   ) AS newCoordinatesResample
  ) AS newimage
LEFT JOIN (
 SELECT r, c, temp039, temp108
 FROM cropped) AS orgimage
ON newimage.resampleRow = orgimage.r AND
   newimage.resampleCol = orgimage.c
) AS georef
ORDER BY rr, cc;

INSERT INTO georef1
SEELECT rr], [cc], y, x, temp039, temp108
FROM XRIT_2_georef(datadir, satname, imgdate, imgtime,
                   ncols, nrows, xydim );
```

Classification

The fire classification module of the processing chain receives as input the cropped and georeferenced image with two pixel temperatures derived from the IR bands 3.9, and 10.8. According to the classification algorithm [5], each pixel is classified by computing the standard deviation of the temperatures inside a 3x3 window. Figure 5 shows the classification algorithm in SciQL.

The query first computes for each pixel the standard deviation for each of the two bands. It uses the structural grouping capabilities of the SciQL, in order to gather for each pixel the values of its neighbors inside a 3x3 window. The classification process outputs a per-pixel value of 0, 1, or 2 that denote fire, potential fire and no fire respectively. The decision is based on thresholding. A set of 4 thresholds, one for the temperature of the IR 3.9 band, one for the difference between the temperatures of the IR 3.9 and the IR 10.8 band, and two for the standard deviations of the two temperatures, are used for the classification of the pixel. The actual choice of thresholds used in the figure are those for an image acquired during the day. During the night a different set of thresholds is used.

```
CREATE ARRAY hrit_T039_image_array
      (x INTEGER DIMENSION, y INTEGER DIMENSION, v FLOAT);
CREATE ARRAY hrit_T108_image_array
      (x INTEGER DIMENSION, y INTEGER DIMENSION, v FLOAT);

SELECT [x], [y],
 CASE
  WHEN v039 > 310 AND v039 - v108 > 10 AND
       v039_std_dev > 4   AND v108_std_dev < 2
   THEN 2
  WHEN v039 > 310 AND v039 - v108 >  8 AND
       v039_std_dev > 2.5 AND v108_std_dev < 2
   THEN 1
  ELSE  0
 END AS confidence
FROM (
 SELECT [x], [y], v039, v108,
  SQRT(v039_sqr_mean-v039_mean*v039_mean) AS v039_std_dev,
  SQRT(v108_sqr_mean-v108_mean*v108_mean) AS v108_std_dev
 FROM (
  SELECT [x], [y], v039, v108,
   AVG(v039) AS v039_mean, AVG(v039*v039) AS v039_sqr_mean,
   AVG(v108) AS v018_mean, AVG(v108*v108) AS v108_sqr_mean
  FROM (
   SELECT [T039.x],[T039.y],T039.v AS v039,T108.v AS v108
   FROM hrit_T039_image_array AS T039
   JOIN hrit_T108_image_array AS T108
     ON T039.x = T108.x AND T039.y = T108.y
  ) AS image_array
  GROUP BY image_array[x-1:x+2][y-1:y+2]
 ) AS tmp1;
) AS tmp2
```

Fig. 5. Hotspot detection algorithm in SciQL

5. CONCLUSIONS

In this paper we presented how modern database technologies can provide a solid basis to manage EO data in the years to come. We highlighted the capabilities of MonetDB in providing a substrate for managing, processing and analyzing vast amounts of EO and geospatial data. We presented how MonetDB was used in TELEIOS for developing a Virtual Earth Observatory and we presented how we implemented a fire monitoring service using some of the distinctive features of MonetDB.

6. REFERENCES

[1] Ying Zhang, Martin L. Kersten, Milena Ivanova, and Niels Nes, "SciQL: bridging the gap between science and relational DBMS," in *IDEAS*, 2011.

[2] Milena Ivanova, Martin Kersten, and Stefan Manegold, "Data Vaults: A symbiosis between database technology and scientific file repositories," in *Scientific and Statistical Database Management*. 2012, vol. 7338, Springer.

[3] Foteini Alvanaki, Romulo Goncalves, Milena Ivanova, Martin L. Kersten, and Kostis Kyzirakos, "GIS navigation boosted by column stores," *PVLDB*, 2015.

[4] Kostis Kyzirakos, Manos Karpathiotakis, and Manolis Koubarakis, "Strabon: A semantic geospatial DBMS," in *ISWC*, 2012.

[5] EUMETSAT, "Active fire monitoring with msg algorithm - theoretical basis document," *Technical Report EUM/MET/REP/07/0170*, 2007, Eumetsat: Darmstadt, Germany.

EVOLUTION OF EARTH OBSERVATION ONLINE DATA ACCESS SERVICES

Torsten Heinen[1], Bernhard Buckl[1], Simone Giannecchini[2], Stephan Kiemle[1], Stephan Meissl[3]

(1) German Aerospace Center DLR, German Remote Sensing Data Center, Oberpfaffenhofen, Germany
(2) GeoSolutions S.A.S., Camaiore, Italy
(3) EOX IT Services GmbH, Vienna, Austria

ABSTRACT

With the Heterogeneous Missions Accessibility (HMA) initiative, the OGC standard "Web Coverage Service (WCS) Earth Observation Application Profile" has been developed to harmonize online access to very large primary Earth Observation data holdings. Although its use in web mapping servers has proven valuable capabilities, this standard is not yet widely adopted. Its acceptance for data download by end users is hampered by the lack of interpretation guidelines and its complexity requiring considerable server and client implementation efforts.

In this context, the project "Evolution of EO Online Data Access Services" funded by the European Space Agency (ESA) and presented in this paper analyses relevant scenarios and technologies for data publication and access, identifies potential for improvements of standards and their implementations, prototypes and evaluates selected improvements and proposes standard extensions for future releases. We hope hereby to considerably improve the acceptance of online EO data access services and standards and to promote their evolution and diffusion.

***Index Terms*—** Web Coverage Service, Earth Observation, Online Data Access

1. INTRODUCTION

The challenges for Earth Observation (EO) data providers have increased considerably over the last decade. On the one hand, particularly Payload Data Ground Segments (PDGS) face ever increasing data rates of new sensor generations while maintaining and operating large archives of historical, heterogeneous sensor families. On the other hand, demand and requirements for instant access to large volumes of full resolution EO datasets have increased dramatically with rising processing capabilities of EO users such as Copernicus Downstream Services and value adding service providers in general [1].

In this context, the WCS Earth Observation Application Profile (EO-WCS) has been developed within HMA as a complementary interface to traditional file based download interfaces such as HTTP and FTP. Among other features, EO-WCS allows for efficient access to EO datasets by enabling download of spatiotemporal subsets and band constraints of full datasets and thus save bandwidth and storage requirements on the client.

The ESA technology project "Evolution of EO **O**nline **D**ata **A**ccess **S**ervices" (EVO-ODAS) conducted by DLR, EOX, and GeoSolutions specifically addresses current challenges and limitations of this interface in a broader context by analysing, prototyping, and evaluating different scenarios with stakeholders involved in current real world projects. The EVO-ODAS project has been started in October 2015 with duration of 18 months.

2. OBJECTIVES & METHODOLOGY

The aim of EVO-ODAS is to foster the evolution and usage of existing online data access standards by

- analysing relevant scenarios of data dissemination and related technological state-of-the-art
- identifying potential for improvements in these scenarios, through consideration (and possibly evolution) of the standards and their implementations
- demonstrational prototyping and evaluation of selected improvements and
- proposing standards extensions for future releases.

To meet these objectives the project is split in two phases as depicted in Figure 1. This allows for early and regular feedback from Standards Organization (OGC Technical Committee, Working Groups and other relevant programs) as well as Stakeholders contributing real world challenges in their respective field of work (see chapter 3).

The project workflow is following these main activities.

- Definition of online data access scenarios in dialog with identified project stakeholders using specific user stories and generic use cases
- Derivation of system requirements for prototype design and evaluation
- Identification of evolution potentials, issues, and limitations in existing standards and tools
- Setup of infrastructure for prototyping including state-of-the-art technologies for demonstration and validation

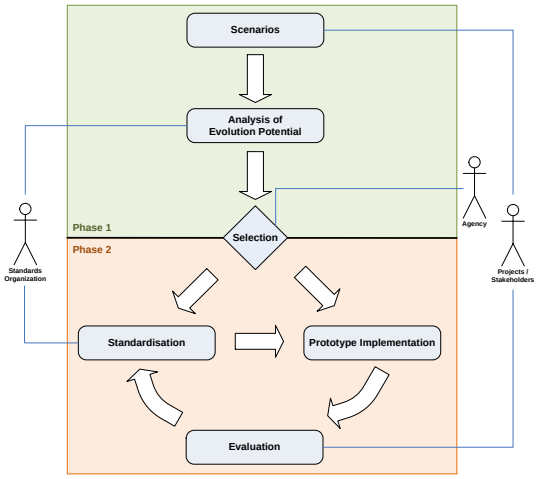

Figure 1: Project workflow and external actor interaction

- Selection of scenarios to be focused on in phase 2 of the project based on the following criteria
  - relevance for one or more use cases,
  - expected benefit in terms of functional improvement, usability, scalability, and performance
  - feasibility for prototype implementation and validation.
- Selection of evolution potentials and issues to be addressed in the prototype
- In parallel solutions to the selected issues are discussed and submitted to OGC
- Design and implementation of prototype including client and server components of different origins
- Prototype validation based on selected scenarios jointly with Projects and Stakeholders

3. STAKEHOLDERS

The project team has identified stakeholder projects serving as source of scenarios and requirements as well as taking profit of the outcome of the EVO-ODAS project and thus allowing a sound evaluation. All the stakeholder projects handle large amounts and long timeseries of EO data. They often have specific requirements concerning the efficient ingestion of data, the discovery of data and access services, the visualization, server-side processing and download of tailored coverages, and the support for federated authentication and content-aware access control and accounting of the end users' activity. A short introduction to the stakeholder's background is given below.

TIMELINE: The DLR project TIMELINE [2] aims at the generation of a 30-year time series of global change relevant variables like surface temperature for Europe. These time series will be used for change and trend detection and a better description and parameterization of land surface processes and their interaction with atmospheric processes.

UKIS: The Environmental and Crisis Information System project (the German abbreviation is UKIS) aims at setting up a framework of modularised and generalized components to build the next-generation information systems. They are integrating access to low-level EO archives and processing chains, extracting and fusing EO-derived information, and visualising air quality forecast simulation results.

EOC-WIS Visualisation System: The work of the Science Communication and Visualization Unit in DLR's Earth Observation Center (EOC-WIS) is primarily intended for non-experts, the media, and scientists active in fields other than remote sensing. Complex science topics are made accessible in a manner appropriate to these target groups and new, intuitively comprehensible forms of presentation are developed.

OPUS: The Scope of OPUS is to implement and demonstrate reusable ODA services integrating large amounts of near real time and offline earth observation data for further usage in collaborative infrastructure activities as planned for Copernicus services [3].

D-SDA: In the context of EVO-ODAS the German Satellite Data Archive (D-SDA) represents a Long Term Archive for EO products. The D-SDA is the entity preserving EO data of many ongoing and historic missions for the long term [4].

4. USE CASES & SCENARIOS

Five use cases describing generic system functions and interactions have been generalized from the stakeholder's user stories and have been aggregated into five end-to-end scenarios.

4.1. Use Cases

EO Data Ingestion: EO data emerge in data centres either in continuous (near real-time) processing streams or in bulk re-processing campaigns. The registration of new datasets and metadata in online data access servers currently follows system-specific methods. A harmonized backend system interface with comprehensive data management functions as well as best practices for its use and integration into EO payload data ground segments are highly desirable.

EO Data Discovery: An efficient data discovery for end users requires metadata to be linked bidirectional to related metadata, dataset series, and services to allow quick evaluation of the datasets different representations by users.

EO Data Access: This use case describes how clients and tools can visualize, download, and perform certain analysis on data, from single datasets up to very large timeseries. Inefficient access to large coverage descriptions, missing asynchronous download requests and basic transformations, and the missing linking of coverage descriptions with other services are examples of already known deficiencies.

Server Side EO Data Processing: Handling big EO data efficiently requires on-the-fly computations as well as the execution of long-running processes. Improvements of standards and implementations include the processing of datasets not available online, the support of asynchronous processing, and the capability to integrate scientific tools and infrastructures for remote processing.

Access Control and Accounting: EO data centres need to account and report on users' access activities for the verification of assured service levels and the identification of potentials performance issues. Access policies need to support the enforcement of region-dependent restrictions. Users should be able to use single-sign-on for overarching services.

4.2. Scenarios

Rolling Data Repository: This scenario demonstrates an online repository that allows discovery of all available product but keeps just a subset available for download (e.g. last 3 month) to cope with limited storage space. Already evicted products can be downloaded by reloading the products on demand.

Timeseries Animation Renderer: In this scenario large spatiotemporal datasets ingested in ODA services are used to feed distributed rendering farms that produce high quality 3D movie presentations.

On-Demand Coverage Analysis: This scenario represents processing of coverages already ingested in ODA services with user specified analytical methods such as spatiotemporal average computations by generating results as new coverages, plots, or videos.

Particle Forecast (and Alerting) Service: This scenario demonstrates server side processing of particle dispersion forecast, like volcanic or fire emissions (Figure 2) based on real-time EO datasets. The 4D (3D + time) dispersion model is published in ODA services and intersected with user specified n-D spaces such as administrative boundaries or flight paths to trigger an alerting service.

Figure 2: Illustration of fire hotspots and their particle dispersion, Source: DLR (CC-BY 3.0)

Interactive Timeseries Browser: This scenario covers the interactive visualization and downloading of large timeseries published in ODA services through web-based tools such as the EVO-ODAS demonstrator client (Figure 3).

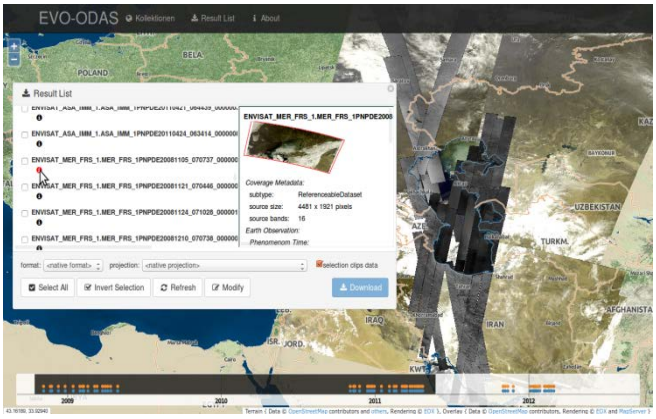

Figure 3: Web-based tool to inspect and browse large EO dataset timeseries. Source: EOX

5. STANDARDIZATION

The adoption of OGC standards for EO online data access services helps to get interoperable access to heterogeneous data holdings without requiring different proprietary tools. Thus, EVO-ODAS carefully analyses existing OGC specifications and documents being relevant for the scenarios identified above in order to find issues, limitations, and evolution potentials to be addressed.

An identified limitation is in ***cross service interaction*** since it is often necessary to link from one service to another, holding different views on the same data. For example, a user might first be presented with an RGB view on some data via WMS before downloading the actual data

via WCS. Another limitation concerns the ***grouping of associated data***. The current coverage definition allows exactly one domainSet per coverage which prohibits, for example, the modelling of an EO satellite scene holding bands with different spatial resolutions, like in Sentinel 2 or Landsat, as a single coverage. Additionally, a typical EO satellite scene is associated with a number of auxiliary data like masks. ***General coverage groupings*** are an identified evolution potential. EO-WCS adds homogeneous and heterogeneous coverage groupings for 2D coverages to WCS. A general coverage grouping mechanism would overcome the 2D limitation and the binding to EO data.

Enhancements and evolutions will be discussed and brought to OGC as necessary. In addition, a prototype ODAS will be set up to demonstrate the feasibility of improvements for the selected limitations and evolutions.

6. IMPLEMENTATION & EVALUATION

For the second phase of the project a subset of the scenarios and requirements will be selected to be addressed. The goal is to implement, deploy, and operate for a consistent period of time a prototype online data access system that demonstrates the results of the standardisation and implementation activities.

The implementation will be mainly performed leveraging on the Open Source tools GeoServer, MapStore, EOxServer, and EOxClient for which the consortium's partners bring together a broad range of complementary skills and real world expertise.

DLR currently operates an Online Data Access Service [5] which delivers atmosphere products containing a complete archive for a variety of EO missions covering a wider range of geophysical parameters (e.g. L2 MetOp A/B GOME-2 from 2007 to 2016). OGC WCS, WMS, WFS and WMTS services are available as well as test instances for WCS-EO and WMS-EO. We foresee that the EVO-ODAS prototype will be deployed on the infrastructure managed by DLR to work in addition to the current infrastructure, sharing some of the input datasets in order to reduce the impact on the storage infrastructure while still accounting for a realistic testing and benchmarking scenario.

The prototype implementation will be validated using large representative EO datasets provided by DLR as part of the selected stakeholders' scenarios. The validation will focus mainly on:

Functionalities: extensions and updates to the standards' interfaces as well as its server side implementations will be verified through the official OGC CITE Team Engine where applicable.

Standards Compliance: the server implementations will be evaluated through a dual setup in the EVO-ODAS prototype. All tests, performance benchmarks, and user acceptance tests will be run against both installations to ensure the syntactical and semantical interoperability of the developed standard extensions. Furthermore off-the-shelf Open Source software such as QGIS and GDAL will be used to validate interoperability.

Performance / Benchmarking: Automated and comparable benchmark tests will be conducted on the prototype to detect, compare, and optimize performance gaps in both server implementations.

With support from ESA, the project partners intend to further accompany the evolution of EO online data access services in the standardization process as well as in the open source implementations.

7. ACKNOWLEDGEMENT

This research is conducted under a programme of and funded by the European Space Agency. The view expressed herein can in no way be taken to reflect the official opinion of the European Space Agency.

8. REFERENCES

[1] Heterogeneous Missions Accessibility (HMA) ESA TM-21, April 2012 ISBN 978-92-9221-883-6, http://esamultimedia.esa.int/multimedia/publications/TM-21/TM-21.pdf

[2] Frey, Corinne and Dietz, Andreas and Bachmann, Martin and Bernhard, Eva-Maria and Ruppert, Thomas and Molch, Katrin and Wolfmüller, Meinhard and Künzer, Claudia and Müller, Andreas and Dech, Stefan (2014) TIMELINE - Processing of AVHRR time series over Europe and North Africa. 34th EARSeL Symposium, 16.-20. Jun. 2014, Warschau. http://elib.dlr.de/92337/

[3] Reck, Christoph and Campuzano, Gina and Dengler, Klaus and Heinen, Torsten (2015) High-Performance Online Data Access for Collaborative Infrastructures. PV 2015 Conference, 3-5 November 2015, Darmstadt, Germany. http://elib.dlr.de/100193/

[4] S. Kiemle, K. Molch, S. Schropp, N. Weiland, E. Mikusch, Big Data Management in Earth Observation – The German Satellite Data Archive at DLR, Big Data From Space BiDS'14, Frascati, ESA/ESRIN, (2014) http://elib.dlr.de/92405/

[5] Heinen, Torsten and Kiemle, Stephan and Buckl, Bernhard and Mikusch, Eberhard and Loyola, Diego (2009) The Geospatial Service Infrastructure for DLR's National Remote Sensing Data Library. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, Volume (4), pp. 260-269. IEEE Geoscience and Remote Sensing Society. http://elib.dlr.de/60172/

CREATING VALUE FROM (BIG) SPATIAL DATA THE BDVA PERSPECTIVE AND ACTIONS

Andrea MANIERI[(1)] *- Anca Liana COSTEA* [(2)] *- Florin SERBAN*[(2)]

[(1)] Engineering - Ingegneria Informatica SpA, Rome, Italy
[(2)] Terrasigna, Bucharest, Romania

ABSTRACT

The digital revolution is creating a new "oil" that can be exploited in a business context: data; Giga-, Tera- or even Peta-bytes of data ranging from satellite signals or telecommunication traffic, till social network behaviour or user profiling. In Europe a set of digital industries constituted a non-for-profit organization, called the Big Data Value Association[1] (BDVA) to exploit the momentum and jointly address the issues that are still open and may hinder the full exploitation of the data "oil". This paper introduces the BDVA activities with particular attention to the Spatial Data for Earth Observation, providing a roadmap that is proposed as part of the research agenda of the BDVA.

***Index Terms*—** Big Data, BDVA, Earth Observation Data, Research Agenda.

1. INTRODUCTION

The BDVA is a fully self-financed non–for-profit organisation under Belgian law. Currently there is a total of 122 members between full and associate, from large companies, SMEs and Academies. The BDVA represents an industry-led contractual counterpart to the European Commission for the implementation of the Big Data Public-Private-Partnership (a contractual Public-Private-Partnership (PPP)[2]). The basic principles are openness, transparency and inclusiveness.

The main role of the BDVA is to provide input to the Big Data PPP strategic research agenda (SRIA)[9] and its regular updates, defining and monitoring the metrics of the PPP and joining the European Commission on the PPP board.

The objectives of the Association are to boost European Big Data research, development and innovation and to foster a positive perception of Big Data opportunities. It aims at:

- Strengthening EU competitiveness and ensuring industrial leadership of providers and end-users of Big Data technology-based systems and services;
- Promoting the widest and best uptake of Big Data technologies and services for professional and private use;
- Establishing the excellence of the EU science base for the creation of value from available Big Data sources.

The work of BDVA is organized in Task Force's (TF), each focusing on one specific topic, where members contribute on a voluntary basis. The focus of a specific Task Force is on "Big Data Applications" with subgroups addressing specific market sectors and needs of scientific domains; one of the most relevant domains is related to the exploitation of Earth Observation data.

2. EARTH OBSERVATION DATA

Our Earth is facing unprecedented climatic, geomorphologic, environmental and anthropogenic changes, which require observation and monitoring at local, continental and global scales, thus resulting in a multitude of new orbital and suborbital Earth Observation (EO) satellite sensors.

EO data volumes are increasing at a rate of several Terabytes per day: for as example, Sentinel-1A only, launched on the 3rd of April 2014, already delivers high-resolution SAR global data every 12 days at rate of 2.5TB per day. The promoted Copernicus[3] free and open access policy, as well as the expected resilience of data sets[4], creates unprecedented opportunities for both industry and

[1] http://www.bdva.eu/

[2] The Public-Private Partnerships (PPPs) instrument or contractual Public-Private Partnership (cPPP) emerged as part of the European Economic Recovery Plan (EERP), initiated by the European Commission. These public-private partnerships aim at the development of innovative technologies in the key industries of Europe. Basically, the public sector funds 50% of the PPPs, the other half of the funding comes from the private sector. From the public sector, the European Commission takes part in the PPPs by providing variable shares of the EU budget for each partnership respectively..

[3] http://www.copernicus.eu/

[4] Ensured by the structure of the constellation, made of two orbital satellites and the availability of a third on the ground, ready to substitute one of the two in case of failure.

scientific communities to fully exploit these data to provide services to end-users and citizens.

However, not only does the volume of data represent a challenge, but also the variety of sensors (acquiring the data) and data formats. Due to the high complexity of the interpretation process which involves image analysis, physical parameter retrieval, model assimilation, multi-temporal, multi-sensor analysis, integration of sensor parameters and other sources of information and knowledge, only few experts can fully exploit the available Copernicus Data. In this context, the major challenges are the democratization of access to data, their efficient exploration and the timely delivery of meaningful extracted information: as easy as any start-up can develop its own mobile app onto open datasets, such as Google Maps™.

The BDVA identified a number of research challenges for the whole technological stack that will be further addressed in the coming years: from enhanced development in Exascale hardware and related computing models till multi-temporal data analysis, data management or information extraction tools; from general analytical methods for the exploitation of the information contained in Satellite Image Time Series till Content Based Image Retrieval technologies and visual data mining techniques.

3. EUROPEAN JOINT EFFORTS ON EO DATA

Together with these difficulties, several opportunities for business development arise in the EO data markets (the so-called midstream and downstream markets). The resilience of the Copernicus services and the open data policy for the free access to the (low-res) data, could boost the adoption of EO data [10] by other markets or players, nurturing new ideas and services, and promoting cross-fertilisation between domains. At the same time, the proposed Directive of the European Parliament and the Council on the "dissemination of Earth observation satellite data for commercial purposes", approved by European Commission on June 2014[11] and now under discussions with the Member States within the EU Council, represent the attempt to remove the latest legal obstacles to clearing the market for High Resolution Space Data (HRSD). Given the importance of the Big Data aspects in space domain, as of June 1st 2015, a new Unit - Space data for Societal Challenges and Growth has been established within the DG Internal Market, Industry, Entrepreneurship and SMEs/Space Policy, Copernicus and Defence.

The EC, in the last years supported several initiatives in the most promising field for the exploitation of EO data, namely: Land, Marine and Atmosphere monitoring, Emergency Management, Security, and Climate Change. Each of them is now operational with services offered to the midstream operators. For example, the Land Monitoring services, initiated by the geoland2[5] project (57 partners from 20 European nations) demonstrating an impressive portfolio of about 200 different products and applications, then further enhanced by the GIO-land project[6], supported by the European Environmental Agency.

Another European intergovernmental organization closely contributing to this action is the European Space Agency (ESA). Given EO data volumes are increasing due to the new ESA launched Sentinel missions within the European COPERNICUS programme, more and more initiatives regarding exploitation of Big Data technologies and EO data are funded under different ESA programmes, such as the Thematic Exploitation Platforms[3].

This challenge is far from being completely addressed, opening important opportunities for exploiting EO data using new (Big Data) technology both by industry, at large, and by EO service providers, thus boosting both the midstream and downstream markets. BDVA identified four dimensions that relate Big Data to the Space domain: i) using big data technology to fully exploit the potentials of EO data, ii) the impact of big data instruments and methods for the next generation of space technologies, iii) the legal implications of the data deluge and iv) the education gaps in future researchers and professional carrier paths.

4. EXPLOITING THE POTENTIAL OF EO DATA

Exploiting Big Data in a Company means understanding "what Big Data means to that company and its customers". The concept of Big Data is then put into the context of the market and the business processes. Understanding the relevance of Big Data in the business context encompasses the following:

- The ability of an organization to access unimaginable amounts of structured and unstructured data both internally (IT systems and interconnected Infrastructures) and through external resources (e.g. data brokers, affiliates or partners).
- The value generated to the company business by capturing, structuring and analyzing high data volumes, and understanding the relationships within and between data.
- The specialized tools and specialized employees (e.g. data scientists) to enable such capture, curation, storage, search, sharing and analysis of the data in a way that is valuable to the organization.
- Analyzing and addressing the potential limitations and legal risks and issues associated the collection, analysis and use of Big Data (and the insights derived from it).

[5] http://esdac.jrc.ec.europa.eu/projects/Geoland2/
[6] http://www.eea.europa.eu/themes/landuse/gio-land

Current exploitation paths for EO data are related to i) the access to data sets in a bundle or ii) the acquisition of final high-resolution information (such as maps). Within the downstream market other opportunities arise from the exploration of available (low-resolution) data sets and the ability to programmatically integrate EO data with other in-situ data, sensors, social network data and alike.

As for any other Big Data source, the value that can be extracted from EO data is not certain a-priori: a set of services enabling exploration of (low-resolution) data, aiming at confirming some business hypothesis, may be an incentive, and then offering the use of (high-resolution) data on a commercial basis, when the expected benefits are proven.

Similarly, the need for uniformly accessing data from different satellites (known as SuperTEP [3]) and other data sources (e.g. the Web, Internet of Things, Open and linked Public Data, etc.) is foreseen as an incentive for the market that will reduce the up-front investment and boost the creativity of applications (web or mobile) developers, like Google with the Google App Engine) [14].

5. BIG DATA IMPACT ON SPACE TECHNOLOGIES

A number of Space Technologies are impacted by the data deluge and deserve some attention. The following areas have been identified to be the most promising to be included in the Big Data PPP Strategic Research Agenda.

5.1 Satellite Image Time Series (SITS)

Considering the continuous, yet, irregular, data acquisition, the opportunity to generate and evaluate SITS became necessary to be exploited, for extracting significant information regarding the complex transformation and evolution processes on ground, these techniques providing solutions for the automatic discovery of regularities, relationships and especially temporal interdependencies, and leading to a better and easier understanding of the underlying processes that causes the detected changes [4].

Presently, there are huge amounts of data, suitable for SITS data mining, covering more than 20 years, with new data added every day. The new generation of high resolution sensors, giving access to detailed image structures, should be better exploited. The current functionalities of the EO Payload Data Ground Segment (PDGS) systems have to be enlarged with extraction and access to the information content of SITS, for enhanced and long-term data exploitation of already acquired and forthcoming data (e.g. SENTINEL 2 data).

5.2 Content Based Image Retrieval (CBIR)

The main idea in content-based retrieval consists of searching semantic similarities. However, the analysis of the collection by a machine is only able to provide similarity by data processing, resulting in a key issue because the meaning of an image is rarely self-evident. This can define a semantic difference between the information that one can extract from the visual data and the interpretation that the same data have for a user in a given situation [1]. Building ontologies to express a semantic meaning of images is considered to be a really challenging issue repeatedly approached in the field of image retrieval and automatic semantic annotation [2].

5.3 Visual Data Mining(VDM)

VDM is a generic term to refer to visual data exploration tools, which allows interactive data presentation in order to increase users' capabilities to understand the information content of large data sets of images and extract meaningful, relevant semantic clusters, together with quantitative measurements presented in a suggestive, visual way.

The VDM tools are able to offer a preliminary insight into a set of optical or radar data, by revealing its semantic structure (natural groups/clusters of resembling scenes/tiles) and some quantitative estimations regarding that structure, through suggestive, simple visual representations. Based on VDM maps, charts and diagrams, the user can make better and faster a priori estimations on the feasibility of the desired image processing.

4. BIG DATA – BIG LEGAL ISSUES

Big Data's increasing economic importance raises a number of legal and privacy related issues still to be properly understood and addressed by regulations (national, European or international).

In the contest of EO Data, the main discussion points are around the ownership of data (its fair use and the commercialization of High Resolution Space Data (HRSD)[11]), the exploitation of EU data by third countries, and the dual use (public vs. security) of sensible geographical information and HRSD[12].

However, other issues may also arise, when EO data are put in context and integrated with other publicly available data[7]. As results BDVA is pursuing research initiatives in the field of privacy preserving methods (computer science) and statistical disclosure limitation (statistical science) in the attempt to promote methods for protecting public release data (e.g. Aggregation, suppression, Data swapping etc.)[13].

[7] A well know sample is Latanya Sweeney showing in her PhD thesis at MIT that 97% of the records in publicly available voter registration lists for Cambridge, MA, could be uniquely identified using birth date and nine digit zip code. By matching on the information in these lists, she was able to identify Governor William Weld in an anonymized medical database.

The high potential of a data-driven economy and the need to strengthen the data value is already acknowledged at a European level. Therefore, effort in spent regarding the sharing, use and interoperability of data, based on common standards and to build an environment ensuring appropriate protection.

7. BIG DATA – EDUCATION

In order to fully exploit the potential of Big Data (including the huge amount of EO data), a key challenge for Europe is to ensure the availability of highly qualified professionals with the right set of skills.

Such professionals, often called Data Scientists, are needed to ensure the ability to extract business or research value from the vast amount of data of different nature available. The National Institute for Science and Technology (NIST) [9] defined initial groups of skills required/expected from Data Scientists as a mix of domain experience, statistics and data mining, and engineering skills.

The qualified Data Scientist is expected to be capable of working in different projects and organisations covering different roles such as Data Engineer, Data Analyst or Architect, Data Steward, etc., and possess the necessary skills to effectively operate components of the complex data and processing infrastructure through all stages of the Data lifecycle.

There are some Data Science programs offered in Europe and the US, mainly as Certificates. Some of them are based on a slight modification or simple re-branding of former curricula on Business Analytics, Data Analytics or Machine Learning. Currently, in Europe, very few are addressing the specificities of the Space domain[8] .

Some initiatives, are attempting to promote the adoption of Data Science curricula (e.g. European Data Science Academy - EDSA[9]) or create consensus on the Data Scientist profile and related professional certification [5] (EDISON project[10]). BDVA is supporting these initiatives in the attempt to promote the industry perspectives in term of skills and knowledge that should be owned by that professional, including specific competences depending on the research or industry domain.

8. CONCLUSIONS

BDVA is at the forefront of the digital revolution by promoting an open forum where all stakeholders can meet, discuss and cooperate on how to make the digital economy a reality for the benefit of the EU and worldwide society. The BDVA promotes the exploitation of Earth Observation Data in conjunction with other open and linked data by supporting the joint resolution of identified gaps and the promotion of joint research initiatives.

Special attention is given to the Space and Earth Observation domains, for their strong implication in the Digital Single Market and the growth of the EU Economy. The BDVA is promoting a Task Force, specifically tailored to addressing the Space and EO domains bridging this community with the wider arena of Big Data Players.

10. REFERENCES

[1] A.W.M. Smeulders et.al., Content-based image retrieval at the end of the early years, IEEE Transactions on Pattern Analysis and Machine Intelligence, 2000

[2] Michaels Lew et.al., Content-based Multimedia Information Retrieval: State of the Art and Challenges, ACM Transactions on Multimedia Computing, Communications, and Applications, 2006

[3] S. Loekken, J. Farres, ESA Earth Observation Big Data R&D Past, Present, & Future Activities, Ground Segment and Mission Operations Department, Earth Observation Programmes Directorate, 2014

[4] S. Laxman and P. S. Sastry, A survey of temporal data mining , SADHANA, Academy Proceedings in Engineering Sciences, 2006

[5]A. Manieri et al, Teaching domain-driven Data Science, to appear in CloudCom'15 proceedings, 2015

[6] N. Savage, Bioinformatics: Big data versus the big C, Nature, 2014,

[7] V. Dhar, Data science and prediction, Communications of the ACM, 2013

[8] Big Data Value Europe, 'European Big Data Value Strategic Research & Innovation Agenda', 2015

[9] NIST Big Data Working Group (NBD-WG). [Online]. Available: http://bigdatawg.nist.gov/home.php

[10] Space-Tec, Assessing the Economic Value of Copernicus: European Earth Observation and Copernicus Downstream Services Market Study. COPERNICUS GIO LOT3, 2012

[11] EC, COM(2014) 344 final, 2014

[12] EU, EDA, CSG, ESA Joint Task Force Civil-Military synergies in the field of Earth Observation , 2010

[13] A. Machanavajjhala Privacy in Data Publishing: Tutorial IEEE ICDE, 2010

[14] https://cloud.google.com/appengine/docs

[8] http://datascience.i3s.uniroma1.it/it
[9] http://www.edsa-project.eu
[10] http://www.edison-project.eu

IMAGES RETRIEVAL FROM BIG DATA ARCHIVES FOR COPERNICUS LAND MONITORING SERVICES

Lukas Brodsky[a], Vaclav Vobora[b], Ondrej Nalevka[b], Marketa Jindrova[b], Adrian Stoica[c], Claudiu Buciu[c], Dan Birtas[c]

[a] Lukas Brodsky, [b]Gisat, [c]Terrasigna

ABSTRACT

This paper introduces an approach of images retrieval from big data archives as Sentinel and Landsat for Copernicus land monitoring services. The specific aim is to retrieve optimal subset of EO image products suitable for the defined user activities (mapping and monitoring tasks within Copernicus), avoiding both manual filtering and transfer of useless data.

1. INTRODUCTION

The pan-European component of Copernicus Land Monitoring Services (http://land.copernicus.eu/) produces land cover / land use information in the CORINE Land Cover data, and the High Resolution Layers. The CORINE data set is mapped over 39 European countries. Five HRLs describe some of the main land cover characteristics as complementary information to CORINE.

It is based on high-resolution imagery, optimally from systematic acquisitions of Sentinel-2 in near future. The required input data is currently collected by different sources: Landsat, SPOT and IRS requiring multi-spectral images with SWIR band. The data used up to now consists of only optical images.

The continental mapping (CORINE and HRL) is applied to 39 EEA countries, which covers <6Mio km^2. Table 1 summarizes estimated data volume acquired in archives for the EEA39 area. The total data volume of virtual constellation Sentinel-Landsat is about 77TB/year.

Theoretical data volume of one cloud-free image coverage accounts to 0.58TB. It can be estimated that one spatial cloud-free coverage for CORINE/HRL mapping amounts 2.7% of all Sentinel-2AB images acquired to archive and 1.1% of all Sentinel-Landsat images acquired to archive within one year.

Table 1: Statistics of yearly archive data volume for EEA39

	Sentinel-2A	Sentinel-2AB	Landsat-7,8	Sentinel-Landsat
No. scenes	36 180	72 360	26 972	99 332
Image volume	18 432 GB	36 864 GB	40 960 GB	77 824 GB
CC volume	3.6 GB	7.2 GB	2.7 GB	10 GB
CC/Image	0.019 %	0.0019 %	0.006 %	0.012 %

Currently the EO metadata archives allow searching and image retrieval based on these database parameters: area-of-interest, time window, collection (sensor) and overall cloud cover (CC) estimation. The CC layers are still not available in the metadata catalogues, even though Sentinel-2 and Landsat-8 have these CC layers in their products. This is probably due to the fact that there is a lack of analytical tools to search images from the catalogues by images content but providing only database search capabilities.

doi: 10.2788/854791

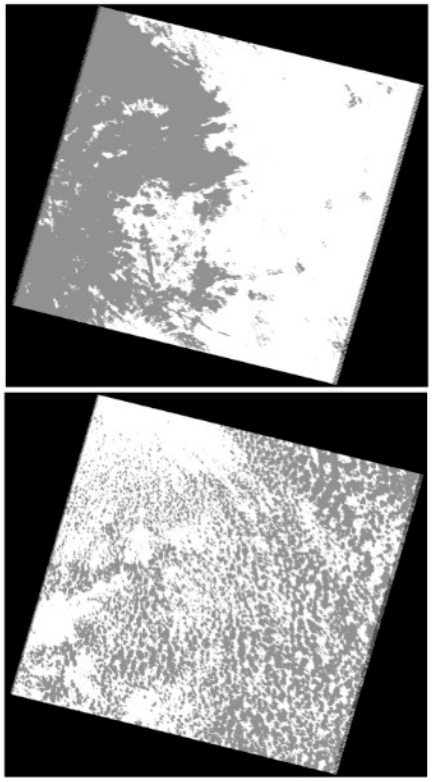

Figure 1: Scene cloud cover layers (grey – clouds, white – valid ground observation, black – no observation)

Sentine-2 L1C product [1] provides cloud mask (in vector format) as metadata, besides the percentage of cloud coverage of the product. This can be considered as major and important step in provision of high-resolution optical images. Such information, cloud cover mask, should be available at best in metadata catalogues providing means for searching capabilities.

2. CONTENT-BASED IMAGE SELECTION

Comparison of the information content of overall scene cloud cover information (CC=40.1% top, CC=45.1% down), as provided in metadata catalogues, and image information content (spatial distribution of the clouds) in CC layers is provided as illustration in Figure 1. The spatial distribution of clouds in the two layers reveals the suitability of the two images for consistent operational mapping.

High complexity of cloud cover distribution in the images prevent from effective use for mapping even they seemingly contain clouds below selected threshold. Such situations are typically identified only when images are downloaded for production phase and evaluated for their content. It leads to delays in production, increase complexity of mapping task or even create gaps in final products as the worst case [2] (see for instance feedback from GIO http://land.copernicus.eu/event or feedback from Urban Atlas Final Report 2011 on rejection of suitable images). Lessons should be taken from such experiences so that future work of a similar nature, especially in case of Copernicus program, can be carried out under better conditions. Moreover, the new Sentinel missions providing systematic acquisitions will result in data explosion on archive side and create new challenges for selecting suitable subsets needed for the operational mapping, which is the purpose of the images acquisitions. On the other hand, the high temporal systematic acquisitions are the only ways forward to acquire required image coverage for the defined Copernicus services.

The aim of this work is to demonstrate how to retrieve optimal EO image products sets suitable for the defined user activities for mapping and monitoring tasks (e.g. Copernicus Land).

We can define the task as optimal subset of images retrieval on user-defined parameters, given the large database of georeferenced satellite images with cloud cover masks available for each pixel indicating validity of land surface (or cloud) observation.

The subset of selected images should be captured in the given time period – at best from similar period of year (typically high vegetation season), within the area of interest (AOI), each pixel of the AOI is cloud-free in at least one image of the subset (or cloud has to be accepted), resulting working units (cloud-free regions) are mostly compact and continuous, and the number of selected images in the subset is small.

A greedy optimization algorithm can be initially selected. It chooses at each iteration the largest working unit of yet uncovered part of AOI. Such algorithm can lead to locally optimal solution. Another approach can be so-called move making approximation algorithm [3] (e.g. alpha expansion). It tries to improve subset selection iteratively by considering a sufficient large neighborhood of the initial selection (greedy algorithm) and solving the optimization task restricted to the neighborhood. The approach can be characterized as best labeling search [4].

Both algorithms were implemented and tested to select best suitable spatial cloud-free coverage from set of available images, based on cloud cover metadata stored in CSW catalogue.

The move making algorithm was initially tested on small regions ~100.000km^2 and finally successfully applied to European scale (EEA39 in Figure 2) providing effective solution for image retrieval from big data archives. The scenarios to be used in operation can include single scene selection with the best cloud-free coverage, temporal coverage, spatial cloud-free coverage to be further extended to spatio-temporal coverage model.

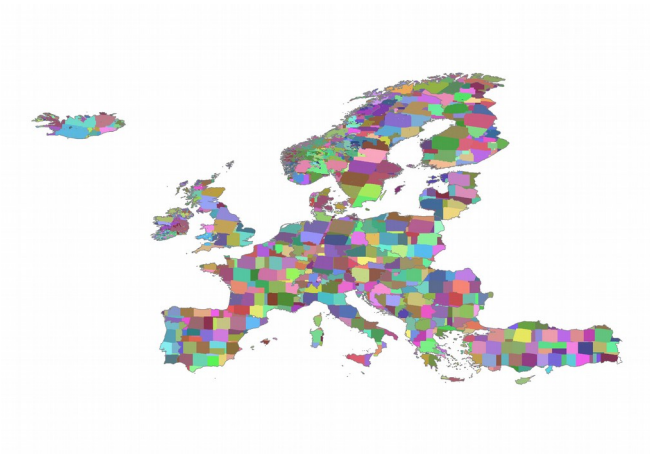

Figure 2 Example of spatial coverage by Landsat for reference year 2012 over EEA39 countries

Figure 3 illustrates simplified work-flow of the overall process retrieving images from the big data archive through metadata collection with cloud cover masks. The CC mask is the key metadata element providing information on the big data archive content. Such work-flow allows content-based analysis of the big data archive based on relatively small metadata collections. The suitability analysis, based on the optimization algorithm, is an intermediate step between user interface and data archives.

3. CONCLUSION

This new analytical approach allows searching metadata archives (if CC layers provided) and selecting suitable subset to be downloaded for mapping purposes. The suitability analysis tool is developed within the ESA GSTP SUCE project for the EO community under Open Source license. The developed analytics should significantly reduce data transfer from big data archives, replace subjective manual work, provide quantitative metadata indicators and hopefully reduce also gaps in mapping products.

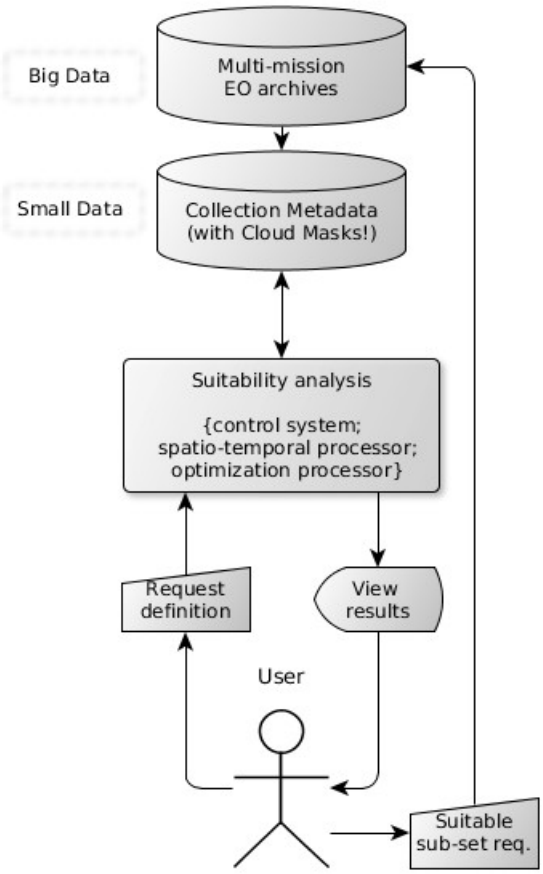

Figure 3: Simplified work-flow of the images retrieval with suitability analysis

ACKNOWLEDGEMENT

The authors acknowledge the financial support of the ESA GSTP project SUCE (Suitability Coverage Engine) with reference 40415001. The authors are also grateful to Boris Flach from Czech Technical University in Prague for his highly valuable contribution in solving the optimization problem.

4. REFERENCES

[1] Sentinel-2 PDGS Project Team, GMES Space Component - Sentinel-2 Payload Data Ground Segment (PDGS) - Products Definition" Reference: GMES-GSEG-EOPG-TN-09-0029.

[2] L. Mequignon, "Urban Atlas Final Report, DELIVERY OF LAND USE/COVER MAPS OF MAJOR EUROPEAN URBAN AGGLOMERATIONS" European Commision, Final_Report_Urban_Atlas _v1.0_20110930.

[3] Andrew Delong, Anton Osokin, Hossam Isack, and Yuri Boykov. Fast approximate energy minimization with label costs. International Journal of Computer Vision, 96:1–27, Jan. 2012.

[4] Boris Flach and Dmitrij Schlesinger. Best labeling search for a class of higher order Gibbs models. Journ. f. Pattern Recognition and Image Analysis, 14(2):249–254, 2004.

Data and Tools: the Thematic Exploitation Platform Concept

S. Clerc[(1)], N. Gilles[(1)], G. Barrot[(1)], D. Illuzzi[(2)], V. Craciunescu[(3)], M. Aspetsberger[(4)], B. Calmet[(5)], M. Datcu[(6)], J. Gasperi[(7)], R. Leone[(8)]

(1) ACRI-ST, France, (2) Planetek, Italy, (3) Terrasigna, Romania, (4) Catalysts, Austria, (5) Thales Alenia Space, France, (6) DLR, Germany, (7) CNES, France, (8) ESA/ESOC, Germany

ABSTRACT

This paper attempts to provide some best practices for the design of Earth Observation (EO) data exploitation services, collected from the observation of successful realization and trends. On-line exploitation, in particular through dedicated thematic platforms, is considered a key concept to allow common users to harness the power of Big Data in Earth Observation. Some key features for a successful thematic exploitation platform are presented.

***Index Terms*—** Thematic Exploitation Platform, Remote Processing, Data Mining

1. INTRODUCTION

With the start of the Copernicus program and the multiplication of commercial initiatives, Earth Observation is entering the world of Big Data [1, 2, 4, 5]. Traditional download-based approaches for data exploitation become unsustainable. It has long been recognized that moving processing closer to the data is needed to overcome data transfer bottlenecks and take full advantage of long time series and near-real time observations (e.g. [3, 6, 7]). Another aspect concerns the difficulty for users to deal with complex and different data types and formats. User-friendly exploitation tools are needed for data mining (visualization, extraction, format conversion, statistical and temporal analysis, classification and semantic analysis, etc.)

Data access services must therefore evolve and propose exploitation tools to users. These tools should be available on-line, supported by a powerful processing capability directly connected to the data archive. This leads to the definition of a Thematic Exploitation Platform (TEP, [8]), i.e. a service providing a set of data and tools for a thematic community of users. We report on some best practices for this approach.

2. BEST PRACTICE #1: MANAGE AREAS AND PERIOD OF INTEREST

While the Earth Observation data can have various type and formats, they all share the notion of a geographic and temporal footprint. The spatio-temporal footprint can therefore be used to connect different data sets or link with semantic feature databases. While most data access services support queries using a spatial and temporal selection, few services take full advantage of the spatial and temporal metadata of data sets. Extracting and saving the footprint in a file, then using this file to perform future queries can be a powerful tool for data mining, see figure below.

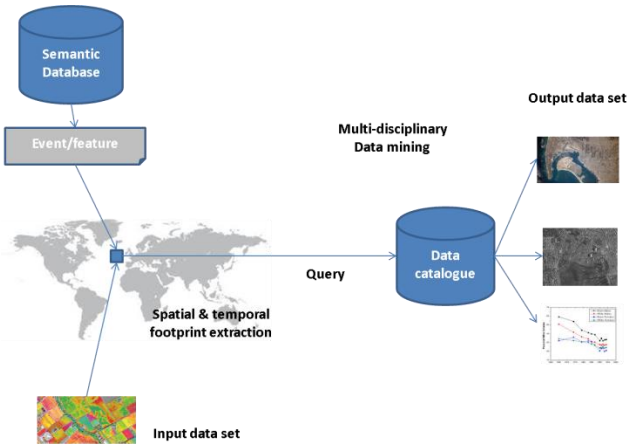

Figure 1: Extracted spatial-temporal footprint used as an input for a catalogue query

For instance, a user can look for an oil-spill event in a specific database. This event can be characterized by a spatial and temporal footprint which can be saved in a file (for instance KML). Allowing the search engine to use this footprint as an input enables the user to retrieve all information available for this event (e.g. satellite radar or optical images, in-situ measurements…). This approach is more flexible and powerful than recording semantic tags as metadata in each EO product. Similarly, spatio-temporal footprint management can support automatic retrieval of match-ups between in-situ measurements and EO data for validation purposes.

In the future, spatial and temporal footprints used for queries from different users could be exploited by data mining techniques to provide relevance rankings [9].

3. BEST PRACTICE # 2: PROVIDE USER-DATA MANAGEMENT TOOLS

Users will stop downloading data only if they can manage them on-line as easily as on their desktop computer. A step in this direction is offered by Google Earth Engine: users can add or delete data in their personal workspace.

Future TEPs should go one step further and offer complete data-cloud services (similar to Dropbox, Google Drive…) with a secured, persistent personal data storage. Users should have the possibility to perform all common manipulations (create, rename, delete), upload and download files easily, and to share files with other users. Obviously, the TEP shall also ensure the confidentiality and security of users' data.

Voluminous data files need not be copied in the user storage area if they can be accessed rapidly for future exploitation: a bookmark (url) to the file is sufficient. A physical copy is needed only if the user intends to modify the data.

The availability of personal data management tools is considered a key feature to convince users to switch from a "download-based" approach to on-line exploitation.

4. BEST PRACTICE #3: SIMPLIFY THE DEVELOPMENT PROCESS

The success of the TEP concept relies on the availability of processing tools which exactly match the needs of end-users. This can only be achieved if the development process is simple enough to enable most users to implement their own processor in the remote execution environment. This requires a dedicated development environment for prototyping. This service is offered by several exploitation platforms (Google Earth Engine, CloudEO workbench, ESA CloudToolbox). Another operational example is the Cal-Val Thematic Collaborative Platform currently deployed for the Sentinel 3 Mission Performance Center [10].

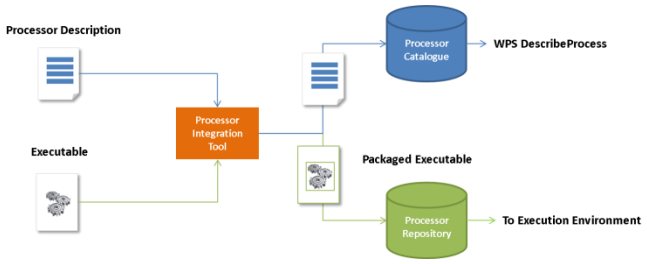

Figure 2: An Integration tool can be used to record the description of the processor and to package the executable.

After development, the integration can be facilitated by an automatic integration tool with a graphical user interface. Users describe the interfaces of their processor using a form, and then upload the directory containing the processor, the associated libraries and configuration files. The processor is automatically wrapped in a container for execution in the remote processing environment. Finally, users shall be able to decide if they want to keep their tools for their own usage, or if they want to share it with other users.

5. BEST PRACTICE #4: SUPPORT QUALITY AND TRACEABILITY

Users want a clear vision of the quality status of the available data and tools. This can be supported by an easily accessible link to data source and reference documentation, digital object identifiers, but also by quality indicators based on user feedback ("stars" rating or user comments). This feature is offered by e.g. CloudEO for products and processors or the Matlab Central repository for code sharing.

Conversely, traceability of usage, enforcement of citation policy and collection of user feedback is important for tool contributors and data providers.

6. CONCLUSION AND POSSIBLE ARCHITECTURE

Thematic Exploitation Platforms can bring a significant change to the way Earth Observation data is used and enable massive data mining through multi-disciplinary data collections.

The TEP should include the following elements:

- A comprehensive internal database of EO, in-situ or model data, as well as a thematic, semantic database of features and events
- Efficient connections to external databases
- A search engine with advanced functionalities, such as manipulation of spatial and temporal footprints
- A storage area for personal data and associated management interfaces
- A remote-processing environment accessible through a web interface,
- Supported by a compatible development environment, and integration tools to simplify the deployment process.

7. REFERENCES

[1] G. Sawyer, "Copernicus & Big Data: A Perspective from the European EO Services Industry," Copernicus Big Data workshop, March 2014.

[2] A. Annoni, "Copernicus and Big Data: Challenges and Opportunities", Copernicus Big Data workshop, March 2014.

[3] G. Barrot, “Big Data only for Big Users?: New challenges for the Sentinel-3 SLSTR and SYN Land PAC”, Big Data from Space Conference, Frascati, Italy, 2013.

[4] H. Jeanjean, “How to meet access requirements for the deluge of Sentinel data: the Sentinel Products Exploitation Platform (PEPS)”, Copernicus Big Data workshop, March 2014.

[5] J. de La Mar, “Taking the data to the users - advanced networks and future Evolution”, Copernicus Big Data workshop, March 2014.

[6] E. Mondon, “Bringing the users to the data (+ applications and services)”, Copernicus Big Data workshop, March 2014.

[7] H. Caumont, F. Brito, E. Boissier, “Big Earth Sciences & the new ‘Platform economy”, Big Data From Space Conference, Frascati, Italy, 2014.

[8] S. Loekken, J. Farres, “ESA Earth Observation Big Data R&D Past, Present, & Future Activities”, CEOS WGISS Meeting #37 Cocoa Beach, USA, 2014.

[9] M. Daoud, J.X. Huang, “Mining Query-driven contexts for geographic and temporal search”, Int. J. of Geographical Information Sci., **27**, 8, 2013.

[10] J. Bruniquel, P. Féménias, Ph. Goryl, H. Bonekamp, “Sentinel-3 Mission Performance Center: paving the way of high-quality controlled data”, EGU General Assembly, Vienna, 2015.

THE (SLOW) MIGRATION FROM IMAGE-BASED INFORMATION EXTRACTION TO DATA STREAM INFORMATION EXTRACTION

Stefano Natali (1,2), Marco Folegani (1,2), Simone Mantovani (1,2)

(1) SISTEMA GmbH, Vienna, Austria
(2) MEEO Srl, Ferrara, Italy

ABSTRACT

The use of satellite data and satellite-derived products has become, in the last 20 years, extensive in many different fields. This has not been a linear process in all scientific and application areas: the "penetration" of satellite-based products has been strongly related to the maturity and diffusion of traditional techniques. In some cases satellite-based products have been used only to replace / substitute existing measurements (e.g. the use of temperature and humidity profiles in weather prediction models) while, in other cases, remotely-sensed information has provided an innovative point of view to complement traditional techniques.

In most cases, the acceptance of the use of satellite-based products has taken years, during which the evolution of space-borne sensors and the data availability has grown faster, with the result that the accepted products and techniques have become obsolete if intended as single product/image. Often the Earth Observation (EO) products were based on the processing of a single image, bi/temporal (intra-seasonal) analyses, and strong human interaction. The current (and mainly the future) status is that the EO data user shall not wonder anymore if a specific image from a specific sensor has been collected on a specific location for a specific time: due to the enormous data availability, the new challenge is how to use the large amount of data available to extract the maximum information content.

Most of the application fields are experiencing / facing a new challenge: moving from a single / small number of images to be processed, toward an intensive / massive multi-dimensional data exploitation scenario, where dimensions are not limited to space and time, but to observational frequencies, various sensors, products, forecast times and so on: streams can be defined as extractions of subsets from the data hypercube, to be processed. This processing shall be guided by theoretical approaches and supported by technological solutions.

The current work is mainly focused on the latter aspect, providing an overview of the available technological solutions to provide tools for massive data.

***Index Terms*—** Big Data, Data hypercubes, array databases, standardization

1. INTRODUCTION

Remotely sensed data from space-borne sensors are nowadays widely used in most of the economic sectors: applications and services related to environmental monitoring, territorial planning, energy market trade, commercial, emergency / crisis management are just few examples. Satellite-based information are used either to replace or integrate other measurement sources, or to provide a further point of view. Few examples of usage of satellite data follow.

1. The International Charter, established in 2000, is a worldwide collaboration among space agencies to make satellite data available for the benefit of disaster management authorities during the response phase of an emergency [1]. Pre- and post-event images (optical and radar) are provided by the data producers and are used by experts to e.g. identify flooded areas and perform damage assessment. Most of the work is done manually to ensure rapid response and the highest level of precision. Since the first activations in 2002, the data availability has grown in resolution and sensors availability, from few mid-resolution products (e.g. Landsat, SPOT, Radarsat) to many sub-metric ones (WorldView-2/3, CartoSat-2, Pleiades, TerraSAR-X, COSMO-SkyMed, …).

2. Satellite applications have always been used to support local and global initiatives for biosphere changes monitoring: as an example, in the framework of the United Nation programme "Reducing Emissions from Deforestation and Forest Degradation (REDD)", remote sensing is used for land cover change / land use mapping and deforestation monitoring (e.g. [3]): single / multi-sensor datasets are used with seasonal and sub-seasonal time series in order to integrate in-field measurements.

3. As last example, to collect a complete dataset of satellite-based SO2 atmospheric concentrations in the 2007-2008 timeframe, the user would have to download five different types of data:

- AIRS from AQUA (data format HDF, 50 x 50 km resolution, 730 files)
- SCIAMACHY data from ENVISAT (data format: N1 or netCDF, 30 x 60 Km resolution, about 120 files)
- OMI data from AURA (data format: HDF, 13 x 24 km resolution, 730 files)
- GOME 2 data from Metop (data format: HDF, 40 x 80 km resolution, about 480 files)
- IASI data from Metop (data format: EPS generic product format, 50 x 50 km resolution, about 1450 files)

It can be easily noted that data download and preparation would be an extremely work-intensive task with a very limited level of automation, that has to be repeated each time a new dataset is added.

These are just few examples to show how the data handling practices implemented so far are inadequate to manage current and future data availability. The following practices become nowadays close to obsolescence:

- Data access to full products from single sensor;
- Bi-temporal / seasonal applications for land cover change;
- Strong human interaction for products creation.

In the following sections a reliable approach to manage the huge availability of data is provided, starting from the definition of data hypercube (Section 2), moving to hypercube storage (section 3) and access (Section 4) solutions, then discussing data processing options (Section 5) and finally outlining the role of standardization in the new concept (section 6).

2. DATA AVAILABILITY AND DATA HYPERCUBES

In 1999, US Vice-President Al Gore articulated a vision of 'Digital Earth' as a multi-resolution, three-dimensional representation of the planet that would make it possible to find, visualise and make sense of vast amounts of geo-referenced information on physical and social environments. Such a system would allow users to navigate through space and time, accessing historical data as well as future predictions (based for example on environmental models), and would support its use by scientists, policy-makers and children alike ([8]). As consequence, there is a need of a new concepts and implementations for the exploitation of this huge and heterogamous amount of data, to implement the so-called massive data-exploitation based approach (see [4]).

A data hypercube is then defined as a multidimensional (spatial, temporal, forecast time, spectral, thematic fields) data collection that describes in the most complete way the data availability in a multidimensional domain.

As example the SO2 dataset as described in Section 1 is a subset of a wider data hypercube (multidimensional data structure) representing all atmospheric science data, for which this subset is an extraction in the "field" dimension for SO2 measurements.

3. DATA STORAGE

A paradigm has to be changed in scientific data exploration: local data storage is not possible anymore. Large, well-maintained and well-connected infrastructures shall be made available to the scientific community, hosting "thematic" data hypercubes. A notable example is the National research Infrastructure for Australia (NCRIS, see, e.g. [6]), that contains satellite data (a large Landsat data collection over Australia since 1973) as well as numerical model data and so on, and supports a wide range of applications (from pure science to industrial innovation) in the field of marine and land cover applications. NCRIS does not only provide storage and data access capabilities, but also processing resources that are needed for data access and exploitation (see next subsections). This is an example as well as many other (see e.g. [5], [7], [9], [10]) that shows that storage and computation resources cannot be decoupled anymore.

4. DATA ACCESS

Once data are stored, there is a need to fast and effective access. Main issues to be faced are:

- Large variety of data;
- Multiple data formats;
- Multi projection / resolution;
- Large data volumes.

Data access is the core of hyper-volumetric data exploitation: an off-line data collection and preparation tool to allow creating user-defined "static" datasets is the main option considered so far; a real time system to interactively navigate the available collections and identify the best dataset for further operations (like e.g. download, processing, …) is the new optimal solution for creating "dynamic" datasets. Data access shall be, from the user point of view, transparent thus not seeing data variety / format / projection issues (seamless moving from one collection to another, contemporary handling multiple collections, …).

Satellite and satellite-based products are traditionally organized as spatial domain arrays, optionally with multiple bands representing spectral channels, vertical layers, multiple products, ancillary data. And this cannot be changed. Various technological approaches are currently used to manage (organize and access) Big Data. Ingesting the various products into a multi-dimensional database structure might be useful to speed up multi-dimensional queries, but has the drawback of at least duplicating the data, taking long data ingestion and data provision time. Solutions not specifically developed for geospatial data (e.g. Hadoop, NoSQL) have proven not to be efficient in accessing and sub setting geospatial data on the fly [12].

Array databases (e.g. rasdaman, SciDB, PostGIS Raster, Teradata) are currently the most adequate solution, with the care of leaving data products stored in the original files, allowing high level of parallelization, achieved via distribution of data into different storage areas, and parallelizing queries on the same data access infrastructure (e.g. making use of the Amazon auto-scaling functionality or any other PaaS).

Standardized interface is a further key issue for a hypercube data access infrastructure: exposing e.g. WCS functionalities, each WCS client can connect and retrieve the requested datasets, via collection, bounding box, timeframe.

In summary, there is a very strong need of a concept / implementation to simplify data access, to automate at the maximum extent all data format issues to reduce the data preparation time and workload, to allow users to focus on data processing more than on data preparation.

5. DATA PROCESSING

A data access platform serving hypercubes subsets permits developing massive (stream) data processing, treating data as multi-dimensional structures. A multi-dimensional dataset can be made by e.g. products from different sensors representing the same field in a defined spatial-temporal domain (e.g. the SO2 example described in Section 2), or different fields (e.g. temperature, humidity, NDVI) in a defined spatial-temporal domain. Homogenizing and processing these data hypercubes means processing at the same time all there data sources in all available dimensions: data homogenization means multi-dimensional interpolation - that can be spatial, temporal, spectral - spatial matching, temporal integration, conversion of unit of measurement.

Data stream processing can be dimensionality reduction, indexes extraction, rule-based analysis or cluster analysis: the global vegetation change maps created by google in 2012 can be considered as simplified data stream processing, since a single-sensor collection (Landsat vegetation index images) has been massively processed along the time dimension do identify vegetation changes related to reforestation, deforestation ([2]). A more complete data stream approach was implemented for the Earth Observation for Climate-related Health risk in Africa (EOCHA) project ([10]), where a data hypervolume formed by collections of precipitation, soil moisture, air temperature, vegetation index are contemporary processed to relate climate change with malaria diffusion in Africa.

In order to implement stream processing tools, the most effective solution is making use of further server-side functionalities, invoked e.g. via WPS / WCPS queries; main difference between WCS and WCPS capabilities are that while WCS allows invoking known and limited number of processing functionalities, WCPS represents a query and processing programming language itself [13].

6. ROLE OF STANDARDIZATION

The role of Open Geospatial Consortium OCG standards [11] is fundamental in order to allow interoperation among the different modules of the same platform, as well as among different platforms. Sharing data access capabilities (e.g. via WCS) permits to access, from a single client, to a multiplicity of data sources as well as the possibility to activate server-side processing via WPS and/or WCPS. The impact of Big Data on standards evolution is also witnessed by the on/going extension of ISO 9075 (Information technology - Database languages - SQL) to include formalization of Multi/dimensional Arrays (SQL/MDA).

7. CONCLUSIONS

Current and future data availability as well as new technological developments allow making real time access and processing of huge data amounts: the paradigm moves from data modeling to massive data exploitation. Data processing services are not anymore based on processing of few images with strong operator interaction: the new challenge is to extract as much meaningful information as possible from a so large data availability that is impossible to be processed / analyzed by hand. The current work aims at providing a conceptual solution to the need of technological implementations to face the continuously growing availability of satellite-based data, paving the way to develop new theoretical approach that exploit the new well accessible data hypercubes.

Various initiatives are on going to implement platforms for massive (stream) data processing, e.g. the ESA Thematic EO Platforms (TEPs) data provision and processing), EarthServer (thematic applications, see [7]) and TAMP (Thematic for Atmosphere, see [9]). As well the standards are continuously adapting to support scientific and technological growth. Once implemented, these platforms will allow

More in general, there is a need of public / private initiatives to implement the large data processing scenario: this is of interest of both data producers to and data users.

10. REFERENCES

[1] The International Charter Home Page - https://www.disasterscharter.org/web/guest/home. Visited on Feb. 12 2016.

[2] M. C. Hansen1, P. V. Potapov, R. Moore, M. Hancher, S. A. Turubanova, A. Tyukavina, D. Thau, S. V. Stehman, S. J. Goetz, T. R. Loveland, A. Kommareddy, A. Egorov, L. Chini, C. O.

Justice, J. R. G. Townshend. High-Resolution Global Maps of 21st-Century Forest Cover Change. Science 15 Nov 2013: Vol. 342, Issue 6160, pp. 850-853 DOI: 10.1126/science.1244693.

[3] E. Hamunyela, J. Verbesselt, M. Herold. Using spatial context to improve early detection of deforestation from Landsat time series. Remote Sensing of Environment 172 (2016) 126–138

[4] T. Hey, S. Tansley, and K. Tolle (Eds.). The Fourth Paradigm: Data-Intensive Scientific Discovery Microsoft Research, Redmond, Washington, (2009) http://research.microsoft.com/en-us/collaboration/fourthparadigm/

[5] http://eodatacube.eu/ (visited on 12.02.2016)

[6] Lesley Wyborn and Benjamin J. K. Evans. Integrating 'Big' geoscience data into the petascale national environmental research interoperability platform (NERDIP): Successes and unforeseen challenges. 2015 IEEE International Conference on Big Data, Santa Clara, CA, USA, October 29 - November 1, 205, 2005—2009.

[7] P. Baumann et al: Big Data Analytics for Earth Sciences: the EarthServer approach. International Journal of Digital Earth, 2015 http://dx.doi.org/10.1080/17538947.2014.1003106
© 2015 Taylor & Francis.

[8] Gore, A., 1999. The Digital Earth: understanding our planet in the 21st century. Photogrammetric Engineering and Remote Sensing, 65 (5), 528

[9] Technology and Atmospheric Mission Platform (TAMP). http://vtpip.zamg.ac.at/ (visited on 12.02.2016)

[10] EOCHA Project portal http://eocha.services.meeo.it/ (visited on 12.02.2016)

[11] Open Geospatial Consortium Standards http://www.opengeospatial.org/standards/ (visited on 12.02.2016)

[12] P. Yue, P. Baumann, K. Bugbee, L. Jiang. Towards intelligent GIServices. Earth Science Informatics 8(3) · June 2015.

[13] OGC WCPS Standard http://www.opengeospatial.org/standards/wcps (visited on 12.02.2016)

MISSION EXPLOITATION PLATFORM PROBA-V

Erwin Goor, Jeroen Dries, Dirk Daems

VITO, Boeretang 200, 2400 Mol, Belgium, http://www.vito.be

ABSTRACT

VITO and partners developed recently an end-to-end solution to drastically improve the exploitation of the PROBA-V EO-data archive [1] and a selection of derived vegetation parameters from the Copernicus Global Land Service [2] by researchers, service providers and thematic users. The analysis of time series of data (+1PByte) is addressed, as well as the large scale on-demand processing of the complete archive, including near real-time data.

From November 2015 an **operational Mission Exploitation Platform (MEP)** on these EO-data available at VITO will be gradually deployed. Several applications will be released to the users, e.g. a time series viewer, a full resolution viewing service, pre-defined on-demand processing chains and virtual machines with powerful tools and access to the data. After an initial release in January 2016, accessible from http://proba-v-mep.esa.int/ [3], a research platform will gradually be deployed allowing users to design, debug and test applications on the platform. From the MEP PROBA-V, access to e.g. Landsat-7/8 and Sentinel-2/3 data will be addressed as well.

***Index Terms*—** MEP Mission Exploitation Platform, PROBA-V, vegetation, data analytics, on-demand processing

1. OBJECTIVES AND BENEFITS

The PROBA-V MEP builds further on the R&D results, i.e. currently existing prototypes, from the ESA/GSTP 'ESE' project, which were further refined in several other projects (EC/FP7 and BELSPO funded) thanks to the active involvement of these projects in the ESE pilots activities.

The PROBA-V MEP has the ambition to complement the PROBA-V user segment by building an operational Exploitation Platform (EP) on these data, complementary data and derived products, addressing hereby the wider vegetation user community with the final aim to ease and increase the use of PROBA-V data. The data offering will consist of the complete archive from SPOT-VEGETATION, PROBA-V and bio-geophysical parameters from the Copernicus Global Land Service.

The reasons for deploying a MEP dedicated to the PROBA-V mission are numerous:

- The data and specifically the time series of daily / ten-daily data from 1998 till present are too big to be downloaded to and processed on the users' premises, at least for the majority of the users.
- On top of the EO-data mentioned above, the platform can co-locate as well complementary data in a way that it is easy accessible. Furthermore tools, libraries and applications which can be used by the large community will be provided. This includes as well the data needed for Cal/Val activities.
- The platform can stimulate collaboration between the users, as we bring together services from various users on the same platform with a number of tools to support the publishing of and to provide feedback on these services. A further focus on documentation, knowledge sharing and user support complements this.
- The platform goes beyond offering standard products by offering in a first place applications to visualise and analyse large time series of data and pre-defined on-demand processing services which deliver user-tailored products. In a next step we will deploy gradually a Virtual Research Environment, being a platform which allows users to develop – debug – test an application on an infrastructure at VITO with access to the complete data archive. Successful applications from third-parties can then be offered as an operational on-demand processing services to the user community on the same platform.
- As an Exploitation Platform (EP) with a focus on open interfaces, we position the PROBA-V mission in an ecosystems of TEPs (Thematic EPs), REPs (Regional EPs) and other MEPs. In the future the PROBA-V MEP can be integrated gradually in a federation of different platforms, including as well Sentinel Collaborative Ground Segments, in line with the current ESA strategy on the 'EO Ground Segment Evolution'.

During the PROBA-V MEP project, which will at least last till the end of the PROBA-V mission in May 2018, several third-party service projects will develop and operate applications on the operation MEP platform. We will address their user requirements to implement the shift of paradigm from "data to user " to "user to data", bridging the gap between the traditional EO ground segment and the scientist or value added industry by providing a one stop shop for access to the full PROBA-V Mission data

(including derived parameters) and to external repositories of similar missions/sensors (including Landsat and Sentinel).

2. TECHNICAL SOLUTION

The PROBA-V MEP will provide scalable processing facilities with access to the complete data archive and a rich set of processing algorithms, models, open source processing libraries/toolboxes and public/collaborative software. The platform becomes the hub processing infrastructure of the mission by functioning as a powerhouse system and open access development environment.

To realise this the platform consists of the following components:

- The existing Product Distribution Facilities [4] and [5], are serving the access to the data archive, both via a Web portal as well as standardised discovery, viewing and data access interfaces. More evolutions on these standardised machine-to-machine interfaces are planned in the near future.
- Hadoop, as a software framework for data-intensive distributed applications, is designed to process large amounts of data by separating the data into smaller chunks and performing large numbers of small parallel operations on the data. It is applied often for processing big data and is applied in this context for the on-demand processing of EO-data, as prototyped successfully in the ESE project. Oozie is used as a workflow processing engine to design an EO-application as a workflow of multiple processes. Spark is used intensively to allow analytics on large time series of data. The Hadoop ecosystem provides furthermore a rich and still growing set of tools which are used to provide fast access to the data in a format needed by the specific application.
- The EO raster data is accessible via NFS and possibly uploaded to the Hadoop Distributed Filesystem (HDFS) using a DataManager which integrates with several catalogues implementing different protocols, so that as well third party-data can be ingested in the platform when needed by a specific user.
- Cloud computing technology enables dynamic resource provisioning and is therefore providing a performing and scalable solution. OpenStack is chosen as cloud middleware. Pre-configured virtual machines will be offered and can run on the OpenStack cluster at VITO, providing the environment needed for users to work with the data and develop/deploy applications on the platform, i.e. containing IDE's, a rich set of tools and access to the complete data archive.
- Interactive Web-based dashboards are designed to provide user-tailored information from the EO-data archives of VITO and other providers, by combining existing components such as AngularJS, Javascript libraries and GIS components into one single solution. The combination of these different components and allowing interactions between these, applied on data available in disparate data stores, offers powerful Web portals to the users in order to make vast amounts of data understandable. We can easily design user-tailored Web-based dashboards which offer at any time near real-time information for the regional extent of interest for the user and in the format chosen by the user.
- A Web portal provides access to all applications and tools offered by the PROBA-V MEP and to the cloud consoles. Furthermore the portal provides all information on the data and components available on the platform and offers tools for e-collaboration and knowledge sharing amongst the users.
- A main concern is security since we allow user to develop and execute their applications on the platform. Their IPR shall be properly protected and the activities of individual users cannot influence the stability of the system and the work of other users. Single sign-on and proper monitoring of used resources are further requirements.

3. CONCLUSION

The platform was launched on 26 January 2016 at the PROBA-V conference in Ghent, Belgium and is accessible from http://proba-v-mep.esa.int/. Three iterations are planned to gradually expand the capabilities of the system and provide new features, in close collaboration with the first third-party projects working on the platform.

The impact of this PROBA-V MEP on the user community will be high and will completely change the way of working with the data and hence open the large time series to a larger community of users. The operational platform is based on recent R&D activities and is in line with the new ESA strategy on the 'EO Ground Segment Evolution'.

4. REFERENCES

[1] http://proba-v.vgt.vito.be/.

[2] http://land.copernicus.eu/global/.

[3] http://proba-v-mep.esa.int/.

[4] http://www.vito-eodata.be.

[5] http://land.copernicus.vgt.vito.be/PDF/.

GERMAN COPERNICUS DATA ACCESS AND EXPLOITATION INFRASTRUCTURE

Christoph Reck, Gina Campuzano, Klaus Dengler, Torsten Heinen, Mario Winkler

German Aerospace Center (DLR), Earth Observation Center (EOC), Oberpfaffenhofen, Germany

ABSTRACT

We present the architecture for the planned German Copernicus collaborative data access and exploitation infrastructure. It shall enable the use of operational earth observation data from the Copernicus ground segment for applications and services developed by commercial, public and scientific institutions in Germany.

Copernicus – the European Earth Observation programme Global Monitoring for Environment and Security – is an ambitious undertaking of the European Commission, the European Space Agency, EUMETSAT and all its member states, to support European citizens, decision makers, scientists and industry with a constant and reliable stream of up-to-date information measured by Earth orbiting satellites.

Due to the high data volume of Sentinel and other future missions, access to Sentinel data in form of download and data distribution service is not enough. The infrastructure also has to provide hosted processing capacity following the Big-Data paradigm of "transferring the software to the data" that can be cost-effectively used by many projects and services in Germany and beyond.

Index Terms – Earth Observation, Online Data Access, Copernicus, Sentinel mission datasets, Big-Data, CODE-DE

1. INTRODUCTION

In 2013, the Space Administration of the German Aerospace Center (DLR) asked the DLR Earth Observation Center (EOC) to conceive an architectural concept for a national platform to mirror incoming Sentinel data for use and offering computing services to its users.

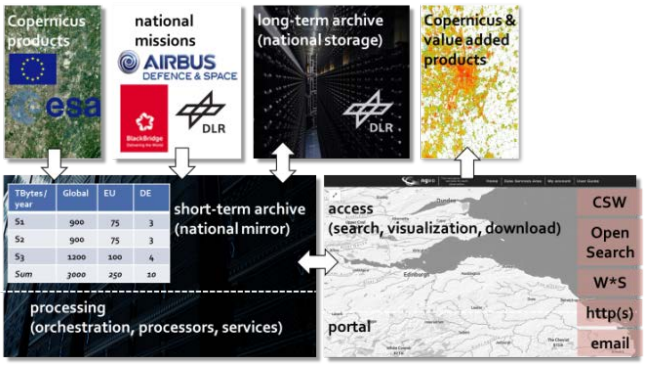

Figure 1: Context of the German national Copernicus data access and exploitation collaborative infrastructure

The primary goal of the Copernicus collaborative data access and exploitation infrastructure is to provide a platform to exploit the possibilities of the continuous data stream of "free, full and open" Copernicus Sentinel data with the following elements:

- Ingestion and Archive
- Search and Access
- Portal and User Management
- Processing and Value Added Products
- Monitoring and Reporting

2. CHALLENGE

Data rates of all incoming Sentinel-1, Sentinel-2, Sentinel-3 and Sentinel-5p earth observation user products itself pose the biggest challenge:

	2014	2015	2016	2017	2018	2019	2020
Yearly volume [TB]	180	966	4 490	6 591	7 250	7 469	8 127
Average Data Rate [Mbit/s]	194	257	1 194	1 753	1 928	1 987	2 162

Figure 2: Estimated data volume and rates for ingesting all Sentinel user products

However, the outbound rate is far larger when this data is systematically processed with different algorithms and accessed by several end users. This requires extremely performant infrastructure and data access methods.

3. ARCHITECTURE

The proposed architecture is highly modular being composed of individual services, which form a general-purpose system, largely based on existing software components, put together in an innovative manner to serve as a high-performant and scalable platform. Figure 3 below depicts the complete architectural concept that is briefly described in the next paragraphs.

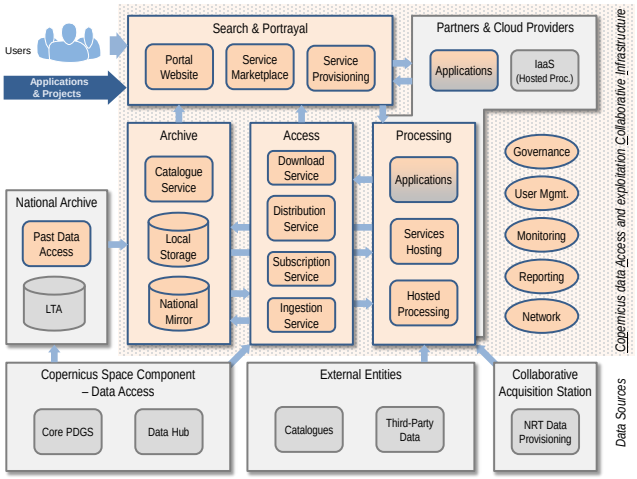

Figure 3: Architectural Concept

Search and Portrayal

Portal Website provides the access point for the collaborative infrastructure elements, portrayed from the registered services, such as the catalogue of the archived products, application provisioning monitoring, etc.

A *Service Marketplace* allows the collaborative infrastructure, external users and projects to register and advertise their accessible services (e.g. a third-party catalogue or an algorithm).

The *Service Provisioning* enables the users to deploy and launch their applications.

Archive

National Mirror provides fast access to data products for global, regional or time-series processing and re-processing. The products provided over the national mirror are the same as disseminated by ESA to the member states.

Local Storage archives the value-added and convenience products of the collaborative infrastructure.

Catalogue Service for the inventory and metadata of the Copernicus Core Services and convenience products available in the Archive. It shall permit dataset discovery visualization and selection, with pointers to the download location.

Access

Ingestion Service is in charge of populating the National Mirror from the Data Hub.

Subscription Service manages registered filters on incoming datasets to provide triggering capabilities for data-driven applications. A trigger may fire an email, dissemination to the users, or drive a data driven workflow.

Distribution Service is in charge of dissemination with equal priority the Sentinel data to the application modules (push service). This contrasts to the capability for direct access to the National Mirror or Local Storage which is subject to polling and workflow delays (pull service).

Download Service provides the means for the users to access products stored in the online archive. The download will provide direct HTTP access to the data products [OGC 13-043] or tailoring and subsetting capabilities via the OGC WCS standard.

Processing

Applications are deployed and run within the collaborative infrastructure. Among them are services and processors for the systematic and on-demand generation of convenience products from Sentinel data.

Services Hosting provides the collaborative infrastructure or deploying and running system applications (e.g. all of the above listed services) and hosted applications for the users. In principle this is a cloud environment with virtual machines and suitable software tools and web interface.

Hosted Processing provides a platform for massive-parallel earth observation data processing, for deploying processors in a cost- and resource-effective way on a cluster close to the data, and job handling with the data selection and processor execution. Also monitoring and accounting functionality is attributed to this element.

Supporting Services

Other cross-cutting services like *Governance*, *User Management*, *Monitoring* and *Reporting*, *Backup* and *Network* infrastructure complement the architecture to its completeness.

4. SOFTWARE SYSTEMS

DLR EOC designed a system utilizing known best breed software and hardware systems assembled to a streamlined simple, scalable and performant architecture covering all interfaces from **Discovery** over **Visualization** to **Download** for users with novel clients:

- Fast catalogue with HMA CSW and OGC OpenSearch interfaces
- Flexible dataset browsing with OGC Web Map Service (WMS)
- High performance data access using HTTP protocol
- Advanced data access using OGC Web Coverage Service (WCS)
- Parallel file system on an online storage attached network (SAN)
- Redundant hardware

The system shall additionally enable retrieval of historical data from the archive.

Discovery

The *EOWEB®* catalogue is based on a database, a metadata-model, ingestion and operation interfaces and user service interfaces compiled for performance [1]. It is configured to hold OGC EOP metadata [4] and provides an HMA CSW standard compliant interface [5] that allows novel clients and user interfaces to comfortably search for data (*EOWEB GeoPortal*, *FEDEO*, *mapshup*, *EOxClient*).

Visualization

Geospatial Data Access Service (GDAS) provides the tools and components to register, describe, access, search and retrieve geospatial data. It is manly composed of:

- *GeoServer* (http://geoserver.org) is an open source server for sharing, processing and editing geospatial data. It offers all major OGC standards services needed for this project (WMS, WFS, WCS, CSW), and also can be setup for INSPIRE conformance.
- *GeoWebCache* (http://geowebcache.org/) is used to cache geospatial data (e.g. for WMS) and therefore speed up the access of this data for the clients.
- *PostgreSQL* (http://www.postgresql.org) is an object-relational database with the PostGIS extension it handles spatial-referenced data.
- *Nginx* (http://nginx.org/) is a web server with a strong focus on high concurrency, performance and low memory usage. It is also used for the proxy and download server.

Catalog Client

The envisaged EOX catalog client provides a as "zoom-in on the data" solution to discover, view, and download available EO data. A timeline and advanced search capabilities allow additional filtering, e.g. by sun angle or cloud coverage.

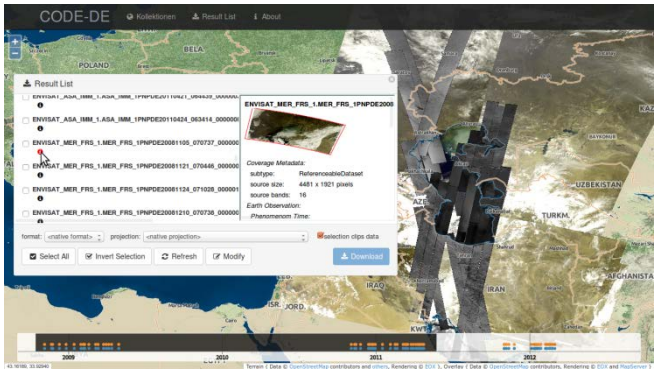

Figure 4: Catalog Client with result details and download links

Download

The large storage allows on-line access to the data products. We have benchmarked the Nginx server and the underlying filesystem to prove the capability of serving files at a total rate of more than 2 GBytes/sec using parallel transfers. *Nginx* was chosen for its performance as well as its simple configuration model that provides existing extension modules as load-balancing, access control and on-the-fly unzipping and content retrieval.

Historical data from the long term archive will be made equally accessible such that these appear to be nearly on-line.

Distribution

To minimize the load accessing recently ingested data, we are designing an efficient data distribution service based on *UFTP* (http://uftp-multicast.sourceforge.net/) for an multicast file transfer mechanism. Each distributed file is preceded by a metadata record that allows the client side filter element to decide whether a data product is needed. At the end of the reception, the client library launches a command that has been configured for the filter condition.

Processing

Bringing the processing to the data is enabled by directly interfacing a processing cluster to the storage area. Processing control and the user interface is allocated in the virtualized or cloud environments. Projects and users can process on the platform interacting on different levels ranging from

- Use of data processors and generic processing workflows and toolbox operators available via the Web GUI, producing products from Copernicus data.
- Simple data access from their VMs, interactive work with tools available (e.g. Sentinel Toolbox) or provided by the project.
- Submission of processing jobs to the cluster, control of processing workflows (orchestration) with platform tools (scripting) or project-specific workflow engines; use of infrastructure services (e.g. subscription) for systematic data-driven processing of new input data; small-scale processing on the VM or (bulk) processing on the cluster.
- Deployment of own processors in Docker containers, or as a service in a VM.
- Provision of data or processing service to other users with publishing in the Service Marketplace.

The processing environment is based on Apache Hadoop using the *Calvalus* environment designed by Brockman Consult GmbH, enhanced with *Docker* (deployment), *Apache Mesos* (resource manager), *Marathon* (service control and scheduling) and *Chronos* (batch processing).

User Management

A central use database and authentication system avoids redundancy and provides single sign on (SSO) capability to the users for accessing infrastructure services after a single login. The user management service consists of the following subsystems:

- LDAP Administration Service, which allows an operator to manage user accounts.
- The User Registration Service for new users.
- The User Self-Management Service allowing a user to manage his own data.
- The CAS Server providing interfaces for authentication of users and services, as well as issuing and validating tickets for the SSO process.

4. HARDWARE INFRASTRUCTURE

The infrastructure hardware configuration consists of a virtualized environment, a high performance computing cluster, and the storage server:

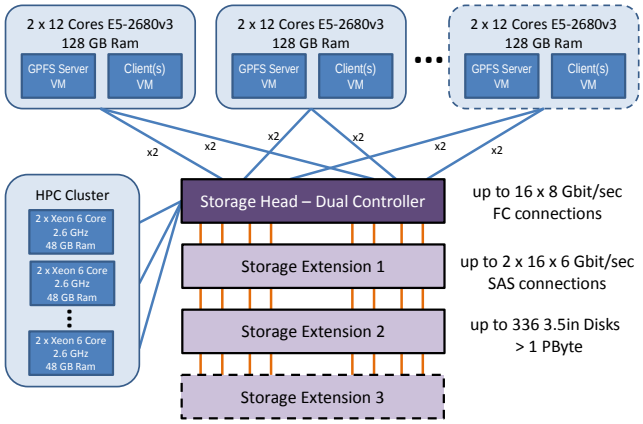

Figure 5: High-Performance server environment

- 3 servers with 2 Intel Xeon E5-2680v3 12 Core 2.5GHz, 128 GB RAM, 2 * 300 GB 10k SAS Disks, 2 FC HBA 8 Gbit/s, Quad Port 10 GE NIC, Quad Port 1 GE NIC, 2 hot swappable 750 watt power supplies.
- 7 HPC servers with 2 Intel Xeon 6 Core 2.6 GHz, 48 GB RAM, 500 GB SATA Disk, Dual Port 10 GE NIC, Dual Port 1 GE NIC, 2 hot swappable power supplies.
- Storage using a GPFS file system with 835 TB in RAID 6 over 282 4 TB NL-SAS disks, dual controllers in one head. Expandable to 12 Petabyte with 4 heads.
- Public Cloud within the DSI vCloud environment.

Network and interfacing components are not depicted. The hosts are connected via 10 Gbit/s network segments. The external connection starts with a dedicated 3.5 Gbit/s connection interfaced to the Internet and GEANT networks by the DFN X-Win provider.

5. CONCLUSION

The German Copernicus Data Access and Exploitation Infrastructure is based on the architecture, software elements and hardware infrastructure outlined in this paper. Most of the system components are re-using existing open-source software, components developed for the DLR multi-mission German Remote Sensing Data Center and the evolution complemented by the OPUS-GMES project that prototyped the:

- utilization of near real time and offline (catalogued) earth observation data
- development and implementation of end to end value adding services
- establishment of state of the art product and service-access
- demonstration of representative production ready service chains

The hardware infrastructure has been setup and tested within the ESA Data Hub Relay project to ensure its capacity and suitability for the high throughput of the envisaged system. Benchmarks showed the capability of serving data to more than 50 clients with a total data rate beyond 2 GBytes/second in the initial minimal setup.

Therefore the presented infrastructure elements have been tested and the risk has been minimized in preparation for the upcoming official German CODE-DE (Copernicus Data and Exploitation Platform – Deutschland) project. Remaining challenges in setting up the operational infrastructure are harmonizing the means for configuration handling of the individual elements, interfacing with the SSO, and the additional effort on the extensions for quota management, prioritization, and ensuring that the security requirements are fulfilled.

SELECTED REFERENCES

[1] P. Harms, S. Kiemle, D. Dietrich, *Extensible Earth Observation Data Catalogues with multiple Interfaces*, PV2007

[2] K. Dengler, T. Heinen, A. Huber, K. Molch, E. Mikusch, *The EOC Geoservice: Standardized Access to Earth Observation Data Sets and Value Added Products* , PV2013

[3] C. Reck, K. Molch, E. Mikusch, J. Hoffmann, *Copernicus Data Access and Exploitation Collaborative Infrastructure – COPACI System Concept*, DLR 2014

[4] *OGC Earth Observation Metadata profile of Observations & Measurements (EOP O&M)*, OGC 10-157r3, Version 1.0.0, 2012-06-12

[5] European Space Agency (ESA), (2012) *Heterogeneous Missions Accessibility (HMA)*, ISBN 978-92-9221-883-6, http://esamultimedia.esa.int/multimedia/publications/TM-21/TM-21.pdf

THE TECHNOLOGY AND ATMOSPHERIC MISSION PLATFORM (TAMP)

Stefano Natali[(1)], Simone Mantovani[(1)], Gerhard Triebnig[(2)], Daniel Santillan[(2)], Marcus Hirtl[(3)], Thorsten Fehr[(4)]

[(1)] SISTEMA GmbH, Vienna, Austria
[(2)] EOX IT Services, Vienna, Austria
[(3)] Zentralanstalt für Meteorologie und Geodynamik (ZAMG), Vienna, Austria
[(4)] ESA ESTEC, Noordwijk, The Netherlands

ABSTRACT

The scientific and industrial communities are being confronted with a strong increase of satellite missions and related data. This is in particular the case for the Atmospheric Sciences communities, with the upcoming Copernicus Sentinel-5 Precursor, Sentinel-4, -5 and -3, and ESA's Earth Explorers scientific satellites ADM-Aeolus and EarthCARE. The challenge is not only to manage the large volume of data generated by each mission / sensor, but to process and analyze the data streams. Creating synergies among the different datasets will be key to exploit the full potential of the available information. Integrating Earth-Observation data with ground based observations and numerical models, is the basis for a new data exploitation paradigm which opens new research and commercial opportunities.

As a preparation activity supporting scientific data exploitation for Earth Explorer and Sentinel atmospheric missions, ESA is funding the technology study and prototype implementation of the "Technology and Atmospheric Mission Platform" (TAMP). The services and tools are developed along use cases defined with users from different scientific and operational fields and implemented according to their requirements to ensure acceptance of TAMP platform by the atmospheric community.

The current work provides summary information about the collected requirements, the TAMP system design and implementation.

Index Terms— Atmospheric Applications, Big Data, Collaborative e-Infrastructure, Data Visualization, Virtual Research Environment

1. INTRODUCTION

New generations of satellite-borne sensors (e.g. Atmospheric Sentinels, Earth Explorer missions) will provide an unprecedented amount and variety of data to be used by the Atmospheric Sciences community and to complement / to be assimilated in numerical weather predictions. As in many other fields, Atmospheric Sciences are experiencing / facing a new challenge: moving from the observational science, to the theoretical science, then to the simulation science, reaching the intensive / massive data exploitation science (.e.g see the Fourth Paradigm, [1]).

Atmospheric Sciences will get a huge benefit from the combined and effective use of a large variety and volume of data that completely fit with the concept of big data (see the 4Vs of Big Data, [2]). As example, Hirtl et al. demonstrate the improved performance in forecasting particulate matter concentration assimilating satellite-based Aerosol Optical Thickness (AOT) maps into numerical models [3].

In order to be ready to exploit upcoming satellite datasets, there is a need to set up and demonstrate the feasibility of a new approach, enriching as much as possible the Atmospheric Sciences scenario.

The role of the potential system users is crucial in the development of a new platform: users can define needs and potential limitation of the system, and its usefulness in their daily work. Furthermore there is a need to establish new collaboration modalities and tools among users: the concept of virtual laboratory shall be implemented, promoting large centralized infrastructure investments than local small ones,

Technological challenges to be faced for the implementation of the TAMP platform span from the big data management related to multi-source and multi-dimensional datasets (see e.g. [8], [9], [10], [11], [12]), to multi-dimensional data visualization needs (see, as example, the V-MANIP project portal [7]), to remote and massive data processing, e.g. managed via message brokers and task dispatched technologies (see, as example, [12], [13]).

In the current work, the development of the TAMP platform is described, highlighting on the one hand the fundamental role of users in the definition of the platform requirements and in its validation, on the other the strong technological effort of the project partners in using cutting-edge technologies to implement the demonstrator.

2. REQUIREMENTS DEFINITON

The definition of the system requirements represents the fundament on which the platform is built. Requirements have

been defined following user, data, and technological needs and constraints. A set of high level European scientific groups (hereafter referred as Scientific and Technical Forum, STF) have been selected to define the user requirements to be successively translated into system requirements. The STF members belong the European Centre for Medium-range Weather Forecast (ECMWF, UK), the German Aerospace Center (DLR, Germany), the Belgian Institute for Space Aeronomy (BIRA, Belgium), the Royal Netherlands Meteorological Institute (KNMI, The Netherlands), the National Institute of Research and Development for Optoelectronics (INOE, Romania), the Austrian Meteorological Service (ZAMG, Austria).

The STF members have been requested to describe a potential use of the TAMP platform defining:

- Which data to use (satellite-based data, numerical models, validation datasets
- How these products shall be prepared to be used for their specific scopes (pre-processing activities such as subsetting, remapping, …)
- How these products can be used together (e.g. comparison between two products providing the same field, combined processing, …)
- How these products / processing results could be effectively visualized and manipulated via advanced multi-dimensional graphic user interfaces (GUI)
- How the results of the analysis / processing shall be exported or directly published onto specific media.

As result of the consultation, a set of 272 unique user requirements have been collected, to be translated into system requirements. The full traceability of the technical (system and modules) requirements has been achieved by means of cross reference matrices to ensure that any change in user or technical requirements can be correctly addressed. Finally eight use cases have been selected to be supported by the implemented software to validate the system performance.

3. SYSTEM DESIGN

The implemented TAMP platform will be deployed on a remote infrastructure providing access via web browser as well as via direct machine connection (e.g. secure shell connection). The platform shall store a large amount of data and shall make available computational resources to perform massive data processing.

Three main logical layers are defined to represent the TAMP platform:

- The Graphic User Interface GUI) layer, containing:
  - The Public Information Portal (PIP): it shows the project main features, providing a social-like page where posts from the system administrator (e.g. availability of new datasets) and users (e.g. specific results achieved via the platform) are available
  - A user page, from which each user can configure preference, post news and access to the Data Analysis and Visualization Environment (DAVE)
  - A data ingestion page, from which a user can upload owned data to be used within the platform
  - The Data Analysis and Visualization Environment (DAVE), from which the user can access to the data, visualize, select, run data assessment tools and save the processing results.
- The server-side layer, that contains five main modules:
  - The processing libraries, a set of data access and processing libraries to support processing modules
  - The data ingestion module, to pre-process and store system retrieved data as well as user provided data
  - The Web Coverage Service (WCS) data access module, to allow internal and external modules to retrieve the stored and processed data via standardized interfaces
  - Data processing modules, where data assessment and processing modules are deployed
  - User virtual machines: these are computational resources available for the users, from which the users can access to all system-available data, deploy owned processors, run the processors on system data and visualize the results via the DAVE GUI.
- The data storage layer, that contains both the system collected or user provided data and the processing results, in two separated areas. Both storage areas can be accessed via the WCS module.

4. SYSTEM IMPLEMENTATION

To pursue the virtual laboratory approach, a dedicated infrastructure to host the TAMP platform has been made available by the Austrian Meteorological Service: 64 processors, 128 GB RAM and 5TB storage are shared amongst the system modules and user virtual machines.

Data to be collected are summarized as follows:

- Satellite-based Ozone products (total column, profile) from SCIAMACHY, MIPAS, GOMOS, GOME, SCISAT, ODIN
- Satellite-based Aerosol products (total column, profile) from SCIAMACHY, MIPAS, GOMOS, GOME, OMI, MODIS, CALIPSO
- Satellite-based NO2 products (total column, profile) from SCIAMACHY, MIPAS, GOMOS, GOME, OMI
- Satellite-based CH4 products from MIPAS

- Satellite-based SO2 products from OMI
- Numerical weather prediction models data coupled with chemistry (MACC model data, BASCOE Reanalysis, WRF/CHEM simulations) low resolution/ global coverage and high resolution/ regional coverage
- Correlative data from AERONET, EARLINET, CLOUDNET, PANDONIA, EMEP emission inventory, TNO emission inventory, ESA EVDC, NASA AVDC, ACTRIS (for SO2)

The platform makes available the following functionalities for data (pre-)processing and data assessment:

- Pre-processing and processing tools
  - Remap data onto a pre-defined geographic grid
  - Additions / subtractions of data
  - Units Conversion
  - Vertical integration
  - Combined spatial-temporal means
  - Extraction/ use of quality flags
  - Filtering of level 2 data according to quality flag
  - Chemical speciation monitor
  - Level 3 product generations including averaging kernel (by users)
- Data Assessment
  - Compute basic statistical tests (subset of the DELTA A&P benchmarking tool)

The users are allowed to update and run software packages and scripts in the following languages:

- Fortran (90)
- C / C++
- Python
- Shell scripts

Data can be extracted from the platform in netCDF(-CF) format; plots can be exported in png / ps / pdf format.

The TAMP platform is, at this stage, a demonstration platform freely accessible by STF members and associated scientists of the same institutions. Nevertheless its design and implementation technologies allow a quick system upscale and transfer to operation to accommodate requests from a wider community. The future evolution of the platform will be defined based on the STF members' feedback and ESA follow-on strategies.

5. RESULTS AND CONCLUSIONS

The upcoming scenario of data availability is driving the development of new concepts of research environments, where large data storage and processing capabilities are available to a wide range of users that can jointly develop and validate new algorithms. The TAMP project aims at demonstrating the approach for Atmospheric Sciences users, where a strong increase of satellite-based products is expected in the near future, and where multi-sensor, multi-temporal and multi-dimensional data can be used in synergy to improve current and future products, paving the way for a future data assessment operational platform.

6. AKNOWLEDGEMENTS

The current study is funded by the European Space Agency under the contract: Technology and Atmospheric Mission Platform (C.n. 4000112832/14/I-CS).

7. REFERENCES

[1] Tony Hey, Stewart Tansley, and Kristin Tolle (Eds.). The Fourth Paradigm: Data-Intensive Scientific Discovery, Microsoft Research, Redmond, Washington, (2009).

[2] http://en.wikipedia.org/wiki/Big_data Big Data Wikipedia Page (visited on 06.11.2015).

[3] Hirtl, M., Mantovani, S., Krüger, B. C., Triebnig, G., Flandorfer, C., Bottoni, M., & Cavicchi, M.: Improvement of air quality forecasts with satellite and ground based particulate matter observations. Atmospheric Environment, 84, 20-27 (2014)

[4] Sentinel-3 Team GMES-S3OP-EOPG-TN-13-0001 Sentinel-3 User Handbook Issue: 1.0'Date: 2 Sept. 2013

[5] Paul Ingmann, B. Veihelmann, J. Langen, D. Lamarre, H. Stark, G. BazalgetteCourrèges-Lacoste: Requirements for the GMES Atmosphere Service and ESA's implementation concept: Sentinels-4/-5 and -5p. Remote Sensing of Envi-ronment 120 (2012) 58–69.

[6] Various Authors: ADM-Aeolus Science Report. European Space Agency publication, SP-1311, 'ADM-Aeolus' April 2008, ISBN 978-92-9221-404-3, ISSN 0379-6566

[7] http://demo.v-manip.eox.at/ Interactive geospatial explorations (V-MANIP Demo Applications) (visited on 12.02.2016).

[8] http://www.esa.int/esaLP/LPearthcare.html EarthCARE Web Portal (visited on 12.02.2016).

[9] http://www.esa-aerosol-cci.org/ ESA CCI Aerosol web site (visited on 12.02.2016).

[10] http://www.esa-cloud-cci.org/ ESA CCI Cloud web site (visited on 12.02.2016)

[11] http://www.esa-ozone-cci.org/ ESA CCI ozone web site (visited on 12.02.2016).

[12] http://www.esa-ghg-cci.org/ ESA CCI greenhouse gases web site (visited on 12.02.2016).

[13] http://www.celeryproject.org/ Celery home page (visited on 12.02.2016).

[14] http://www.rabbitmq.com/ RabbitMQ home page (visited on 12.02.2016)

EXPLOITATION PLATFORMS OPEN ARCHITECTURE

Salvatore Pinto (1), Sveinung Loekken (2)

(1) RHEA Group S.A., (2) ESA

ABSTRACT

The present ESA support to EO data exploitation, driven by the evolution of technology and corresponding shifts in user expectations, is currently seeing the emergence of a new and complementary operations concept, based on the availability of collaborative Exploitation Platforms.

Focusing on the ESA experience with the Thematic Exploitation Platforms projects and their precursors, ESA is currently defining the concept of an Exploitation Platform Open Architecture, which includes the definition of the architecture key components, interfaces and related interface standards. The aim of this activity is to provide harmonization, guidance and recommendations to the current and future Exploitation Platforms projects.

In this paper, we will summarize the status of this work, discussing the current Exploitation Platform Open Architecture draft, the status of the standardization of the Exploitation Platform interfaces and the plans for the evolution of the Exploitation Platform harmonization activity.

***Index Terms*—** Exploitation Platforms, Architecture

1. INTRODUCTION

An Exploitation Platform (EP) is a virtual open and collaborative environment, which brings together EO and non-EO data, computing resources, hosted processing, collaborative tools (processing tools, data mining tools, user tools, …), development tools, test bench functions, application shops, market place functionalities, communication tools (social network), documentation, accounting framework and operations tools to manage resources.

An EP targets a wide set of users, here divided in three major classes:

- Final User (FU), consumer of the platform processing services and/or products stored and generated via the platform.
- Scientific user or service provider, here referred as Principal Investigator (PI), performing development and integration of algorithms within the platform, to serve the FU.
- Data Provider (DP), providing input data products (EO or non-EO) to the platform and owning the IPR on the data.

In addition to these users, the EP includes Operators (Ops) users, responsible for the operational tasks on the platform.

An Exploitation Platform can provide different type of processing services to the users. In the EP architecture and this paper, these are divided into Apps and Workflows. An App represents an interactive applications, managed completely by the user via a dedicated GUI, while a Workflow represents a processing based on pre-defined input and output data types, pre-defined processing steps and orchestration logic, which is controlled by the user via the selection of processing parameters and input files.

This paper summaries the content of the Exploitation Platform Open Architecture document [1], released under Creative Commons Attribution-ShareAlike license in initial draft version. All the information contained here can be then accessed in a more detailed form into the document.

2. EXPLOITATION PLATFORM KEY FUNCTIONALITIES

An Exploitation Platform is a complex environment, where several actors are involved with different roles and several functionalities are offered to the users, tailored to their specific roles. The key functionalities of an EP can be summarized as the following:

- “Data discovery”, which allows the FU to search for data products (input or published products) based on data metadata and custom search rules.
- “Data management”, for the DP to upload data into the platform, edit and delete it, setup data metadata, setup rights for data access and other data management related tasks.
- “On-demand processing”, which allows the FU to run processing services on-demand in the form of a Workflow or an App.
- “Massive processing”, which allows the FU to perform bulk processing over large datasets or systematic processing upon new data acquisitions.
- “Development and integration system”, for the PI to integrate his software into the platform as a new Workflow or Apps or to update existing versions of the software.
- “Collaboration tools”, for the PI to share processing services or processing results, within the EP or within external communication tools (e.g. social networks).
- “Documentation and support tools” for the operator to provide support to the other users.

• "Services marketplace", for the FU to discover and compare all the platform services.

• "Operator interface", where the operator users can perform all the standard operation tasks.

3. EXPLOITATION PLATFORM ARCHITECTURE DESIGN PRINCIPLES

The current draft design of an Exploitation Platform Architecture [1] comes from ESA's experience acquired in several Exploitation Platform precursor projects, such as the SuperSites Exploitation Platform (SSEP) [2] and the Geohazards Thematic Exploitation Platform prototype (TEP-QW) [3], as well as other scientific exploitation processing environments, such as the Grid Processing On-Demand (G-POD) [4].

During the EP design, the following principles have been followed:

• Infrastructure independence, to allow the deployment of the EP within different ICT technologies, such as Cloud, HPC, HTC or bare-metal.

• Scalability, according to requests or manually performed by the operator, allowing easy addition or removal of ICT resources to optimize operational costs.

• Cost effectiveness, for optimizing resource usage, reducing the platform overhead and maximizing the utilization of the ICT resources.

• Interoperability, achieved via common interfaces standards, to share data, processing services, algorithms and ICT resources between different EPs.

• Distributed approach, with separation of roles of the key components (e.g. storage, processing, …), allowing multiple separate deployments of each components over a wide area network.

• Flexibility in service integration, allowing PI to integrate existing software applications with minimal adaptation regardless of their programming language or software dependencies constraints.

• Common authentication and authorization, integrated quota and accounting, with the possibility to extend the EP functionalities into a pay per use scenario or federate multiple EPs.

• Flexibility to adapt the user interface to support thematic applications and thematic user communities without the need to customize the back-end.

4. OPEN ARCHITECTURE

To give a general overview of the Exploitation Platform system, the Exploitation Platform Architecture is split among four macro-components.

A macro-component is a logical collection of components that have similar or strictly related functions

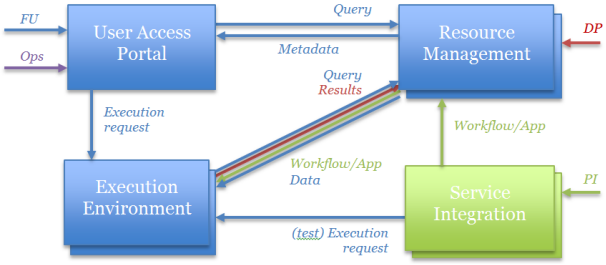

Figure 1. Macro-components of the EP Open Architecture

and implements a homogeneous set of Exploitation Platform services.

Each macro-component may exist separately from the others, offering its functionalities directly to the users. Also, a macro-component can be deployed in multiple instances, tailored to serve different communities and federated across different resource providers.

The four macro-components implemented in the architecture, with a basic view of their relations, are depicted in Figure 1 and briefly described in the following paragraphs.

At the end of this chapter, we will provide an overview of the key internal EP components. A detailed description of each macro-component and its internal components and interfaces is provided in [1].

4.1 User Access Portal macro-component

The User Access Portal macro-component provides the interface to the FUs and the Ops. In addition, this macro-component implements a set of common underlying services used by the other components, such as authentication, accounting, monitoring and collaboration tools.

The UAP macro-component connects to the Resource Management macro-component for retrieving information about data, App and Workflow resources required for a given processing service and to the Execution Environment macro-component for executing processing services.

Moreover, the UAP retrieves and visualize the processing results, permitting to share them among the community and within external communication networks (e.g. social networks).

Being the main interface to the FUs, the UAP macro-component is the one more related to the particular Exploitation Platforms project and it is usually heavily customized to fit the EP user community needs.

4.2 Resource Management macro-component

The resource management macro-component handles the resources available in the platform. With the generic term resource, it is considered both a processing service, which

can be in the form of a Workflow or an App, input Data or processing Results.

This component implements the data discovery and data management functionalities, together with the management of the processing services and processing results. The resource management macro-component takes care of the resource storage, resource location resolution, resource catalogue, resource metadata harvest, resource quick-look generation and all the other resource management tasks.

The resource management macro-component takes as input a Query, which is the request for discovery of a resource fulfilling different selection criteria, and returns the resource itself, in form of one or multiple Data products, processing Results, App or Workflow packages.

The resource management macro-component considers also as input a Publish request, which is the request, performed by the DP or the PI, to insert one or more resource into the system (e.g. new data or new Workflows or new Apps).

The resource management macro-component can be deployed in multiple instances within the platform and distributed over different sites. Moreover, an EP can connect to Resource Management components managed by different exploitation platforms, enabling the effective sharing of resources between EPs.

4.3 Service Integration macro-component

The Service Integration macro-component provides the PI with a framework to integrate his own application, algorithm and/or software into the platform as a new processing service (Workflow or App).

The PI uses this macro-component to describe the application logic and ultimately package the application into a Workflow or an App. This macro-component does not include processing functionalities, but, for testing and debug purposes, sends directly Execution Requests to the Execution Environment macro-component.

Output of this component is a Workflow or App package, constructed according to a given standard, which is published by the PI into the Resource Management macro-component. This allows the possibility to share processing algorithms within several Exploitation Platforms.

As per the Resource Management macro-component, this component can be distributed across different EP and different deployment sites. Moreover, it can be installed at the user premises, to allow local integration of the application without the need to consume platform resources.

4.4 Execution Environment macro-component

The Execution Environment macro-component provides an environment to run processing services, implementing on-demand processing and massive processing functionalities.

The Execution Environment macro-component takes as input an Execution Request, which can be for both a Workflow and an App processing.

In case of a Workflow processing, the macro-component will perform the pre-defined processing operations according to the specified processing parameters, without further interaction with the FU, then publish the Results into the Resource Management macro-component.

In case of a an App processing, the Execution Environment will interact with the FU during the entire processing via the UAP macro-component, for the FU to interactively perform the processing operations and the publication of the results.

This macro-component receives the Workflow, App and input Data required for the processing execution from the Resource Management macro-component. The Execution Environment macro-component can also perform Queries to the Resource Management to look for additional resources required by the processing (e.g. auxiliary data).

As per the Resource Management macro-component, this component can be deployed into multiple sites. This feature is very important for processing execution performances, allowing to deploy the processing clusters close to the data storage, thus minimizing the time to access the input datasets.

4.5 Internal components overview

The macro-components represent a logical collection of several EP internal components. Of the overall EP architecture, a set of components can be considered fundamental for the building of an EP and provide the EP key functionalities. These components are also related to the EP main interfaces for resources and processing services access. In summary, an EP should have, at least a:

- Catalogue, for discovering of platform resources (eg. data, results, processing services), according to metadata, classification, ranking, etc…
- Resource Access Gateway, for data, Workflow, App and processing Results access, download and usage accounting.
- Execution Gateway, for submission of processing requests of both Apps and Workflows.
- Workflow/App manager, which interprets and executes a Workflow or an App into the platform.
- Geo Resource Browser, which provides the possibility to visualize and share all resources involved in a processing (eg. data, results, processing services). This component is the key for ensuring collaboration, reproducibility and sharing of processing results.
- AAI, which manages the Authentication ad Authorization for the users of the platform.

Moreover, several other components are important in meeting the design principles of an EP, such as:

• Workspace, for the FU to submit and manage processing requests.

• Systematic Processing component, to orchestrate massive processing operations.

• DevBox to provide the PI with a framework for the integration of new services into the platform.

• Capacity Manager, to adapt dynamically the size of the processing cluster and ensure optimization of ICT usage.

• Resource Ingestion, to upload new data into the platform.

• Management console, for the Ops to manage the platform.

A detailed description of these components is reported in [1].

5. INTERFACES AND STANDARDS

Interoperability between different EPs is based on a set of standard machine-to-machine interfaces for the key components of the platform, which allow the sharing and federation of the main functionalities of the EP.

In particular, from the EP key components, the catalogue component is associated with the Resource Search interface, the resource access gateway with the Resource Access interface, the execution gateway with the Processing Execution interface, the Workflow/App manager with the Processing Service Package interface, the geo resource browser with the Processing Container interface and the AAI component with the AAI interface.

For each one of these interfaces, a selection process of the best open standard for a given interface is currently ongoing. When the existing standards do not meet the EP requirements, a proposal for future evolutions of existing standards is currently under definition.

The selection and standards evolution process follows a set of guidelines, including preference to web/developer oriented API standards, to enable seamless integration and simple client implementation. The guidelines translates then in a set of recommendations, such as, for example, preference towards JSON over XML, REST over SOAP, pre-defined namespace vs custom namespaces, standards which clear associated implementation best practices.

These recommendations, together with the particular interface requirements coming from the EP design, drive the selection standards selection and evolution processes.

In the current EP architecture draft, the following baseline standards are proposed for each one of the key components main interfaces:

• Resource Search interface: OGC OpenSearch standards [5] and BPs [6]

• Resource Access: OGC Download Services for EO [7]

• Processing Service Execution: OGC WPS [8]

• Processing Service Packaging: new proposal, based on XPDL/Docker

• Processing Container: OGC OWS Context [9]

• AAI: OGC User Management Interfaces for EO [10]

6. FUTURE ACTIVITIES

The EP Open Architecture and related interfaces described in this paper are the initial result of an on-going harmonization activities within the ESA Exploitation Platform projects. In parallel to this activity, ESA is currently developing separate Thematic Exploitation Platforms and Mission Exploitation Platforms, respectively tailored to the need of particular user communities and a particular satellite mission. These projects have parallel architecture design activities, with the aim to explore multiple different solutions in relation to the EP users community needs.

Thus, the current EP Open Architecture draft, originally provided as guideline to the TEPs, will be evolved based on knowledge and lessons learnt acquired from the TEPs projects parallel design.

The TEP projects will also contribute, together with other ESA and non-ESA activities related to Identity Management and standardization, such as HMA and OGC test-beds, to the final definition and standardization of the Exploitation Platform interfaces.

Moreover, the current harmonization activities plan includes, always with the help of the TEP projects, the definition of recommendations for the implementation of the EP architecture core components, based on existing Open Source implementations.

7. REFERENCES

[1] S. Pinto, "Exploitation Platform Open Architecture", DRAFT2, https://tep.eo.esa.int/webdav/guest/document_library/Public/Exploitation%20Platforms%20Open%20Architecture%20DRAFT2.

[2] ESA SuperSites Exploitation Platform, http://ssep.eo.esa.int

[3] ESA GeoHazards Exploitation Platform precursor, http://geohazards-tep.eo.esa.int

[4] ESA Grid Processing On-Demand, http://gpod.eo.esa.int

[5] P. Gonçalves, OGC® OpenSearch Geo and Time Extensions, 10-032r8

[6] Y. Kudo, J. Gasperi, CEOS OpenSearch Best Practices Document, v1.0.1, CEOS-OPENSEARCH-BP-V1.0.1

[7] D. Marchionni, R. Cafini, OGC Download Service for Earth Observation Products Best Practice, 13-043

[8] P. Schut, OGC Web Processing Service, 05-007r7

[9] R. Bakin, P. Gonçalves, OGC OWS Context Atom Encoding Standard, 12-084r2

[10] P. Denis, P. Jacques, OGC User Management Interfaces for Earth Observation Services, 07-118r9

FACILITATING COOPERATIVE RESEARCH, DEVELOPMENT AND OPERATIONS IN EARTH OBSERVATION

Christian Briese[1], Werner Mücke[1] & Michael Aspetsberger[2], Wolfgang Wagner[3,1]
[1]EODC Earth Observation Data Centre for Water Resources Monitoring GmbH, Vienna, Austria
[2]Catalysts GmbH, Linz, Austria
[3] Department of Geodesy and Geoinformation, Vienna University of Technology, Vienna, Austria

ABSTRACT

The Sentinels generate huge amounts of data with a high spatio-temporal resolution. This leads to new challenges in research and in the processing and archiving of big EO data. Based on the EODC concept this contribution provides an insight in the practical realisation of the ideas within the EODC framework. Next to the aspect of collaboration it focuses on the information exploitation within EODC and the distribution of results to the users. Finally, in order to fully exploit the richness of EO data, the necessary federation of data centres is discussed.

Index Terms— Big EO data, processing, collaboration, virtual research environment, federation of data centres

1. CHANGING THE PARADIGM

With the advent of the Copernicus space programme and its novel Sentinel satellite missions, both the temporal and spatial resolutions of the produced Earth Observation (EO) data are dramatically increasing, which consequently leads to Big Data volumes and unprecedented challenges in their exploitation. In the medium-term, a paradigm change in the EO landscape is to be expected: there is not going to be a single data centre for processing, storing and distributing of EO data, instead we are moving towards a federation of multiple interconnected data centres spreading and handling the workload among each other. Additionally, due to the inherently increasing complexity of observed phenomena and retrieval algorithms, the association of interdisciplinary working groups will be advantageous. Furthermore, contrary to current practice, the methods, algorithms and software will have to be transferred to where the data are located, in order to avoid bottlenecks due to inadequate data transfer capabilities. This contribution introduces the EODC concept and focuses on its realisation. EODC aims at addressing the above mentioned challenges, thereby facilitating cooperation between science, industry and public sectors, giving the users the means to access Sentinel data and work with them in an efficient and economic way.

2. RESEARCH, DEVELOPMENT AND OPERATIONS ON THE EODC PLATFORM

2.1. Introducing the basic concept

EODC builds its IT capacities on three basic pillars (see Figure 1): (1) NORA, which stands for "Near real-time operations and rolling archive", (2) SIDP, the "Science Integration and Development Platform", and (3) GTR, the "Global testing and reprocessing facility". NORA offers four services: data input from external satellite data archives, output of these data to EODC storage, status monitoring of system performance and data transfer, as well as near real-time processing of selected products. SIDP is EODC's fully equipped and flexible cloud infrastructure, where tailor-made pre-configured virtual machines and other services (e.g. continuous integration and deployment) are being hosted, supporting the remote development and testing of methods and source code. GTR is a top500-ranking supercomputer, namely the Vienna Scientific Cluster 3 (VSC-3), intended for large scale processing. These three components are complemented by a Petabyte-scale data archive (i.e. the EODC data pool), which is physically co-located with SIDP and GTR and connected via InfiniBand network technology to minimise transfer times and increase I/O.

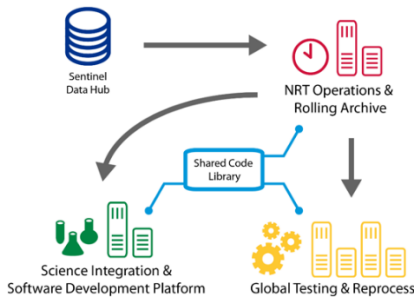

Figure 1: The three pillars of the EODC IT concept. [1]

2.2 Connecting the community

Within the EODC, individuals join together into communities sharing similar foci or research and development goals. The idea is for the single researcher or organisation to take advantage of the experience and skills of all involved partners and to participate in collaborative development processes, while at the same time gaining increased external visibility via a larger group. In order to support community building, EODC organises regular community events for information exchange. In an effort to pro-actively stimulate cooperation, several tools are

implemented on the EODC IT platform, such as a shared code library (based on GitLab) or a knowledge base (similar to wiki).

2.3. Accessing the data pool

EODC has been acquiring Sentinel-1 (S-1) from day one and is providing them through its long-term data archive. On average, the data are available in the archive approximately 2.5 hours after their processing time and 6.25 hours after acquisition time. Currently, a total of 66556 S-1 acquisitions are hosted in the data pool (status October 2015). Once a user is logged on to the EODC network, he may access the data either through SIDP or GTR. Search and discovery are provided via a meta database that supports spatio-temporal search functions and scripting. Conventional access is established by means of file system access and http-export. Evaluations are ongoing to provide OGC WM(T)S and WCS interfaces for selected products.

2.4. Exploiting the information

The workflow from raw or pre-processed to refined and value-added data is as follows: Starting on SIDP, users or communities develop their methods, while testing it on small data sets (typically hundreds of Mega- and up to few Gigabytes). As soon as their code base has reached a certain state of maturity, they transfer to GTR, where the same codebase is available for testing on larger areas, usually regional or even global scale. In order to exploit the computational capacities of the super computer to its full potential, programmers are advised to either provide parallelized code or achieve parallelization through the respective application, e.g. parallel processing of many data files at once. If a near real-time product is targeted, the developers move to NORA for the establishment of an operational NRT service. While working with SIDP and NORA is basically similar to working on personal computers, as both offer the possibility to host user-definable virtual machines, accessing the supercomputer GTR is a completely different environment. It involves a minimum level of understanding for high performance computing, parallelisation and certain programming experience. To create an environment that supports a variety of users with differing skill levels, EODC offers a simplified processing submission and task scheduling interface (see Figure 2). A browser-based graphical user interface allows the users to select pre-defined processing chains from the EODC code library, use their own configuration files to define settings for included algorithms, and select the number of computing nodes they would like to employ. After the processing is triggered by the user, everything else is done automatically: the respective codebase is pulled from the library and together with the configuration files processing containers are built, representing sand-boxed environments including all software packages necessary for successful computation. During data processing on GTR, which is handled by the open source workload manager SLURM, the browser-based job monitor gives access to principle information of the tasks, such as estimated run time. On job completion, the users are notified and may inspect their results using the available data viewer.

2.5. Distributing the results

Value-added data produced on the EODC platform may be stored in the archive for later usage, and additionally published or redistributed. The EODC interactive delivery platform provides multi-channel online access and basic analysis features for various kinds of EO data and is intended to offer users tailor-made solutions to deliver their information products to a wider audience.

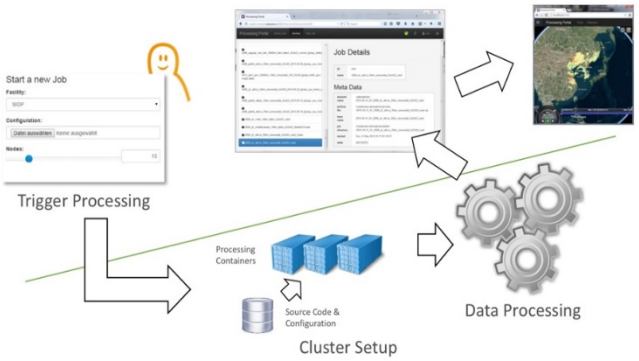

Figure 2: Illustration of simplified scheduling of processing jobs on the VSC-3 supercomputer.

3. OUTLOOK

The ever-increasing amounts of EO data to be expected in the upcoming years demand for a sophisticated expansion strategy for computational power, as well as for storage space. EODC is planning to receive and provide Sentinel 1, 2 and 3 data and therefore aiming to significantly extend its storage. In parallel, the virtual research environment SIDP is being equipped to host 100+ users in the same time frame. Given the diversity of EO products in Europe, and the data centres producing them, EODC is strongly working towards a federation of data centres or other similar initiatives across Europe, in order to be able to exploit the information richness of up-to-date EO data to its full potential.

4. REFERENCES

[1] Wagner, W. et al., 2014: Addressing grand challenges in earth observation science: The Earth Observation Data Centre for Water Resources Monitoring. In: ISPRS Technical Commission VII Mid-term Symposium 2014, Volume II-7, Istanbul, Turkey, pp. 81-88.

TOWARDS A JRC EARTH OBSERVATION DATA AND PROCESSING PLATFORM

P. Soille[1], A. Burger[2], D. Rodriguez[1], V. Syrris[1], and V. Vasilev[2]

European Commission, Joint Research Centre (JRC)
[1]Institute for the Protection and Security of the Citizens, Global Security and Crisis Management Unit
[2]Institute for Environment and Sustainability, Digital Earth and Reference Data Unit

ABSTRACT

The Copernicus programme of the European Union with its fleet of Sentinel satellites operated by the European Space Agency are effectively making Earth Observation (EO) entering the big data era. Consequently, most application projects at continental or global scale cannot be addressed with conventional techniques. That is, the EO data revolution brought in by Copernicus needs to be matched by a processing revolution. Existing approaches such as those based on the processing of massive archives of Landsat data are reviewed and the concept of the Joint Research Centre Earth Observation Data and Processing platform is briefly presented.

***Index Terms*—** Earth Observation, Sentinel, Copernicus, Infrastructure

1. INTRODUCTION

To date, the United States (U.S.) Government is the largest provider of environmental and Earth system data in the world[1]. A first data revolution happened in 2008 when the U.S. Geological Survey decided to release for free to the public its Landsat archive which is the worlds largest collection of Earth imagery [11]. Still, the European Commission, with its ambitious Copernicus programme and associated Sentinel missions (S1 to S6 satellite series) operated by the European Space Agency and complemented by a range of contributing missions, is on the way to become the main provider of global EO data with a free, full, and open access data policy. With expected data *volumes* of 10 TB per day (when *all* Sentinel series will reach full operational capacity), data *velocity* highlighted by the production of global coverage with repeat time as short as 2 days for Sentinel-3, and data *variety* resulting from sensors in the optical and radar ranges at various spatial, spectral, and temporal resolutions, the Copernicus programme is a game changer making EO data effectively entering the big data era [10]. Figure 1 shows the overall estimated data throughput for the Sentinel 1–3 missions compared to those delivered by the Landsat 8/MODIS satellites.

[1]http://dels.nas.edu/resources/static-assets/besr/miscellaneous/Stryker.pdf

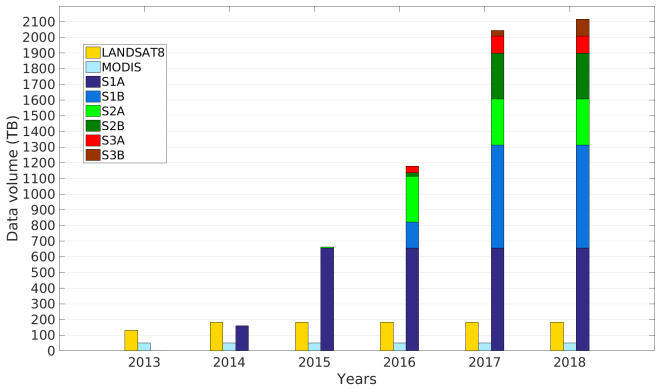

Fig. 1. Yearly data flow estimates from Sentinel 1–3 (assuming full operational capacity) compared to MODIS and Landsat 8 data flows.

Whilst the European Union (EU) is making EO entering the big data era in terms of data production, innovative developments need to be pursued to fully exploit the potential of the generated data whether for academic, institutional, or commercial applications. This also applies to the Joint Research Centre where the current fragmented approach of EO data storage and processing is no longer sustainable.

2. THE SENTINELS AND THE BIG EO DATA ERA

The evolution of the cumulative data produced by the Landsat missions and Sentinel 1-2-3 with estimations until summer 2018 is shown in Fig. 2. The underlying calculations are based on the following assumptions and considerations:

1. The data volume of Landsat 1-6 missions is $\sim$120 TB[2];
2. The data volume of Landsat 7 and 8 is estimated from the relating metadata files provided by USGS[3];
3. MODIS Terra and Aqua generate 70GB/day each [7];

[2]http://academic.emporia.edu/aberjame/remote/landsat/landsat.htm
[3]http://landsat.usgs.gov/metadatalist.php

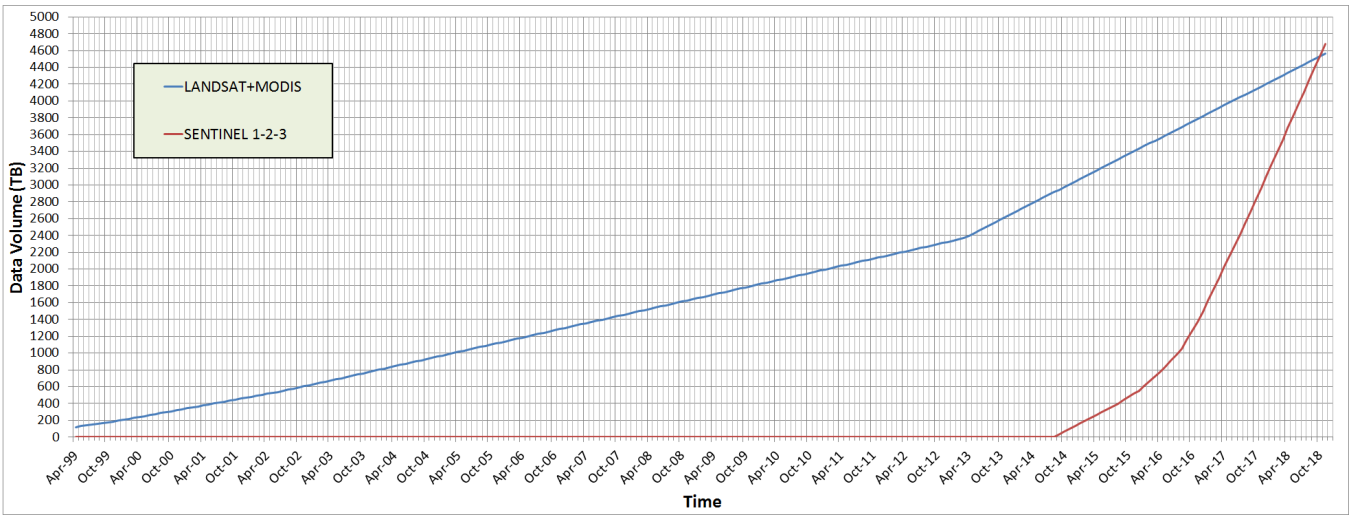

Fig. 2. Data volume: evolution of the data produced by the MODIS/Landsat missions and Sentinel 1-2-3 (estimations based on data throughput in full operational capacity after six months of the scheduled launch dates).

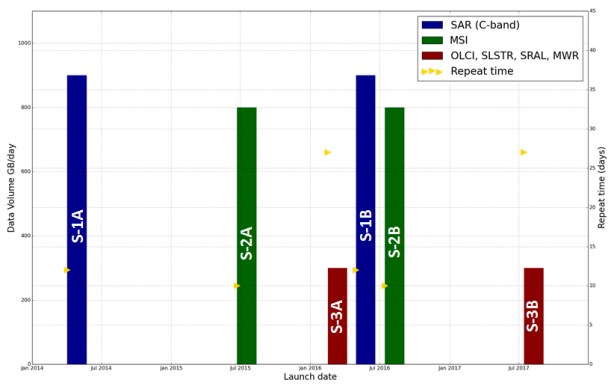

Fig. 3. Data velocity of Sentinel 1–3: daily data throughput in full operational capacity and orbit repeat time (in days).

4. The data volume of Sentinel 1, 2 and 3 in their twin constellations is approximately 3.6 TB/day, 1.6 TB/day and 0.6 TB/day respectively[4]; all starting dates are considered to be around six months after the respective mission launch (assuming that data are generated in full operational capacity by that date).

These estimations indicate that the cumulated data generated by the Sentinel 1–3 missions will exceed the data generated by all Landsat missions during 2018.

The data velocity daily of S1–3/A-B is highlighted by their data throughput and orbit repeat times, see Fig. 3. Finally, data variety of the Sentinel 1 to 3 satellite series is summarised in Table 1. Revisit times are not indicated since they increase with latitude.

[4] https://www.ffg.at/sites/default/files/esa_sentinel_core_products_overview.pdf

3. PLATFORMS FOR BIG EO DATA

This section presents a brief survey of some existing platforms and other initiatives to address the needs of big EO data. Given the available space for this short paper, it is by no means comprehensive. In particular, platforms mostly devoted to data dissemination are not considered here.

3.1. Examples from public sector

- NASA Earth Exchange (NEX)[5] is a platform for scientific collaboration, knowledge sharing and research in the Earth science community;
- ESA Exploitation Platforms (EP): Thematic Exploitation Platforms (TEPs), Earth Exploitation Platforms (EOPs) and Sentinel Application Platform (SNAP);
- DLR GeoFarm hardware organisation follows the cloud-like vitalisation of processing hardware, see also [4].
- The Theia Land Data Centre is a French national inter-agency organisation designed to foster the use of images issued from the space observation of land surfaces [6];
- Earth Observation Data Center (EODC)[6] public/private initiative [2] that combines HPC with data provision and collaborative development;
- Australian Geoscience Data Cube (AGDC) at National Computational Infrastructure in Canberra, Australia [5].

[5] https://nex.nasa.gov
[6] https://www.eodc.eu

Table 1. Sentinel 1 to 3 data variety: spectral, temporal, and spatial (repeat cycle/global coverage for 1 satellite).

Mission	Sensors	Applications	Repeat cycle/ Global coverage	Resolution	Formats
Sentinel-1	C-Band SAR	Monitoring: sea ice, oil spills, marine winds and waves, land-use change, respond emergencies such as floods and earthquakes	12 days/12 days	Strip map mode 80km swath 5x5m; Interferometric wide swath 240km 5x20m; Extra wide swath 400km 25x100m; Wave mode 20x20km at 5x20m	.SAFE with GEOTIFF, XML, PNG, XSD, HTML and netCDF files
Sentinel-2	MSI (13 bands from 443 nm to 2,190 nm)	Monitoring agriculture, forests, land-use change, land-cover change; mapping biophysical variables; monitoring coastal and inland waters; risk mapping and disaster mapping	10 days/10 days	10m, 20m, and 60m spatial resolution to identify spatial details consistent with 1ha minimum mapping unit)	.SAFE with JPG2000, XML, GML and HTML files
Sentinel-3	OLCI (21 bands from 400 nm to 1,020 nm), SLSTR (9 bands from 555 nm to 10,850 nm), SRAL, MWR	Monitoring sea (ocean circulation, tides), coastal zone, inland waters and land, ice and sea-ice, climate geodesy and geophysics, and land topography	27 days/3 days	OLCI 300m, SLSTR (500m for solar reflectance, 1km for thermal infrared bands)	.SAFE with GEOTIFF, XML, PNG, XSD, HTML and netCDF files

3.2. Examples from private sector

Several companies are hosting large amounts of EO data in combination with processing capabilities, e.g., Google Earth Engine (GEE)[7] for a web-based platform with dedicated web API, Amazon Web Services (AWS) with availability of Sentinel-2[8] and Landsat-8[9], and CloudEO[10] that proposes a geo-infrastructure as a service.

3.3. Big EO data initiatives

In parallel to the development of public and private platforms, a number of governmental and research initiatives aiming at addressing the needs of big geospatial data are flourishing; see for example:

- BEDI is a U.S.A. government big data initiative on civil Earth observation[11];
- The Big SkyEarth EU COST action[12];
- The EarthServer: http://www.earthserver.eu/ see also [1];
- *EarthCube* is a joint initiative between the Division of Advanced Cyberinfrastructure (ACI) and the Geosciences Directorate (GEO) of the US National Science Foundation (NSF).
- The Earth System Grid Federation: An Open Infrastructure for Access to Distributed Geospatial Data;
- NOAA's Big Data Project within which a set of Data Alliances are being formed with providers of IaaS.

[7] https://earthengine.google.com
[8] http://sentinel-pds.s3-website.eu-central-1.amazonaws.com
[9] http://aws.amazon.com/public-data-sets/landsat
[10] http://www.cloudeo-ag.com
[11] http://www.whitehouse.gov/sites/default/files/microsites/ostp/NSTC/national_plan_for_civil_earth_observations_-_july_2014.pdf
[12] http://www.cost.eu/COST_Actions/TDP/Actions/TD1403

4. JRC EO DATA AND PROCESSING PLATFORM

A number of JRC projects are exploiting Earth Observation data to achieve their goals. In the Sentinel era, this can only be addressed by an integrated approach combining data storage and data processing. This leads to the concept of JRC EO Data and Processing Platform (JEODPP) developed in the framework of the JRC EO&Social Sensing Big Data (EO&SS@BD) pilot project.

The envisaged architecture consists of processing servers accessing the data provided by a series of storage servers and their directly attached storage (Just a Bunch of Disks or JBODs) in a distributed file system environment. The I/O bottleneck typically observed with network attached storage is avoided by considering appropriate high speed server intercommunication topology (switched fabric in fibre channel). This topology has the best scalability compared to arbitrated loop and point-to-point alternatives. Storage servers are automatically populated with the data requested by the applications. For example, the automatic download of Sentinel-2A data is achieved by using a time-based job scheduler launching OpenSearch and OpenData (ODat) scripts taking into account user requirements (geographical areas, cloud coverage, seasonality, etc.).

Processing can be performed at various levels through a sandbox environment. In its simplest utilisation, users have direct access to a batch job scheduler allowing the automatic and distributed processing of EO data at continental or global scale. For example, the detection of clouds on the full S2A

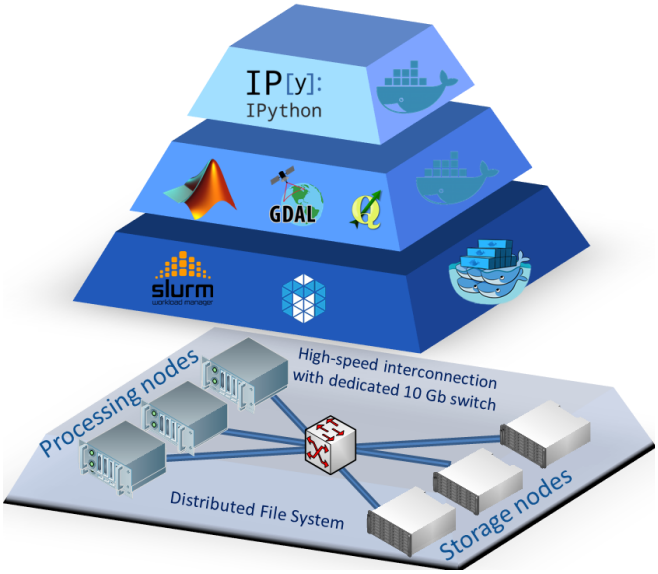

Fig. 4. Components of the planned JRC EO data and processing platform (JEODPP).

expert user data (about 30 TB of JPEG2000 compressed imagery) in only 5 days on a computing cluster with 10 processing nodes. Similarly, the automatic mosaicing of 2.5m SPOT (Copernicus CORE_3 data set) covering the whole territory of the EU plus 11 additional states (European Environment Agency member and associate members) [9] as well as the computation of Global Human Settlement Layers [8] from 4 multitemporal global Landsat data sets are achieved within one week.

The sandbox will offer different levels of user interaction with the platform:

- direct access to the batch job scheduler (slurm or HTCondor);
- virtualisation/operating-system-level virtualisation of the desired environment for prototyping (based on Docker Linux containers);
- interactive EO visualisation/processing capabilities;
- interactive data science and scientific computing through Jupyter web application (IPython notebooks), see also [3].

The various components of the planned EO data and processing platform are sketched in Figure 4.

5. CONCLUDING REMARKS

The EU Copernicus programme with its series of Sentinel missions acts as a game changer by bringing EO into the big data era. The value of the produced data depends on our capacity to extract information from it. The velocity, variety, and volume of the generated data combined with the need for using other non EO data sources calls for innovative approaches in data storage and processing. The scope of the JRC EO&Social Sensing Big Data (EO&SS@BD) pilot project is to propose innovative solutions addressing the needs of JRC projects. We are currently testing and optimising the various components of the JEODPP to ensure its scalability.

6. REFERENCES

[1] P. Baumann et al. Big data analytics for earth sciences: the earthserver approach. *International Journal of Digital Earth*, 1–27, 2015. doi: 10.1080/17538947.2014.1003106.

[2] C. Briese, et al. Challenges in the exploitation of big earth observation data. In *Proc. of BiDS'14*, 49–52. 2014. doi: 10.2788/1823.

[3] A. Burger et al. Towards an infrastucture for interactive Earth Observation data analysis and processing. In B. Chan, editor, *Cloud Services for Synchronisation and Sharing*, p. 27. ETHZ, 2016. doi: 10.5281/zenodo.44783.

[4] S. Kiemle, et al. Big data management in earth observation - the German satellite data archive at DLR. In *Proc. of BiDS'14*, pages 45–49. doi: 10.2788/1823.

[5] A. Lewis, et al. Iterating Petabyte-Scale Earth Observation Processes in The Australian Geoscience Data Cube. In *Proc. of BiDS'14*. doi: 10.2788/1823.

[6] C. L'Helguen et al. Muscate: Multi-satellites, multi-sensors and multi-temporal theia data centre. In *Proc. of BiDS'14*, 114–118. doi: 10.2788/1823.

[7] E.J. Masuoka, et al. Key characteristics of MODIS data products. *IEEE Transactions on Geoscience and Remote Sensing*, 36(4):1313–1323, 1998. doi: 10.1109/36.701081.

[8] M. Pesaresi, et al. Operating procedure for the production of the global human settlement layer from Landsat data of the epochs 1975, 1990, 2000, and 2014. Technical Report, JRC, 2015. doi: 10.2788/253582.

[9] P. Soille. Seamless mosaicing of very high resolution satellite data at continental scale. In P. Soille et al., editors, *Proc. of BiDS'14*, 222–223. doi: 10.2788/1823.

[10] P. Soille et al. Preface. In *Proc. of the BiDS'14* doi: 10.2788/1823.

[11] The Landsat Science Team. Free access to Landsat imagery. *Science*, 320(5879):1011, 2008. doi: 10.1126/science.320.5879.1011a.

A VERSATILE CLOUD BASED PLATFORM FOR EARTH OBSERVATION INTERDISCIPLINARY ACTIVITIES

Olivier MELET, Ground Segment Architect, CNES, olivier.melet@cnes.fr

Alain MONTMORY, IT Architect at Thales Services, alain.montmory@thalesgroup.com
Eric LI CHAM YON, Software Engineer at Thales Services, eric.lichamyon@thalesgroup.com

ABSTRACT

With the increasing complexity of spatial operations and the schedule constraints to set up operational Ground Segment, test phase should be optimized by allowing multiple teams to integrate and test in parallel their own components, and to perform expertise activities.
Based on these needs, CNES has designed and developed with Thales Services, a multi-purpose and multidisciplinary platform for test and expertise in order to support development and operations of ground segment components for Earth Observation.
This platform is based on the latest cloud technologies and allows multiple teams to perform individually their business activities in a very efficient way. They can set up, in a matter of minutes, a test or an expertise environment in a private, secure and managed execution context.

***Index Terms*—** Cloud, multi-purpose platform, versatile platform, Earth Observation ground segment facilities, HPC Cluster, image processing.

1. INTRODUCTION

With the increasing complexity of spatial operations managed by the French Space Agency (CNES) and the schedule constraints to set up operational Control & Mission Ground Segment, test phase should be optimized by allowing multiple teams (Control Command, Mission Planning, Payload processing) to integrate and test in parallel their own components. In addition, these teams need an expertise environment during the operational phase to perform analysis with their own tools.
Based on these needs, CNES has designed and developed with Thales Services, a multi-purpose and multidisciplinary platform for test and expertise to support development and the operations of ground segment components for Earth Observation.
This test and expertise platform built upon the latest Open Source Cloud Technologies offers the possibility to set up a test environment within minutes. An environment consists of a coherent set of resources such as compute machines, a dedicated network, dedicated storage capacities, access to common support services like remote desktop access, printing, media import, file exchanges between centers, backup, archive facilities and central user management.

2. THE GOAL

This platform has been designed to support activities of interdisciplinary teams working on Earth Observation programs. The users of this platform are the different teams involved in the whole Ground Segment qualification including User Segment, Mission and Planning, Satellite Control, and Payload Processing.
The platform offers the possibility to set up VM (Virtual Machine) hosting different OS (different flavors of Linux and Windows) as well as Baremetal computer (Baremetal: computer hosting direct guest Operating System without hypervisor and managed by the infrastructure as if they were Virtual Machine resources) organized in a coherent environment. This environment has:

- its own private Network,
- an access to the different computers of the environment (VM and/or Baremetal) through a Web Portal with direct desktop display export (including 3D rendering)
- the possibility to set up various components in the environment thanks to DevOps approach (DevOps is a mix between Development and Operation. it is when exploitation phase is taken into account in earlier phase of a project using the appropriate tooling to install and manage component on the target platform).
- the possibility to launch support service (VM backup, import and export data from several type of medias, application log visualization) directly from web portal.

3. THE MEANS

The platform is built using the OpenStack [1] ecosystem embedded in a packaged distribution “Fuel” from Mirantis.
OpenStack controls large pools of compute, storage, and networking resources through a dashboard or via the Openstack API. Key advantages of Openstack are its open source nature (customizations are possible), the ability to manage virtual and Baremetal compute resources, the

support from multiple vendors for compute, network and storage resources via the driver mechanism.
Thales Services added various functionalities which rely on the Open Source Components:

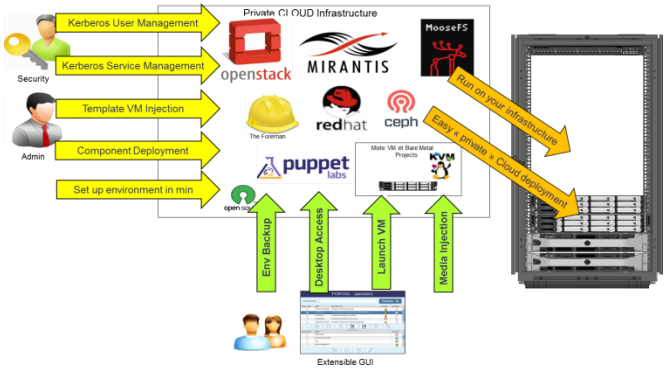

Fig. 1. Main technologies and functionalities of the platform

Each Open Source or Thales developed component brings the following features to the platform:

3.1. Integration with User Management

This integration is brought by RedHat IPA stack (IPA: Identification, Auditing, Policy, a RedHat packaged identification solution which coupled Ldap with Kerberos). So, identifying, auditing or setting a policy is done through one component across the platform.

3.2. Services discovery coupled with user profile

The coupling between IPA, WebPortal and Keystone (Open Stack authentication component) allows managing the "see only what you need" security policy by showing only environments and services for which the user is allowed for.

3.3. Web Portal : the platform access

The WebPortal (based on Vaadin GUI library [6]) access which offers an environment synchronized view with the underlying OpenStack backend environments management. It allows accessing different services like:

- An access to the platform with any of the "Screen/Keyboard" available
- A direct access to the remote desktop of the VM through rdesktop (for Windows VM) or noMachine[7] (for Linux VM)
- View the log of the current environment,
- View the log (Security and Application) of the whole platform with the right user profile
- Save/Restore a VM
- Import/export of media (USB, DVD) via a specific service.

3.4. Unified HOMEDIR management

Each platform user has access to several VM but:

- Its HOMEDIR remains constant as it is physically stores on a SharedFs: MooseFs [5] and its home directory can be viewed from every VM.
- The import/export media service also store there data on this SharedFs in order to be able to inject them into business environment.

MooseFs is a shared Posix File System builds with an aggregation of local disk resources which are viewed as a global namespace thanks to a central "master server" (backed by "metalogger" to avoid Single Point Of Failure).

3.5. Cloud scope log management for privileged user

Privileged platform user has access to the whole platform log.
This display can highlight some system problems, likes an error in an environment which is caused by a system problem (example disk failure) in the Open Stack infrastructure.

3.6. Open Stack deployment specificities

Regarding the OpenStack « private Cluster » deployment, the following choices have been made :

- The start stack is a Mirantis fuel OpenStack distribution [1]
- An active/passive deployment of OpenStack controllers with Heartbeat [8] and HaProxy [9] load balancer has been added
- A full deployment of DevOps tools based on Puppet and Foreman OpenSource components [3][4].
- Storage software to offer storage block to VM with the use of commodity hardware. The ceph rados [2] component is used for that purpose.
- VM images are also managed by ceph storage.
- MooseFs distributed file system to offer shared resources among environments and users.

Ceph and MooseFs do not work at the same level. Ceph is a Block Device whereas MosseFs is a Posix FileSystem layer. CephFs which provide Posix FileSystem layer exists but, at the time of the platform deployment, CephFs was not marked as "production ready".

3.7. DevOps approach for all platform layers

One of the design principles of the platform was to have a DevOps approach. This approach emphasizes the collaboration and communication of both software developers and operational staff by automating both software delivery and infrastructure changes. The combination of Openstack and Puppet is the key in achieving this goal. We are able to easily provision compute resources, configure them, and install business software ready to be tested by the right people (via the central user management)

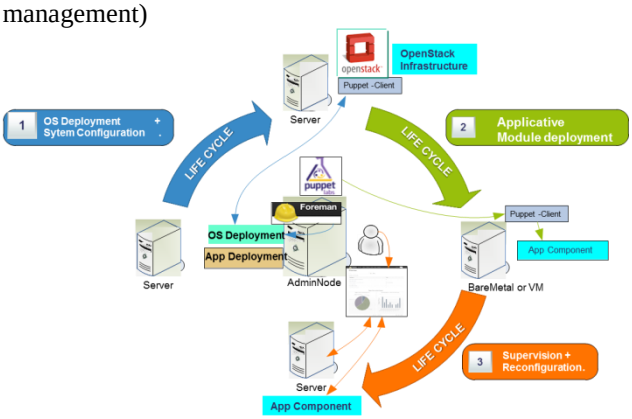

Fig. 2. Platform components Devops lifecycle

3.8. A mix VM and Baremetal deployment

Another strong capability of the platform is its ability to manage Baremetal and traditional VM in the same environment. Baremetal are used for intensive computing where the overhead of a VM is not acceptable.
A Baremetal driver is deployed on an Openstack compute node and is configured to be able to manage Baremetal computers (add PXE and IPMI tools for provisionning). With this technology Baremetal and VM are managed in a transparent way for the user and with the same tool, Web Portal access for user point of view and OpenStack horizon for administrative purpose.

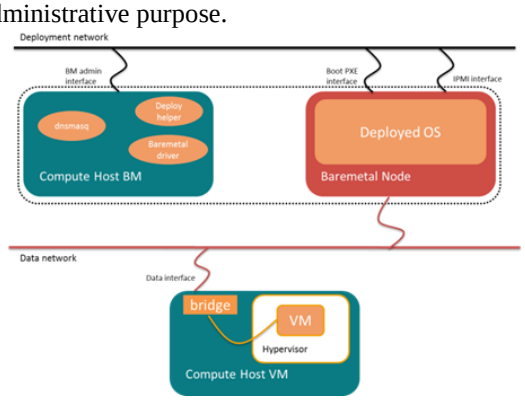

Fig. 3. VM and Baremetal deployment principle

4. THE FUNCTIONNALITIES

Once deployed, the platform offers the following services:

- Set up a new VM template and inject it into OpenStack,
- Set up environment with private network and user managed firewall
- Single entry point to configure machines of the entire platform
- Management of virtual and Baremetal resources indifferently
- Centralized user management
- A friendly user access point with a click and display policy based on a strong (Kerberos) Single Sign On and services access filtering.
- A media import service allow importing into the platform only controlled and mastered media with a registration of all import and export action.

A test environment (an environment with several VM and/or Baremetal) dedicated to the test of a specific component or a complex processing chain could be setup in one hour from scratch when the component deployment plan is already available.

From User point of View the Web portal is the most important component, it is its primary access point to the platform. Full text search in name and environment description allows selecting very quickly the user's environments and then a click on the target VM raise a graphical desktop regardless of the Operating System (Windows or Linux) of the target machine. The WebPortal is coupled with the LDAP profile of the user and then display only the services the user has access to.

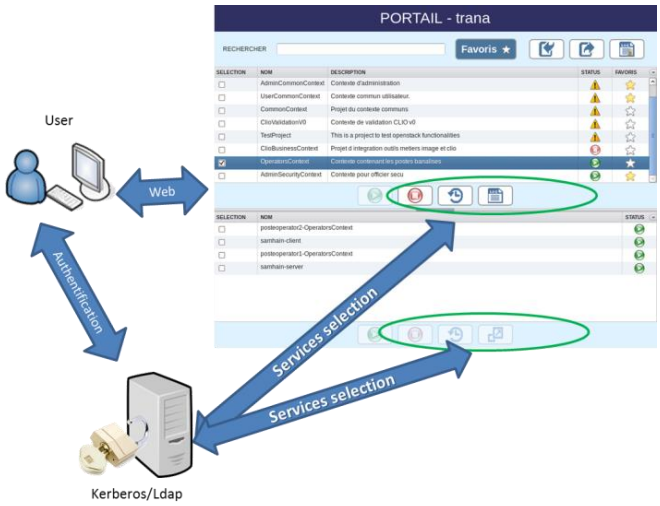

Fig. 4. Platform Web portal as primary access point

doi: 10.2788/854791

5. THE OPERATIONAL FEEDBACK

The activities of test or expertise require deployments of different environments according to the phases of the projects (even for the same ground segment component). It is, for example, necessary to deploy an environment (composed of different computer facilities) to perform the test of a software and another environment to analyze telemetry with the same software. Furthermore, the environments can have very different needs in term of computing resources type (virtual machines, physical machines for massive processing, physical machines with 3D graphic acceleration…). During the development, these constraints are present for every components of the ground segment.
To cope with these needs, it was usually planned to procure dedicated reference platforms for a given project with dedicated software and hardware resources. The sizing of this dedicated platform was therefore not optimized and could not allow sharing computing resources between different environments and business teams. Moreover, the previous architecture of dedicated reference platforms required long time deployments which thus had to be planned well in advance.

Thanks to this new architecture, the deployment of these different environments becomes extremely easy and fast (about an hour, for example, for the deployment of a environment containing several virtual machines and deployment of software prerequisites). This agility allows platform administrators to be very reactive and to optimize the sizing and usage of the platform according to the phases of the projects (it is possible for instance, to stop a whole environment and thus to release resources which can be used by other environments). In addition, it is possible to integrate into each environment, a wide variety of computer resources such as VM, Baremetal with high performance storage to instantiate a computer cluster or even include graphics stations to use software components with 3D views.
The use of the platform is further facilitated by a generic access through a portal web access that gives the access to the different environments personalized according to the profiles of the users.

After several months of operational use, users confirm that the platform fully meets their needs. In addition, several test phases carried out by different teams have already been held and have confirmed the gains in terms of deployment compared with previous architectures.
Several environments have been deployed for different business teams with different design environments and different software pre-requisites.
For instance, a very complete environment was deployed to host performance bench of an Earth Observation image processing. This environment includes virtual machines as well as Baremetal (blade computers) managed by a Distributed Resources Manager (DRM) in addition to a high-performance storage system to address a High Performance Computing (HPC) facility. The Virtual Machines host the processing orchestrator software, the image catalog and other image tools as IDL/ENVI while the HPC cluster performs the image processing. It has been therefore possible to deploy in this environment all the elements representative of an operational image processing center.

6. CONCLUSION AND OUTLOOK

The use of the platform has validated operationally the advantages of such architecture. The rapid deployment of business environments, the ease and flexibility of use is appreciated by all user teams.
This first experience with back office kind of activities (test and expertise) has highlighted the interest of this architecture to accommodate different activities compartmentalized but benefiting from the same pool of computing resources and tools.
In addition, it has been demonstrated that it is possible to deploy representative environments of operational image processing centers. Similarly, the deployment of massively parallel computing technologies such as Spark or Cascading can be considered in just as it had been possible to deploy an HPC cluster.
Finally, those very promising results suggest the possibility to use this type of architecture for operational activities of ground segments and to build multi-mission centers on the basis of shared IT resources with great agility of deployment and high flexibility of use.

10. REFERENCES

[1] OpenStack Mirantis Fuel distribution: https://www.mirantis.com/.
[2] Ceph rados: http://docs.ceph.com/docs/master/rados/
[3] Puppet: https://puppetlabs.com/
[4] Foreman: http://theforeman.org/
[5] MooseFS: https://www.moosefs.com/index.html
[6] Vaadin: https://vaadin.com/home
[7] noMachine: https://www.nomachine.com/
[8] HeartBeat: http://linux-ha.org/wiki/Heartbeat
[9] HaProxy: http://www.haproxy.org/

SPARKINDATA: EARTH OBSERVATION APPLICATION MARKETPLACE HOSTED IN THE CLOUD

C.Jaber[1], T.Habib[1]

[1]Atos, 6, impasse Alice Guy 31024 Toulouse, France

ABSTRACT

Several trends are impacting today's earth observation (EO) landscape. Free and open access to the Copernicus data, the availability of a wide variety of open data, the creation of consortiums bringing organizations together to collaborate on missions, systems (e.g. CEOS's Recovery Observatory, GEO's GEOSS) or partnerships between IT and Telecom actors to provide the computational and storage resources to process the data (e.g. Helix Nebula Cloud). This paradigm shift is also translated by significant contributions from the IT domain into the EO world. The data era we are currently living in is implying the integration of the EO data in the standard streams of exchanged information, widening its user base and increasing the market potential and societal benefits from this global, precise and high quality data. The emphasis in this paper is put on our efforts to build the SparkInData (SID) application marketplace. The objective of SID is to provide a complete ecosystem offering processing algorithms, frameworks and services making EO data processing much more accessible. This is accompanied by a strong partnership with data, algorithm and service providers, which makes our efforts market-driven, aiming at maximizing the value of available data sources.
In this paper we focus on presenting the technical architecture of SID, showing its layered model (IaaS, PaaS and SaaS) [1], its technology stack and how we are tackling the questions of Variety, Volume and Velocity related to this Big Data platform.

Index Terms— SparkInData, Helix Nebula Science Cloud, IaaS, PaaS, SaaS, Kubernetes, Docker, SOA,

1. THE SPARKINDATA PROJECT

The SparkInData project is partially funded by the French General Commission for investment. The partnership which includes today (Atos, Aerospace Valley, BRGM, CNES, Mercator-Ocean, TerraNIS, Geosigweb, Geomatys, IGN, IRIT, and INP Purpan) aims to set up an exploitation platform for EO data with a special focus on making the building, deploying and operating new services for the consortium partners and other new users as easy as possible. The partnership is thus extensible and the consortium boundaries are in continuous expansion to include additional partners each bringing a new service or data source. SID also connects to existing platforms from the different partners (e.g. CNES's PEPS or THEIA) as the objective is to build a federation platform and to leverage existing assets. Our objective is to work on the platform's long term sustainability and future developments. In addition to data and IT level access, the design of SID offers the tools to quickly prototype, test and to push into production an algorithm/service.

2. THE IAAS LAYER

Although the SID software components are totally cloud agnostic, IT resources deployed by the platform are virtual machines (VM) that are today spawned on "Helix Nebula Science Cloud" [2]. The Helix Nebula cloud's "Blue Box" interface enables seamless multi-machine, mutable deployments on cloud providers within the consortium (Atos, Cloud Sigma, Inter-route, T-systems, and EGI-fedCloud). SID aims to be fully independent from the cloud provider thus it is possible to use any cloud provider resources or even to deploy it on a bare-metal machine cluster.

3. THE PAAS LAYER

The SID project adopts the PaaS approach with the advantages it offers:

- Cloud Scalability (resources sufficiency)
- Cloud Elasticity (resources adaptability)
- Cloud on demand model (finer billing granularity)

In the development of SID PaaS we have identified the following challenges:

- Failover Handling
- Scaling (including auto-scaling)
- Fast deployment
- Resources sharing
- Multitenancy (Authentication and Isolation)
- Billing
- Big data handling

These are obviously not specific to SID but are general to platforms of this type, independently from the application domain.
Besides, the aspects inherent to the PaaS construction, the SID context brings its own set of requirements. The number of heterogeneous parties involved in the construction of the platform is high and continuously increasing each with his own legacy tools and methods.
On one hand, each user/partner has its preferred development language or framework, several legacy algorithms and applications can be reused and deployed on the platform as services, thus constraints on the construction and the hosting of a given service has to be as soft as possible.
On the other hand, the overall process of service creation has to be as efficient and simple as possible from the development of the service on the partner/user premises to the deployment on the platform.
The PaaS architecture of SID has been built from the ground up in order to answer these challenges. In this paper, the focus is brought on:

- The **Deployment model** covering matters related to software set up on IaaS resources.
- Applications/Services high level organization fostering platform modularity and extensibility through a **Service Oriented Architecture.**
- **Generic and EO off–the-shelf services** essential for the proper functioning and the quick extension of the platform.

3.1. The deployment model

The very first aim of the SID platform is to offer a way or place where all partners' services will be hosted. Each of these partners will potentially develop applications in different languages or using different frameworks. To achieve this goal, one solution would be to make them all use the same language or framework. Of course this solution is not envisaged as it brings a lot of constraints on the user side and will prevent the platform and community from benefiting from existing performant software modules. In addition, it will prevent any use of new technologies that potentially bring advances in terms of functionalities and performances. Imposing such "unique language" will thus lower the platform's attractiveness for user adoption, increases the integration complexity on the user side and deprives the platform from future interesting software modules.

Since 2013, containers technologies have gained momentum enabling virtualization features offered by the Linux kernel. These features have been around for a long time now but projects like **Docker** have created an ecosystem where containers are easier to create and use [3]. A container provides an abstraction layer between the hosting Operating system and the application running inside the container. Unlike the traditional fully-fledged virtual machines (VM), containers need only to embed the required libraries and applications to run and leverage the hosting machine OS kernel to function normally. This additional abstraction reduces the coupling between the application and the machine on which it is executed and reduces drastically deployment time (less than a minute in average).

Building on this application container approach, the scalability, resilience, isolation and deployment are handled by container orchestration modules. This type of modules provides both enough isolation between the deployed containers and a secure execution environment.
Many orchestration frameworks have been introduced with different philosophies and each one of them has its pros and cons. SparkInData's choice landed on the Google backed technology: **Kubernetes** [4]. This latter introduces concepts allowing in particular:

- Simple deployment and organization of containerized applications across multiple hosting machines.
- Applications interactions are also made easier by providing discovery mechanisms as Domain Name System (DNS).
- Storage Management (permission control, quotas,…)
- Containers scheduling enables smart and **fine grained resources use and sharing** between applications.
- Application failover and rescheduling in case of machines or application failure or termination.

With respect to our goals, the combination of Docker and Kubernetes gives us the necessary tools to execute and manage the potential numerous applications to host and manage on our platform.
Thanks to this "lightweight" virtualization layer, the application delivery process illustrated in Figure 2 is feasible and much more deterministic. In fact, the partner can be sure that the application will work as it worked in his own development environment.

Another important aspect of the deployment model is the split between the services and the persistence layers. In fact, different and separated clusters of VMs host the **distributed file system (DFS)** and various databases needed by the services. This separation simplifies platform resources deployment, management and elasticity.

3.2. The Service Oriented Architecture

SID platform contribution to EO field actors is not only an access to physical resources but aims also to offer an ecosystem where applications interconnect and benefit from each other. Thus, the platform has been conceived following a **service-oriented architecture**. In fact, all features and functionalities are encapsulated in **Services** with the least possible coupling between them. All partners and users intending to participate and make accessible their applications to the rest of the platform shall deliver also their applications following this concept of Service. Obviously a minimal amount of standardization is required. Today there are two options, either expose a **documented REST API** or adopt one **the OGC WxS standards** (if relevant to the service).
A **"Service Management"** service is the entry point supplied by the platform to developers for the definition and deployment of their applications. In particular, the developer will be able to define:

- Each component of its application
- Physical resources (e.g. CPU, RAM, storage) needed.
- Auto-scaling triggers for one or several of his/her application components.

For its part, the service manager enforces the constraints on interface (e.g. documentation), on physical resources and makes all services discoverable. As discussed earlier, **Figure 1** shows the proposed workflow for the delivery and integration of applications into SID. It is important to emphasize that the Service strategy, greatly increases the reach of a developer's application and enables other developers to incorporate it in their processing chains.

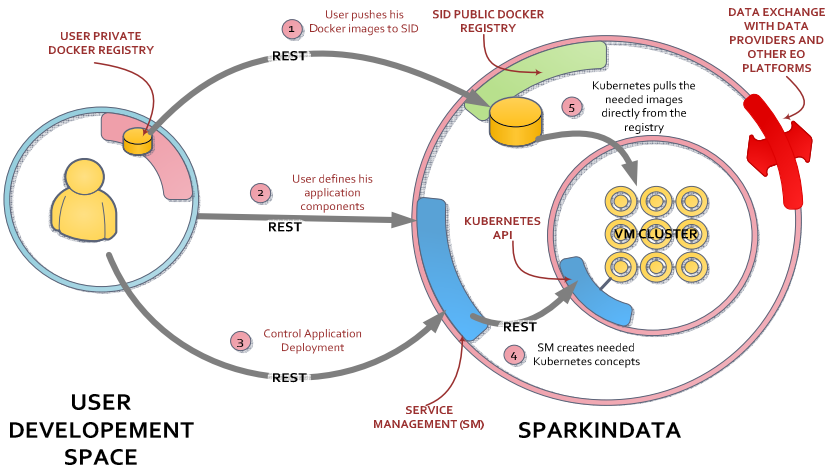

Figure 1: SparkInData delivery model overview

3.3. The Common services

SID features have been designed and thought to simplify EO data mining and handling. Common services (e.g. the services management service, billing service, authentication /authorization service, data ingestion service, Resources management service, computing framework service …) with multitenancy support are provided off-the-shelf. This has the advantage to make service development much quicker and focused on the real added-value. The Common Services provided and implemented by the platform have scalability, multitenancy and failover in mind in order to cope with the volume of data and provide the needed availability for the system. The following is a non-exhaustive list of services that are available on the platform:

- As described in the previous chapter, a **Service Manager** playing the role of an interface between users/services and the platform deployment system.
- An **Ingestion Service** has been developed from scratch for SID. Enabling us to adapt as much as possible to our EO data providers thanks to a plugin system. This service spawns an elastic group of "workers" among which is dispatched the needed data download and post-processing jobs
- Resources usage varies according to services deployed and processing occurring on the platform. SID platform shall spawn new virtual machines or release them in order to be cost-effective. The **Resources Manager** is the intermediary between the platform and the IaaS layer. In fact, the Resources Manager receives feedback from Kubernetes container scheduler to adapt the current amount of available VMs (Scale-up/down). In the same way, the DFS volumes are also monitored by this manager in order to increase or decrease automatically the amount of dedicated VMs and thus the storage space.
- Available datasets on the platform and on data providers' premises metadata and location are available through an OGC/OpenSearch compliant interface. Moreover, an ongoing work on catalogue spatial sharding and semantization will enable respectively to cope with the high amount of metadata handled by the system and to simplify the mapping between SID catalogue model and the various data sources ones.
- Three approaches can be adopted to develop and integrate processing chains in SparkInData. One through the deployment of a Docker containing the processing blocks (or several Dockers for a

complete pipeline) through the Service Manager, the second is through the deployment of Dockers fitted with the OGC WPS interface and the third is through the use of SparkaaS provided as a Common Service by the platform. The first solution would be to create from scratch an application and to deploy it through the Service Manager. This approach is convenient for users not willing to embrace a given paradigm imposed by an already existing service like the **SparkaaS** or the **OGC WPS** ones. The Spark as a Service gives a controlled access to a shared cluster of VMs running Apache Spark. The user submits his/her application package to the service and gets the results (published on the platform) to be visualized or used as an input for another service. The Spark cluster can also be accessed through web based notebooks [5] allowing interactive prototyping in different languages. Finally, developer can leverage the off-the-shelf WPS service. The developer delivers the processing algorithm encapsulated in Docker images implementing a given interface and the WPS service takes the responsibility to exposes them as standard WPS processes that can be chained and to execute them throughout the Kubernetes cluster. This last approach enables other users, if permitted by the developer, to benefit from the processing block and eventually incorporate them in new processing chains. **Figure 2** illustrates the different aspects exposed in the previous sections and gives a high level view of the platform.

3. THE SAAS LAYER

The end-user accesses and instantiate the different published services through a user-friendly interface called: **the marketplace**. The user is then redirected to the service interface (if it exists) or to common workspace interface allowing to call any provided service API and to visualize conveniently the results. In addition to data and PaaS level software blocks, SID partners works on the building of business oriented high level services in various fields going from Agriculture and territorial collectivities to Energy and Oceanography.

4. CONCLUSION

In this paper we have presented a part of the SparkInData platform architecture. The focus was made on some of the key aspects related to the interface with the Infrastructure layer, the data management module, the deployment methodology, the service oriented architecture implementation and finally the end-user interface through the SAAS layer. While there are already operational services running in SID, the work on the SID platform is still ongoing and it is expected to enter the full operational phase with a general public commercial offering by Q3 2016.

5. REFERENCES

[1] Qi Zhang, Lu Cheng and Raouf Boutaba, "Cloud computing: state-of-the-art and research challenges," Journal of Internet Services and Applications, vol. 1, no. 1, pp. 7-18, May 2010. 6
[2]http://www.helix-nebula.eu/
[3]http://www.docker.com/
[4]http://kubernetes.io/
[5]http://jupyter.org/

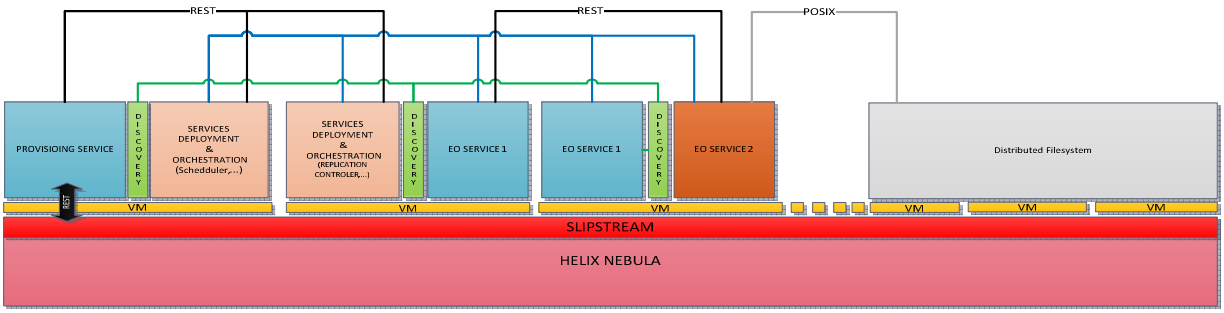

Figure 2: SparkInData Cross layer overview

ENABLING COLLABORATION BETWEEN SPACE AGENCIES USING PRIVATE AND CLOUD BASED CLUSTERS

M. Poncet[1],T. Faure[2],C. Cavet[3],A. Petiteau[3],PM. Brunet[1],E. Keryell-Even[2],S. Gadioux[2],M. Burgaud[2]

[1]CNES, 18, avenue Edouard Belin 31401 TOULOUSE CEDEX 4, France
[2]Atos, 6, impasse Alice Guy 31024 Toulouse, France
[3]APC, AstroParticule et Cosmologie, Université Paris Diderot, CNRS/IN2P3, CEA/Irfu, Observatoire de Paris, Sorbonne Paris Cité, 10, rue Alice Domon et Léonie Duquet, 75205 Paris Cedex 13, France

ABSTRACT

Mission centers need more and more resources to ensure their operational activities. Scientific process requires more CPU and space disk than ever before. Virtualization and Cloud technologies offer both additional flexibility and opportunities for standardization of space agencies infrastructure. Despite these important evolutions in system architecture there still exists some cases where the available amount of resources is insufficient. A solution when preparing a mission center where the available resources are inadequate can be to use external infrastructure. This paper describes alternatives, the foreseen architecture and the remaining challenges to address.

Index Terms— Cluster computing**,** Big Data, Interoperability of cluster, Mesos, Docker, XRootD, Chronos

1. MOTIVATIONS

The French Space Agency is often confronted with a wide variety of architecture problems. For some missions the infrastructure shall be ready to support a large amount of computation occurring once a year for a reasonable set of data. It also exists missions with constant computations on a huge amount of data scattered around the world. To handle all of these situations the French Space Agency frequently analyses the different possible alternatives. Cloud and Virtualization makes technical architecture more flexible and eases the modification of resource allocations [1]. The goal of this study is to determine the right architecture when the set of available resources is not sufficient, e.g. when the French Space Agency has to count on external services like cloud, or when other agencies need to rely on French Space Agency resources.

2. USE CASE

The objective of this study is to analyze some use cases and determine the appropriate architecture for distributed infrastructure. A CNES use case focuses on image processing for mission centers. The objective is to demonstrate that from a pool of images, resources shared through different mission centers, it is possible to run processes transparently on them at different points of the cluster. If for some reasons the cluster is overloaded some machines could be added even if they are located elsewhere.
The use case was defined by examining different CNES projects: PEPS, NGO/eLISA, Euclid, Muscate. This led to the definition of a test case encompassing the key issues, and using Orfeo ToolBox (OTB) algorithms (correlation) on images captured from the Landsat satellite. As it is possible to use these algorithms using command line, this additionally demonstrated the ability to run some legacy code on a cluster.
These tests were performed on the CNES private HPC cluster, with an extension using an Atos Cloud.
The global picture of our use case is presented in the next figure.
After the success of the tests with OTB, the study performed additional experimentation with NGO/eLISA use case.

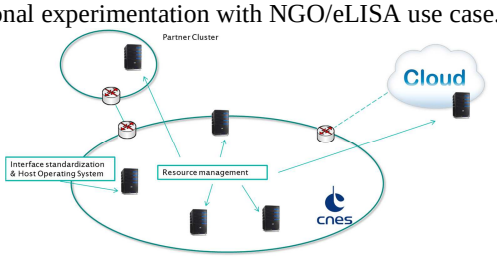

Figure 1 : Global picture of study use case

2.1 NGO/eLISA

The NGO/eLISA[1] (evolve Laser Interferometer Space Antenna) mission has the goal to detect gravitational waves emitted in the mHz frequency range (massive black hole binary, white dwarf…). This ESA L3 mission will be launched in 2034. In order to validate the new technology of the NGO/eLISA mission, the LISAPathfinder satellite was launched December 3, 2015. LISAPathfinder will produce data in order to indentify the characteristic noises in a spatial interferometer. The data analysis code of LISAPathfinder is based on MCMC algorithm. The code is implemented in

[1] eLISA: https://www.elisascience.org/

MATLAB/C++ and is multi-threaded with MPI in order to decrease the computing time. We have built a Docker image (see Sect. 5.1.1) containing the code and the dependences (MPI, LAPACK, BLAS, ATLAS). This image is shared with a private Docker registry in order to preserve the input private data included in the Docker image. Benchmarks of performance based on the computing time of the code were realized in order to validate several workflows. The execution was accomplished on the local Arago@FACe[2] cluster (with MKL) and on a virtual cluster of the StratusLab@LAL[3] cloud (without MKL, deployed with SlipStream[4]).

Workflow	Cores	Computing time
Arago@FACe	8	~2h40
StratusLab@LAL	8	~9h45

Table 1: performance benchmarks

The computing time differences are due to the availability of the MKL library on the Arago cluster and the commodity hardware used in the cloud.

3. METHODOLOGY

This CNES R&D study has been organized around three main stages. The first one named "A Phase", which took place end of 2014, was a "state of the art" step with namely a review of the literature. Its main objective was to determine candidate's technologies selected according to our requirements and our strategy. The second one named "B Phase", during 2015 and the beginning of 2016, is the prototyping step itself with the prototype architecture definition, implementation, test case selection and execution. It main objective is to demonstration the chosen architecture. The last one, named "C Phase", is a step of results validation, analysis, conclusions and outreach that is currently ongoing. The A Phase took 10% of the whole time of the study, the B Phase 85% and the C Phase 5%.

4. EXECUTION

4.1 Platform needs

To support our needs, the platform shall respect a set of technical requirements listed below.
REQ-001: The platform shall run software independently on a bare metal server, on the cloud or on partner infrastructure.
REQ-002: The platform shall work on « standard » data storage locations.
REQ-003: The platform shall be able to run scientific software with high requirements of speed performance.
REQ-004: The platform shall be able to run scientific software with high requirements of data volume performance.
REQ-005: The platform shall be flexible and shall support horizontal and vertical scalability.
REQ-006: The platform shall be able to relaunch software jobs which fail.

4.2 Evaluated Technologies

The technologies evaluated during this study have been categorized in three categories: software standardization, resource management and distributed storage.
Software standardization ensures REQ-001, REQ-003 and REQ-004 satisfaction. The evaluated technologies are separated into two groups. The first concerns technologies with no software adaptations: the virtual machines and the containers. The second one is framework standardization: for example: usage of Map Reduce or Spark.
Resource management ensures REQ-005 satisfaction. The evaluated technologies are popular open-source programs: Mesos[5] and YARN[6]. This section presents a job scheduler compatible with the chosen resource manager satisfying REQ-006.
Distributed Storage ensures REQ-002 satisfaction. The goal is to find a suitable distributed file system compatible with the multi-center approach. Fast and configurable, it shall ensure that files are available through the different clusters. The evaluated distributed storage systems are HDFS[7], XrootD[8], GlusterFS[9] and CephFS[10].

5. ARCHITECTURE

5.1 Static architecture

5.1.1 Software standardization

To standardize the way the software is packaged, executed and deployed, the chosen solution is Docker[11]. Docker

[2] FACe: Centre François Arago, Laboratoire AstroParticule et Cosmologie.
[3] StratusLab: http://www.stratuslab.eu/; LAL: Laboratoire de l'Accélérateur Linéaire.
[4] SlipStream: http://sixsq.com/products/slipstream/.

[5] http://mesos.apache.org/
[6] https://hadoop.apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/YARN.html
[7] https://hadoop.apache.org/docs/r1.2.1/hdfs_user_guide.html
[8] http://xrootd.org/
[9] https://www.gluster.org/
[10] http://docs.ceph.com/docs/master/cephfs/
[11] https://www.docker.io/

allows the creation of containers, lightweight images which run on Linux systems (in our case).
Docker is a solution based on the concept of containers. It makes it possible to run isolated processes inside one Operating System. The only requirement to create a container is to rely on a specific Operating System (e.g. Linux, Windows).
Compared to standardization by frameworks (using Map Reduce or Spark) the containers and the virtual machine are less intrusive and are adapted for legacy software.
Docker, due to its core concept of resource isolation in a kernel, is faster than a virtual machine [2] that's why we chose it. Also Docker is more focused on running processes whereas Virtual Machines are focused on running one or several services. Docker was concerned more relevant for this study as the goal is to execute scientific algorithms.
In our context, Docker will allow the creation of images for legacy scientific software. Docker runs files named "Images" which are built from already defined Docker images and some executables. Once the image is created they execute in the same way on every host machine.

5.1.2 Resource management

To support resource sharing between mission centers an analysis has been made on the existing tools allowing the management of cluster resources. Two cluster resource managers have been evaluated: Mesos and YARN. The chosen resource manager is Mesos. Mesos offers available resources to higher level frameworks. Each framework contains a specific scheduler which will receive the offers from Mesos. According to its list of tasks it accepts or not the resource offer.
The biggest pro of Mesos is that the framework does not have the visibility of the whole cluster, it adapts its scheduling according to the resources offered. Another advantage of Mesos is the availability of several frameworks with various targets. Using Mesos it is possible to run Spark, Yarn or Docker jobs indifferently on the cluster.
A resource offer example described in [3] is presented below:

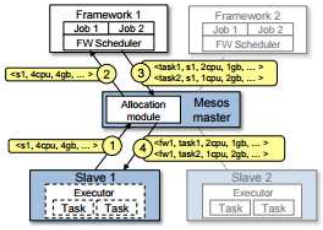

Figure 2 : resource offer example [3]

Our use case needs to run legacy software on Docker. Chronos[12] Mesos Framework is a framework designed by Airbnb and open sourced at the end of 2014. Its role is to schedule Docker jobs. If some big data processes are developed in the future, the Mesos choice would still be relevant as Yarn or Spark frameworks exist.
In the meantime, HPC batch scheduler editors have integrated Docker support in their product. In particular, PBSPro, which is installed into the CNES HPC center, can manage Docker jobs seamlessly in version 13. PBSPro server is able to pull a docker image, to launch it, to export the environment variables needed by the job to run inside the container, to mount the job directory and stagein/stageout directories in the container, to use "docker stop/rm" command to clean up the job specific docker container instances.
Besides, one strength of HPC batch schedulers is their powerfull scheduling policy which is a clear advantage in a shared production platform. Finally the resource management functionality seems to be more effective in PBSPro than Mesos to achieve a high sustained utilization rate. A detailed comparison of Mesos and PBSPro is not in the scope of this study and shall be study in the future.

5.1.3 Distributed storage

XRootD is a technology used by CERN which provides access to data repositories. XRootD is not a distributed file system but a service for performing access to data of any format [4]. XRootD resources can be accessed through FUSE. Using XRootD a client asks for a resource, XRootD works as a router which locates the resource on the cluster, once the resource is located, the client is connected directly to resource and can download or process it.
XrootD is the appropriate choice for this study because it has built-in features for data federation [5].

5.1.4 Operating Systems

No major constraints have been defined for operating system choice. To use Chronos framework, Docker has to be supported. It is supported since Linux 3.10. In addition to Linux based systems like Debian or Red Hat, "atomic" operating systems exist. They are optimized to run containers like Docker. Some examples of Atomic operation systems are CoreOS, Project Atomic, or Snappy Ubuntu Core.

5.1.5 Clusters interconnection

[12] https://mesos.github.io/chronos/

The first step for connecting clusters is made with Mesos, by manually starting additional Mesos slaves on the extension cluster. The interconnection of clusters has been realized through route definition across the different networks.
Further experiments for automating resource scaling will involve modifications to Mesos. The change consists in allowing frameworks to express “wishes”, i.e. resources that the framework would need, so that Mesos can estimate the remaining volume of work and request additional resources if necessary.

5.2 Dynamic architecture

In the selected architecture, Mesos manages a set of nodes located on the same network. The defined network contains machines inside the CNES infrastructure and machines running inside the Helix Nebula cloud. A user is able to submit a job to Chrono. Thanks to Mesos Chronos the jobs on a machine are deployed. The machine which receives the job downloads the Docker image in the distributed Docker registry and then executes the Docker process. Storage folder is mounted from the distributed storage managed by XRootD on the Docker image. Once the job is completed Chronos can deploy another job on the available machine.

5.3 Security

In first versions of Docker, some concerns have been expressed by the community: the most important was about security. In some cases it was possible for a user to grant root privileges and to break the isolation property of the containers. There are still ongoing analyses about Docker security by the IT department but Docker improved its security constantly making it reliable [6]. Nevertheless general recommendations for Docker usage are: use version 1.7 minimum, run containers as “non-privileged”. If there are major security requirements for the system it is still possible to run Docker containers inside a Virtual Machine.

5.4 Installation and configuration

The Operating System was deployed on the CNES cluster using Cobbler, this tool made it possible for a relative beginner to perform a complete installation of CentOS on the nodes of the clusters. The configuration and the installation of Docker, Mesos, and XRootD were made using Ansible, the same configuration has been used for physical machines (CNES infrastructure) and virtual machines (Helix Nebula cloud). The set of available tools and documentation allowed to perform this step in less than one month.

6. CONCLUSION

In conclusion, this study demonstrates the capacity to run predefined scientific programs (NGO/eLISA provided by APC) with no adaptation; the only work performed was limited to the “containerization” of the program. This study also demonstrates the opportunity for a space agency to verify that it is possible in terms of network connection to plug its infrastructure to a cloud.

7. NEXT STEPS

The next step of this study is to evaluate performances and network latency impacts using a cloud infrastructure. If results are relevant some tests can be realized with other space agency mission centers.
It is also important to note that the community is very dynamic and there is frequent news about new Mesos frameworks, Docker schedulers or new technologies. These technologies have been released during our study and should be analyzed.

8. REFERENCES

[1] Pasian, F., Hoar, J., Sauvage, M., Dabin, C., Poncet, M., & Mansutti, O. (2012, September). Science ground segment for the ESA Euclid mission. In SPIE Astronomical Telescopes+ Instrumentation (pp. 845104-845104). International Society for Optics and Photonics.

[2] Felter, W., Ferreira, A., Rajamony, R., & Rubio, J. (2015, March). An updated performance comparison of virtual machines and linux containers. In Performance Analysis of Systems and Software (ISPASS), 2015 IEEE International Symposium On (pp. 171-172). IEEE.

[3] Hindman, B., Konwinski, A., Zaharia, M., Ghodsi, A., Joseph, A. D., Katz, R. H., ... & Stoica, I. (2011, March). Mesos: A Platform for Fine-Grained Resource Sharing in the Data Center. In NSDI (Vol. 11, pp. 22-22).

[4] Dorigo, A., Elmer, P., Furano, F., & Hanushevsky, A. (2005). XROOTD-A Highly scalable architecture for data access. WSEAS Transactions on Computers, 1(4.3).

[5] Colamaria, F., Colella, D., Donvito, G., Elia, D., Franco, A., Luparello, G., ... & Vino, G. (2015). Local storage federation through XRootD architecture for interactive distributed analysis. In Journal of Physics: Conference Series (Vol. 664, No. 4, p. 042013). IOP Publishing.

[6] Bui, T. (2015). Analysis of Docker security. arXiv preprint arXiv:1501.02967.

DATA SCIENCE MASTER DEGREE AT SAPIENZA UNIVERSITY OF ROME: TIGHTENING THE LINKS TO SPACE SCIENCE

Frank S. Marzano [(1)], Mario Montopoli [(2)], Stefano Leonardi [(1)], Giancarlo Rivolta [(3,4)]

[(1)] Sapienza University of Rome, Rome, Italy
[(2)] National Research Council, Rome, Italy
[(3)] ESA Research and Service Support, Frascati (Rome), Italy
[(4)] Progressive Systems Srl, Rome, Italy

ABSTRACT

The remarkable increase in the volume and complexity of available data and new technologies that have been developed to process them requires a combined multi-disciplinary approach to design an overall strategy aimed at transforming data into useful information. Key ingredients to develop a successful strategy are data manipulation and visualization, large scale computing, statistical modelling, learning techniques, algorithmic thinking. The *Laurea Magistrale in Data Science* is a new international 2-year Master degree taught completely in English at the Sapienza University of Rome since September 2015. Among the offered courses, *Earth Observation Data Analysis* is the link of the Master degree with the Space Science domain opening appealing applications for environmental monitoring and geoinformation retrieval.

Index Terms— Data Science, Master degree, Sapienza University of Rome, Earth observation.

1. INTRODUCTION

With the recent technological advances, we have developed the ability to keep track of massive amount of data. in all sectors of society, business as well as scientific domains, there has been an increased availability of complex data [1]. This trend is leading to a critical need for skilled professionals who can mine and interpret the data. in this respect *Data Science* (or *Big Data*) is an emerging discipline which combines the expertise coming from computer science, statistics and optimization and provides the most advanced tools for understanding, designing and implementing the process of collecti*ng a complex amount of data and* transforming it in useful knowledge. Among these tools we can mention data warehousing, data compression, data visualization and high performance computing, probability models, statistical learning, machine learning, predictive analytics, uncertainty modeling. A well-trained data scientist, capable of deploying these tools into a real-world problem, will have improved chance of job opportunities in all kinds of business companies, innovative information and communication technology (ICT) industries well as public sector, research and international organizations from Earth observation (EO) [2] to Space Science [3], from astronomy to other space and application domains [4].

Data Science technologies are emerging and spreading out with a fundamental impact on how research is carried out, how data are shared and how applications are empowered. *Data Science* is considered as main enabler and facilitator of the Open Science initiative for European Research Area (ERA), recently launched by European Commission (EC). The effective use of *Data Science* technologies requires new skills and demands for the new professional profile of a data scientist. The future data scientists should exhibit knowledge, competencies and skills in data mining and analytics, information visualisation and communication, as well as in statistics, engineering and computer science, and acquire experiences in the specific research or industry domain of their future work and specialisation.

The Sapienza University of Rome has recently approved and started on September 2015 a new international 2-year Master degree, called *Laurea Magistrale in Data Science* (LMDS) and taught completely in English. This Master degree, the first in the Italian university panorama and one of the first in Europe, is a challenging initiative which is indeed supported by research institutions, private companies, small-medium enterprises at national and international level. As a part of the educational program, the *Earth Observation Data Analysis* course is also offered and will be the main link to the Space domain of applications. A short overview of the Sapienza *Laurea Magistrale in Data Science* will be given in the next sections.

2. DATA SCIENCE AT SAPIENZA

The *Laurea Magistrale in Data Science* is a joint initiative within the I3S Faculty of the Sapienza university, combining the expertise of four Departments: i) informatics (DI); ii) Department of Computer, Control and Management Engineering (DIAG); iii) information Engineering, Electronics and Telecommunications (DIET); iv) Statistics

doi: 10.2788/854791

(DSS). This Master program provides a solid and comprehensive preparation to understand and manage big data, including acquisition, mining, management, and statistical analysis. Fig. 1 shows the current web site of the LMDS Master degree (http://datascience.i3s.uniroma1.it).

Fig. 1. Web-site front page of the Master degree *Laurea magistrale in Data Science* in Rome (http://datascience.i3s.uniroma1.it).

The LMDS course, having a total of 120 European credit transfer credits (Ects), is organized in 4 semesters (within 2 years) as follows:

Semester I	
f) Algorithmic Methods of Data Mining and Laboratory	9 Ects
f) Fundamentals of Data Science and Laboratory	9 Ects
f) Statistical Methods in Data Science and Laboratory I	6 Ects
a) Intellectual Property Compet. and Data Protection Law	6 Ects
a) Economics of Network industries	6 Ects
Semester II	
f) Networking for Big Data and Laboratory	9 Ects
f) Statistical Methods in Data Science and Laboratory II	6 Ects
b) Data Management for Data Science	6 Ects
b) Cloud Computing	6 Ects
b) Data Mining Technology for Business and Society	6 Ects
b) Data Monitoring Analysis and Communication	6 Ects
c) Statistical Learning	6 Ects
c) Quantitative Models for Economic Analysis and Manag.	6 Ects
Semester III	
b) Data Privacy and Security	6 Ects
b) Social and Behavioral Networks	6 Ects
b) Signal Processing for Big Data	6 Ects
b) Network infrastructures	6 Ects
c) Optimization Methods for Machine Learning	6 Ects
c) Statistical Methods for official Statistics	6 Ects
d) Efficiency and Productivity Analysis	6 Ects
d) Probability and Stochastic Processes for Data Science	6 Ects
d) Digital Epidemiology	6 Ects
Semester IV	
d) Earth Observation Data Analysis	6 Ects
d) Economics of information	6 Ects
d) Bioinformatics	6 Ects

The Master degree plan is such that fundamental courses are those of group f) for 39 Ects, a total of 6 Ects is chosen within the optional group a), a total of 18 Ects is chosen within the optional group b), a total of 6 Ects is chosen within the optional group c) and a total of 12 Ects is chosen within the optional group a). Finally, 12 Ects are left as open choice to the student plus 3 Ects for ability training and 24 Ects for the final Master thesis.

3. EARTH OBSERVATION AT SAPIENZA

Among the offered courses, Earth Observation Data Analysis (EODS) is the link of the LMDS Master degree with the Space science domain. The EODS module aims at providing a general background on the remote sensing systems for Earth Observation from space-borne platforms and on data processing techniques [4]. It describes, using a system approach, the characteristics of the system to be specified to fulfil the final user requirements in different domains of application. Remote sensing basics and simple wave-interaction models useful for data interpretation are reviewed together with technical principles of the main remote sensors. The course also provides an overview of the most important applications and bio-geophysical parameters (of the atmosphere, the ocean and the land) which can be retrieved. The most important techniques for data processing and product generation, also by proposing practical exercises using the computer, are analysed together with an overview of the main Earth Observation satellite missions and the products they provide to the final user.

Joint seminars with the European Space Research Institute (ESRIN) of the European Space agency (ESA) are also foreseen together with stage opportunities. ESRIN is ESA centre for Earth observation in Frascati (Rome, Italy), and manages the ground segment for ESA and third-party Earth-observation satellites. ESRIN also hosts the Eduspace website, routinely used by thousands of schools, some universities and institutions across Europe and the world. ESRIN also hosts the ESA G-POD, a generic GRID-based processing-on-demand operational environment [5]. G-POD provides the necessary flexibility for building an application virtual environment with accessibility to data, computing resources and results, useful for educational training in Data science. Moreover, as part of ESA's Earth Observation Ground Segment Department, Research & Service Support (RSS) has the mission to provide tools and services that support the user community in EO data exploitation [5]. RSS

provides as well this e-collaboration area to find or exchange information, and share ideas on Earth Observation data exploitation related projects, an ideal environment for practicing EO concepts and tools.

4. CONCLUSION

A summary of the Sapienza *Laurea Magistrale in Data Science* has been given by illustrating its organization and teaching topics. Specialization in Earth observation data analysis is also foreseen by stressing the physical-mathematical modeling of wave-interaction processes and geoinformation retrieval. Strong links of LMDS Earth Observation curriculum with ESA-ESRIN are also envisaged by exploiting G-POD and RSS facilities.

5. REFERENCES

[1] National Science Board, D. Simberloff et al., NSB-05-40, *Long-Lived Digital Data Collections Enabling Research and Education in the 21st Century*, Sept. 2005 (http://www.nsf.gov/pubs/2005/nsb0540/nsb0540.pdf).

[2] S. Nativi, P. Mazzetti, M. Santoro, F. Papeschi, M. Craglia and O. Ochiai, “Big Data challenges in building the Global Earth Observation System of Systems”, *Environmental Modelling & Software*, vol. 68, pp. 1–26, June 2015.

[3] F.S. Marzano and G. Visconti, Eds., *Remote sensing of atmosphere and ocean from space: models, instruments and techniques*, Advances in Global Change Research series, Springer, Dordrecht (NL), ISBN 1-4020-0943-7, pp. 246, 2002.

[4] P. Soille and P.G. Marchetti *Proceedings of the 2014 Conference on Big Data from Space (BiDS'14)*, Publications Office of the European Union, 2014, DOI: 10.2788/1823

[5] P.G. Marchetti, G. Rivolta, S. D'Elia, J. Farres, G. Mason and N. Gobron, “A Model for the Scientific Exploitation of Earth Observation Missions: The ESA Research and Service Support”, *IEEE Geoscience and Remote Sensing Newsletter*, vol. 162, pp. 10-18, 2012.

SERVICES AND USER SUPPORT ON THE FORESTRY THEMATIC EXPLOITATION PLATFORM

Renne Tergujeff[1], Tuomas Häme[1], Oleg Antropov[1], Clive Farquhar[2], Rakesh Prithiviraj[2], Indranil Mustafee[2], Victoria Bennett[3,4], Philip Kershaw[3,4], Philippe Ledent[5], Arnaud De Groof[5], Jarno Hämäläinen[6], Frank Martin Seifert[7], Joost van Bemmelen[7]

VTT Technical Research Centre of Finland Ltd (1); CGI Group Inc. (2); Centre for Environmental Data Analysis, STFC Rutherford Appleton Laboratory (3); National Centre for Earth Observation (4); Spacebel S.A. (5); Arbonaut Ltd (6); European Space Agency (7)

ABSTRACT

The objective of the Forestry Thematic Exploitation Platform (F-TEP) is to change the concept of Earth Observation (EO) in forestry, with a vision of a one-stop shop for forestry remote sensing services for both academic and commercial sectors. Three basic usage scenarios are supported: exploitation, service development and product development. F-TEP is one of the in total six thematic EO Exploitation Platforms of the European Space Agency, developed in parallel projects.

F-TEP is organized as combination of a web portal and a full virtual environment. The F-TEP services cater for users from many segments, with varying familiarity with EO data and processing. Each user group is supported by providing an interface that they feel comfortable working with, and which helps them to achieve their goals for forestry data products with ease. As users grow with expertise and confidence and want to pursue more complex tasks, they can proceed to unlock a new range of functionality. With this model, F-TEP can act as a learning platform as well.

The strategy for user support in F-TEP combines design for easiness of use, good textual and visual support material, strong support for communication and collaboration between peer users, and a responsive helpdesk.

***Index Terms*—** Earth Observation, forestry, remote sensing, data processing, exploitation platform

1. FORESTRY THEMATIC EXPLOITATION PLATFORM

Forests, which occupy about 40% of the global land surface, are the dominant terrestrial ecosystem of the Earth. Deforestation, mainly the conversion of forest land to agricultural land, has been and still is a major land cover change process, particularly in the tropics.

The Forestry Thematic Exploitation Platform (F TEP)[1] is one of European Space Agency's EO Exploitation Platforms that are developed in six parallel projects[2]. The objective of F-TEP is to change the concept of Earth Observation in forestry. It aims at improving the current general practices that are based in working individually in scattered projects. The vision of F-TEP is a one-stop shop for forestry remote sensing services and a federation of users in academic and commercial sectors.

The services offered will provide remote access to pre-processed Copernicus and other satellite data, ancillary data and third party data, together with computing resources and tools to process and analyse it. Users are able to produce data products of interest without having to download the source data, which can be tens of gigabytes per image. The F-TEP also aims to be a place to find information about what is happening in forestry earth observation globally and a place to network and discuss this with colleagues.

F-TEP will provide free access with free software tools and data to everyone interested, as well as added-value services and data based on a number of service levels. It will also offer a self-service environment for the users to develop their own tools and services. This will drastically reduce the burden on data acquisition from miscellaneous sources with varying formats and processing levels.

F-TEP is organized as combination of a web portal and a full virtual environment. The portal offers extensive functionality for data exploitation, including: browsing of image data based on area of interest; viewing previews; uploading reference data; selecting, parameterising and running processing services for the data; and managing the defined processing tasks. Figure 1 presents a single screen from the work-in-progress mock-up of the user interface being designed.

[1] http://forestry-tep.eo.esa.int/
[2] https://tep.eo.esa.int/

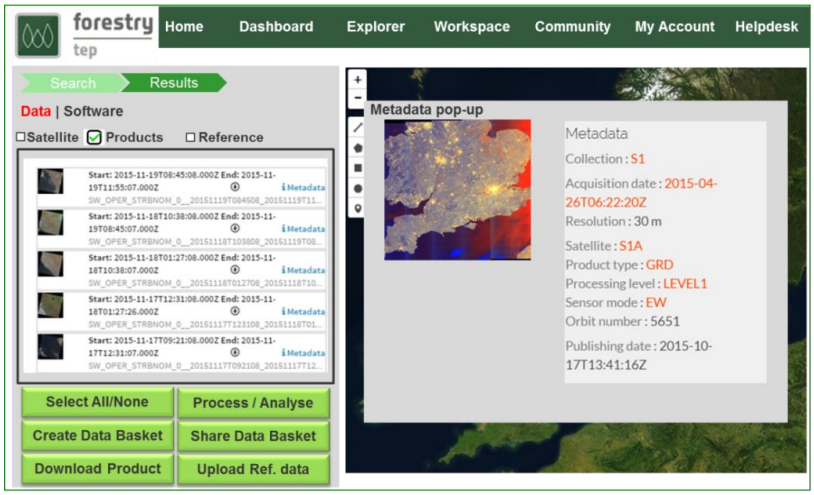

Figure 1. Data selection screen from the F TEP portal interface mock-up, in development.

Through the portal the users can also manage their own data sets and projects and optionally publish their results to other users of the system. A discussion forum and helpdesk are available as well.

The virtual environment, which is also accessed through a web browser, offers full desktop and command line interfaces that enable more advanced processing tasks. It allows efficient access to hosted state-of-the-art toolboxes and the ability to develop and run custom processing algorithms. This approach is suitable for users with advanced skills, who would like to get more control and contribute to the platform development.

It is foreseen that at least the ESA Sentinel toolboxes[3] and the Orfeo ToolBox[4] with the graphical Monteverdi user interface will be offered on free basis. Commercial off-the-shelf software for satellite data pre-processing, ortho-rectification, and radiometric and geometric correction will be made available as well.

Potentially, F-TEP may also be used as a test-bed for methods developed in other projects, such as the ESA projects Innovators III AccuCarbon[5], ESA DUE GlobBiomass[6] and Innovators III SAR for REDD[7].

The F-TEP also fulfils the needs of service providers who make use of the platform to offer services to the platform users. For instance, software providers can offer access to their applications via the F-TEP platform to reach a wide range of audience.

[3] https://sentinel.esa.int/web/sentinel/toolboxes/
[4] https://www.orfeo-toolbox.org/
[5] http://due.esrin.esa.int/page_project153.php
[6] http://globbiomass.org/
[7] http://due.esrin.esa.int/page_project154.php

2. SERVICE SCENARIOS

The Forestry TEP will serve its users through three general usage scenarios: exploitation, service development and product development.

In the *exploitation* scenario, the user employs data, products and tools that already are available on the platform. These tools range from browsing for pre-processed imagery, reference data and existing third party products, to computing forest biomass maps and assessing their accuracy, and finally downloading just the final product for further use, for instance. The exploitation user can be a novice in remote sensing, an institutional user or a scientist who uses remote sensing products in support of research. A user-friendly interface is of central importance to enable effective exploitation for the user.

In the *service development* scenario, the user programs his own service, e.g. for forest biomass estimation, and adds it to the F-TEP environment, with the help of his own tools and tools on the platform. A service established on the platform can then be shared or licensed as Software-as-a-Service, or used by the author himself. This scenario is obviously of added complexity compared to the exploitation scenario.

The *product development* scenario differs from service development in that the final products are computed on the F-TEP platform and distributed to data exploitation users. The distribution can be commercial or offered freely.

3. COMPETENCE NEEDS AND USER SUPPORT

For the success of the F-TEP, acceptance and adoption by the exploitation users is vital. The Forestry TEP will serve users from varying backgrounds, differing in their

familiarity with Earth Observation (EO) data and processing. The users range from academics and EO professionals to novice, non-expert users. A key target user for F-TEP is UN-REDD (the United Nations Programme on Reducing Emissions from Deforestation and Forest Degradation) [1]. Other identified main user groups include: Copernicus core services; Global Forest Observation Initiative (GFOI) and other international programs; national forest inventories; universities and research centres; forest managers; value-adding industry; land use planning and nature conservation agencies; and sustainable development non-governmental organizations. Among these target groups of potential users, an extensive user community survey was performed, and its results have been used in specification of the platform's requirements baseline.

The overall approach employed is to support each user group by providing a view and an interface that they feel comfortable working with, helping them to achieve their goals for forestry data products with ease. The basic interaction model for the simplest routine-like tasks allows the user to easily identify and select the source data of interest and a pre-developed processing or analysis task, such as generation of vegetation indices. With minimal interaction and knowhow the user can achieve an end product with added value.

The initial interface of the platform will be made simple, to present a minimum threshold for an uninitiated user to approach and use the platform. For advanced users, options to open up more demanding functionality are available. As a novice user grows with confidence and expertise and wants to pursue more complex tasks, he too can proceed to unlock new functionality as they become more familiar with the system. With this model, F-TEP can act also as a learning platform.

The strategy for user support in F-TEP combines design for easiness of use, good textual and visual support material, strong support for communication and collaboration between peer users, and a responsive helpdesk. Application of these principles includes: building a portal that is easy to approach and use; preparing a documentation wiki and training videos; and establishing and promoting a discussion forum that encourages interaction among the users.

Collaboration among the platform users is strongly promoted. The users can e.g. work together in generating products and developing services, share them with other users, review products and services shared by others, and discuss them with the authors. To foster collaboration, the platform allows creating groups of users with similar interests, asking questions from forestry experts through the user forum, and sharing processing code and workflow files through provided repositories. Linking with external forestry communities is supported through providing forestry themed news and Twitter feeds.

4. TECHNICAL SOLUTION

The F-TEP platform is a cloud based solution that is being developed to be hosted on JASMIN-CEMS[8], a data analysis facility for the environmental sciences community based in the UK [2]. This provides both a cloud platform and access to a curated archive of EO data, including the latest Sentinel data. Co-locating the F-TEP with the archive provides good network connectivity to relevant data sources.

As a cloud based offering, F-TEP has been designed to utilise a stack of three service layers: Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). The SaaS layer consists of applications delivered over the web to end-users of the platform. The PaaS layer enables application developers to develop and deploy applications in the cloud. IaaS refers to the software interface to configure virtualised abstractions of servers, storage, network and operating systems. The three-layer service stack is shown in Figure 2.

The functional building blocks of F-TEP are introduced in the following:

User Portal acts as an entry point for users of F-TEP and provides the functionality described earlier in this article.

User Management function authenticates and authorizes platform users for their activities. Functionality to foster user collaboration is included.

Core Controller is the heart of the platform as it serves the user requests by delegating the tasks to other functions. It has the complete view of the platform and is responsible for co-ordinating other functions within the platform. For instance, it interacts with the Data Store Manager to locate the data requested by the user.

Accounting Manager is responsible for tracking the resource usage and providing capabilities for pay-per-use model for sustainability, covering EO data, storage, processing power and software usage.

Persistence Manager hosts a working area for the platform. This working area can be considered as a rolling cache for input data and processed output data.

Data Store Manager ingests and retrieves data by interacting with the internal and external data archive.

Resource Manager allocates virtual machines, applications and other resources. It optimizes the resource usage to keep the cost of operation to a minimum level.

Computing Cluster constitutes all the computing nodes responsible for executing the user requests. It represents the actual resource that completes the requested processing tasks.

F-TEP Data Store indexes and archives all the products generated within the platform.

[8] http://www.jasmin.ac.uk

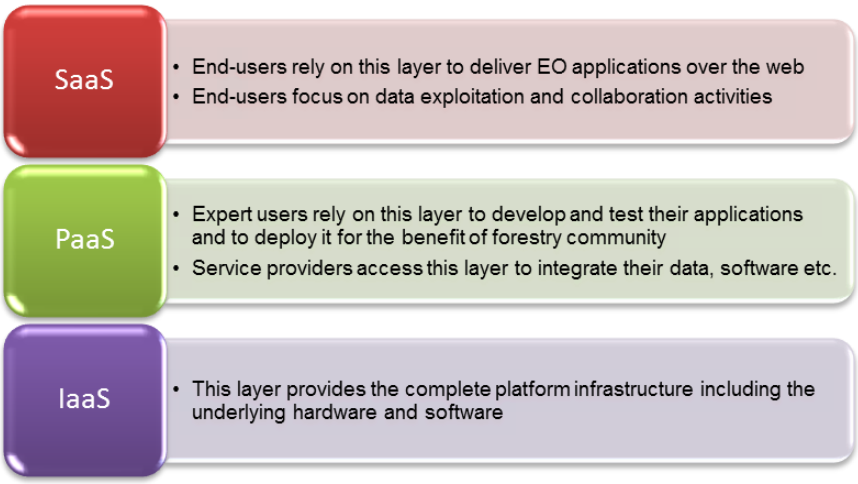

Figure 2. F-TEP platform architecture, designed to offer the software, the platform and the infrastructure as services.

5. THE PROJECT ROADMAP

The Forestry Thematic Exploitation Platform is being developed under a contract of the European Space Agency, and it started as a project in March 2015. It is a collaborative effort conducted by VTT Technical Research Centre of Finland Ltd (coordination, user community, pilots); CGI UK (technical design and implementation); Science and Technology Facilities Council (STFC) (system deployment, JASMIN-CEMS cloud infrastructure and curated data archive); Spacebel (sustainability analysis and planning); and Arbonaut (value adding service development).

The project has finalized the initial design and development phase and entered Phase 2a, which covers detailed design, implementation and setup of the system, services and governance, and preparation of pilot projects.

Two extensive pilots, with national authorities and academic and research institutes as users, will be conducted in 2016-2017: one in tropical and temperate regions of Mexico, and another in boreal forests of Finland. The Mexican pilot concerns monitoring of above ground biomass and quantifying associated carbon stocks for climate change reporting [3][4]. It supports development of the Measurement, Reporting and Verification activity of the REDD+ program (Reduction of Emissions from Deforestation and forest Degradation). It also supports general forest management planning. The Finnish pilot concerns operational forest management and it aims at mapping undesired broadleaved tree shrubs on forest regeneration areas to support their effective maintenance.

The pilots will showcase the usefulness of F-TEP to the forestry user community and validate the developed functionality through actual usage. They also provide a template for later services to follow.

The pre-operational Phase 2b of the project will execute the pilot projects, while also refining the system and services. It is planned that for selected test users, the pre-operational services can be made available around September 2016. For an extended user community the services will come online in early 2017. The candidate pre-operations user group consists of over 30 organizations, 16 of which have been identified as core users.

The Thematic Exploitation Platforms (TEPs) aim to form together an interconnected ecosystem of services, systems and users. Throughout the project, collaboration and alignment with the other TEP projects is supported by a common outreach working group managed by ESA.

6. REFERENCES

[1] UN-REDD Programme Strategic Framework 2016-20, UN-REDD Programme Fourteenth Policy Board Meeting, 20-22 May 2015, Washington, D.C., USA. UNREDD/PB14/2015/III/3.

[2] B.N. Lawrence, V.L. Bennett, J. Churchill, M. Juckes, P. Kershaw, S. Pascoe, S. Pepler, M. Pritchard and A. Stephens, "Storing and manipulating environmental big data with JASMIN", Proceedings of IEEE Big Data 2013, IEEE, Santa Clara, CA, USA, pp. 68-75, 6-9 Oct. 2013. doi:10.1109/BigData.2013.6691556.

[3] T. Häme, Y. Rauste, O. Antropov, H. A. Ahola and J. Kilpi, "Improved mapping of tropical forests with optical and SAR imagery, Part II: Above ground biomass estimation," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 6, no. 1, pp. 92-101, Feb. 2013.

[4] O. Antropov, Y. Rauste, F.M. Seifert and T. Häme, "Selective logging of tropical forests observed using L- and C-band SAR data," Proc. *Intl. Geoscience Remote Sens. Symposium* IGARSS 2015, pp. 3870-3873, 2015.

EDUCATION AND SKILLS: PERSPECTIVES FROM BIG DATA VALUE ADDED ASSOCIATION – BDVA

Florin SERBAN, Anca Liana COSTEA, Mirela CONSTANTIN

TERRASIGNA - 4, Verii Str., 020723, Bucharest, Romania

ABSTRACT

In a data driven economy, handling the huge amount of data and extract the needed and valuable information is a big challenge. Given the unlimited opportunities, the multitude of new creative products and services that can be developed, a large number of actors are involved in big data related activities. They all face similar difficulties, share similar aspects that are to be covered, and, most of all, have the same need of collaboration and constant dialog, between themselves and between them and political / decision making actors. This is why the Big Data Value Association (BDVA) has been created, to boost European Big Data Value research, development and innovation and to foster a positive perception of Big Data Value.

One main mission of BDVA is the development of a Strategic Research and Innovation Agenda. Ensuring the availability of highly skilled people is a key element, covered by BDVA's activities and the SRIA.

***Index Terms*—** Big Data, Strategic Research and Innovation Agenda, dialog, collaboration, education, skills.

1. INTRODUCTION

BDVA was established by 24 founding members, from large and SME industry and research, in 2014 and is the industry-led contractual counterpart to the European Commission for the implementation of the Big Data Value PPP signed in October 2014. Driven by principles like openness, transparency, cooperation, inclusion, fair access, cross domain or cross stakeholders, the association represents a mean by which its members, different actors active in Big Data can become more competitive, can multiply their network of contacts, and become part of the excellence of the science base of creation value from Big Data.

The association is a fully self-financed non–profit organization and, since its establishment, it experienced a constant grow in number of members, starting from the 24 founding ones to more than 120 members (full and associated) in January 2016. Each of its organized events raised the interests of more and more organizations to become part of the initiative.

2. STRATEGIC RESEARCH AND INNOVATION AGENDA - SRIA

One of the main BDVA's responsibilities within the PPP is to draft, to submit to the European Commission and to yearly update a Strategic Research and Innovation Agenda, together with defining and monitoring the metrics - Key Performance Indicators - of the cPPP. Defining the overall goals, main technical and non-technical priorities and a corresponding research and innovation roadmap for the PPP, SRIA addresses multiple dimensions of Big Data, such as data, legal and policy issues, technology, business related aspects, social or education and skills.

3. SPACE ADDRESSED WITHIN BDVA

BDVA is a horizontal network, not dedicated to a specific application area. Therefore, there was a need, raising from its members, that different domains to be clearly defined and addressed. Several domains are already identified and specific working subgroups are created as application areas within one of the major Task Force of the association. Space is among these, and efforts are presently done to define a clear and coherent strategy to address Big Data from Space related issues.

4. EDUCATION AND SKILLS IN BDVA

In the present explosion of available data, one main challenge is to fully exploit this data and to provide the extracted strategic information for a deeper

understanding on environment improving aspects like decision making or productivity. For this, the key represents the human resources, meaning having the people with the necessary skills to develop and use big data technologies to turn it in valuable information. It is the reason why strategic importance is given to education and skills within the BDVA, based on the agreement and understanding that these represent a key element for future leverage of the Big Data potential. A specific Task Force dedicated to education is active within the BDVA, contributing to the dedicated part within the SRIA and implementing different related activities and initiatives. The declared vision for Big Data Value (BDV) in Europe is that in 2020, 'Millions of jobs established for data engineers and scientists, and the Big Data discipline is integrated in technical and business degrees. The European workforce is more and more data-savvy seeing data as an asset.'[1]

In order to tackle the related challenges, data science has become a multidisciplinary area, combining mathematics with computer science, but also with domains like human-technology interaction or business models. Important skills that data scientist will have to acquire in the future have to allow them solving the technological challenges in the 4Vs of Big Data (Velocity, Volume, Variety and Veracity) and to turn the data into actionable information.

5. DATA PROFESSIONALS

Strong points have been identified, such as the existence of innovative technologies and skilled people in several domains, or an important number of high education establishments where people can be educated.

In the same time, weak points have been noted as far as the skills and education in the Big Data is concerned, and these have to do with the lack of specialized education programs for data analysts and not enough skilled people to participate in training programmes.

Opportunities arise in addressing the education related issues by the fact that new areas always come up to be explored, and that innovation is seen as key aspect for trainings in BDV.

All the initiatives to be implemented have to concentrate on the main identified threads, namely the continuous lack of skilled professionals and the brain drain from Europe to other regions.

To address the issues above, not only that data science aims to collect, analyze and interpret data from various sources, but to turn it in operational information and to offer a comprehensive understanding of the environment and context. Therefore, different educational programmes seem to be required next to developing technical skills, emphasizing on application, context and awareness of multiple challenges, focusing on societal impact and relevance and considering the horizontal approach between several domains.

Taking into account all these aspects, BDVA Task Force dedicated to education, have made a classification of the targeted job profiles needed to be developed, in six independent job profiles.

1) Data scientist – represents one of the key targeted profession, having expertise in big data technologies and data analysis in order to gain first rough insight from available data;
2) Data analyst – using the existent big data infrastructure it covers areas such as data mining, machine learning, statistical based data analysis etc;
3) Data engineer – addresses big data technologies and use these to build big data management ecosystems to tackle challenges in terms of velocity, volume, variety and veracity;
4) Visual analytics engineer – applying intuitive techniques, it allows to understand and summarize information comprised within large amount of data, in a short period of time;
5) Process miner – data mining engineer – looks into process event data, to create an overview of different processes, in order to optimize them;
6) Domain analyst – is the profession providing, on top of data analysis, the added value for a specific targeted domain.

All these identified job profiles are easily applicable in big data from space related activities, therefore paying attention and educating these types of specialists would be only bring value in the future in exploiting at maximum the existing and to be available data.

Different European countries have national educational programmes for Data Mining, Analytics, Big Data, and Data Science: Austria, Belgium, Estonia, France , Germany, Greece, Hungary, Ireland, Netherlands, Portugal, Russia, Spain, Sweden, UK. In addition, there are European joint programmes such as:

- Erasmus Mundus Master Course in Data Mining and Knowledge Management (DMKM) [1], based in six universities in four countries: France, Romania, Italy and Spain;
- EIT ICT Data Science Masters Programm[2], at TU/e Eindhoven, UNS Nice Sophia-Antipolis, UPM Madrid, KTH Stockholm, TUB Berlin, and PoliMi Mailand;
- IT4BI (Information Technology for Business Intelligence)[3], a two-year, English-language course in BI, given by 5 leading universities in Belgium, France, Spain, and Germany.

In order to promote the industry interests regarding the needs in human resources, BDVA is following these programmes and is, in addition, supporting present initiatives, such as:

- the EDSA project (e.g. European Data Science Academy - EDSA[4]), aiming at analysing the required sector specific skillsets for data analysts across the main industrial sectors in Europe and developing modular and adaptable data science curricula to meet these needs;
- the EDISON project[5], having as declared aim to accelerate the process of establishing the profession of Data Scientist.

6. CONCLUSIONS

Focusing on enhancing cooperation and targeting the excellence of the science base of creation of value from Big Data, BDVA sets the ground for extensive experience and skills acquiring by working on projects in specific technical priority areas. As follow up, skill development requirements can be identified and, later on, addressed by education institutions and education providers. Therefore, as result of BDVA activities, new educational programmes in data science and data engineering can be established, professional courses to re-skill / up-skill the current work force can be initiated, and curricula updates can be implemented, so that the required skills to be available to support and enrich the know-how.

7. REFERENCES

[1] Big Data Value Europe, 'European Big Data Value Strategic Research & Innovation Agenda', http://www.bdva.eu/sites/default/files/europeanbigdatavaluepartnership_sria__v1_0_final.pdf , January 2015

[1] http://www.em-dmkm.eu/
[2] http://www.masterschool.eitdigital.eu/programmes/data-science/
[3] http://it4bi.univ-tours.fr/it4bi/
[4] http://www.edsa-project.eu
[5] http://www.edison-project.eu

THE D4SCIENCE INFRASTRUCTURE TO SUPPORT ACADEMIC COURSES

Nadia NARDI(1) - Giampaolo CORO(2) - Pasquale PAGANO(2)

(1) Engineering - Ingegneria Informatica SpA, Rome, Italy
(2) CNR - National Research Council, Pisa, Italy

ABSTRACT

To be able to promote and exploit Data, such as Space Data, for societal innovation and education, a clever space is required, rather a training environment, where everything needed to successfully support scientists, trainers, trainees and students needs to exist. The D4Science e-Infrastructure[1] does just this.

D4Science supports the operation of a large set of diverse Initiatives, Communities of Practice, and Projects by offering Virtual Research Environments (VREs) and Services. It aims to offer trainees remote computational facilities to execute experiments and provide facilities to visualize datasets and perform analysis. D4Science supports academic courses by taking a comprehensive approach to what it means to boost education and facilitate knowledge bridging between research and innovation.

***Index Terms*—** Education, D4Science, BlueBRIDGE, VRE, e-Infrastructure.

1. THE EDUCATION LANDSCAPE

Education, together with research and innovation, is one of the priorities of EU investments to boost jobs and growth, as highlighted by the new European Commission's political guidelines [1]. Its importance is particularly relevant in complex scientific contexts, such as observation data collected by space-borne and ground-based sensors, which have an important impact on societal and economical strategies. Education is often performed in university courses and through focused training events and workshops. These are organised by specialised scientific institutes, and address scientists at various stages of their careers.

Unfortunately, trainees hardly have in their academic careers the possibility to work with an operational data platform of any type. Experimentations and academic project works are often done on simulated environments or sample data sets. The role of e-Infrastructures and the European Research Area[2] (ERA) should be regarded also as a means to support scientists during their initial careers in Higher Education Institutions (HEI) by allowing students to explore such operational environments and gain hands-on experience on real scientific challenges.

2. THE D4SCIENCE E-INFRASTRUCTURE

D4Science [2] operates a Hybrid Data Infrastructure that: (a) is designed to integrate "resources" from other e-Infrastructures and commercial vendors by using a system of systems strategy, (b) offers unified access to integrated resources by abstracting from the underlying e-Infrastructures, and (c) is called to serve different Communities of Practices.

Born in 2004 from a series of e-Infrastructure projects (namely DILIGENT, D4Science I and II), it addresses the lack of a data platform to support the manipulation and management of Data of any nature by scientists. From Digital Libraries in support of Environmental Conferences (ESA scenarios in DILIGENT and D4Science) to the computational support of world fishery data (FAO scenarios in D4Science I and II), the D4Science platform – based on the gCube system[3] – has proven its suitability to support Science and Education challenges.

D4Science connects +2000 scientists in 44 countries, integrating +50 heterogeneous data providers, executing +13,000 models and algorithms/month; providing access to over a billion quality records in repositories worldwide, with 99.7% service availability. D4Science hosts +40 Virtual Research Environments (VREs) [3], [4], [5] to serve the biological, ecological, environmental, mining, and statistical communities' worldwide.

D4Science manages several organizations, mapped as Virtual Organizations (VOs), each of which manages several sub-communities, mapped as VREs [6]. VOs and VREs are hereafter called operational context or simply context.

Each VO and each VRE has:

- At least one manager who has the responsibility to manage the users belonging to his/her contexts.
- One key used to encrypt/decrypt sensible information on the different enabled storage devices and to

[1] d4science.org/home
[2] http://ec.europa.eu/research/era/index_en.htm
[3] www.gcube-system.org

encrypt/decrypt access credentials stored in the Wallet service.

- A subset of the resources (applications, platforms, datasets, hosting nodes, storage devices) that can be exploited by the users belonging to that context.

Users and resources can be assigned and withdrawn at any moment to any operational context. Changes are immediate without any human intervention. Users can belong to several contexts, and resources can be shared across operational contexts.

The D4Science e-Infrastructure is logically composed of seven areas: Enabling Layer, Spatial Data Infrastructure, OLAP Infrastructure, Storage Infrastructure, Computing Infrastructure, Analytical Infrastructure, and Registries. Each of the areas exploits several heterogeneous technologies integrating more than 500 software components. All of them are made interoperable through the exploitation of gCube Mediators while common registries make resources and data easy to discover. The publication of resources and data profiles is performed by Mediators that, in a transparent manner to the federated resources, perform harmonization and publication realizing a unified view of the infrastructure resources.

D4Science not only offers services that facilitate the discovery of resources and data, it offers services for seamless access and analysis to a wide spectrum of data including biological and ecological data, environmental and geospatial data, statistical data and semi-structured data from multiple authoritative data providers and information systems.

Services offered include: facilities to manage distributed resources efficiently (monitoring, accounting, alerting, failure-recovery, scale-in/out); facilities for end-users according to the Science 2.0 paradigm (web portal, workspace, messaging system, social networking) [7], [8]; facilities for efficient data management (data and metadata harmonization, metadata generation, data discovery and access); facilities for visualizing large datasets (GIS maps, graphs, tables); facilities to perform data analysis (Artificial Intelligence algorithms; forecasting and signal processing methods); facilities to efficiently perform computations on large datasets (cloud computing and high-throughput computing). These services can be exploited both via web-based graphical user interfaces and web-based protocols for programmatic access, e.g. REST and SOAP. This offering complements specific and community-specific applications and helps to integrate all of them in a unique VRE.

All facilities are completely integrated and users/services

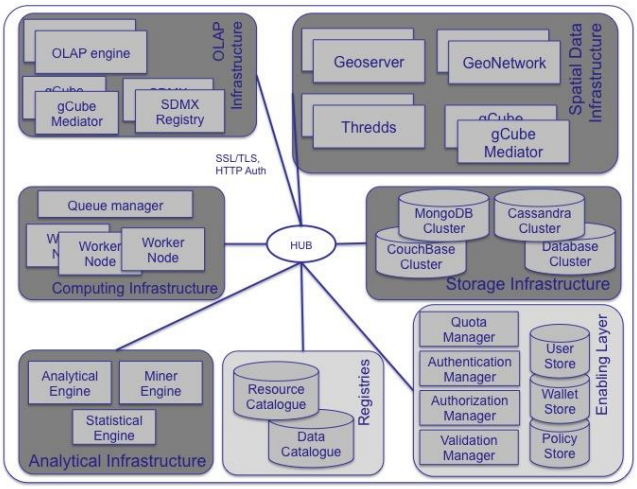

Figure 1- D4Science Architecture

are authorized to exploit them under the limit expressed by the policy in turn expressed by the VO/VRE Manager in the operational context of the call.

The authentication of a user is based on Userid and password while the identity of a service is based on the host certificate. The Authentication service matches information received with data contained in the User Store assigned to the operational context and returns the response. The response could either be a Security Assertion Markup Language (SAML) assertion containing authentication information and user attributes, or an identity token giving permission to get authentication information and user attributes if the authentication succeeds, otherwise an error message is shown. The module is totally decoupled from the rest of the architecture in order to avoid any conflict with the federated resources.

Shibboleth[4] is used for obtaining web browser Single sign-on (SSO) for portal-based authentication, and is integrated with the Enabling Layer in order to obtain Federated Authentication for services. Federated Authentication is here a type of Authentication based on a signed SAML Token indirectly referenced in a SOAP message, containing the authentication information of a user authenticated in the federated and trusted domain. Three exploitation models are available to meet the different user needs:

- As a user, you can securely preserve, access (from anywhere), and confidentially share your data and exploit one or more of the existing applications;
- As a group of users, you can create your own virtual environment and add the applications you need;
- As a community, you can manage the communities and offer them a set of virtual environments.

[4] https://shibboleth.net/

3. THE CHALLENGE

The current situation in different educational institutions is very fragmented, heterogeneous and characterised by a lack of technological support for scientists, trainers, trainees and students. All the steps in preparing courses and workshops require manual work, which is repeated each time a new course is held. These steps often need the training environment to be reset and reconfigured for the next course. For example, available data processing services and models come under heterogeneous programming languages, which usually require installing complex software on users' computers. Furthermore, executing models on users' data usually requires a long phase of data preparation and powerful hardware, not always available in the laboratories of the institutions. Sharing data, parameters and results of experiments is difficult and collaboration between trainers and trainees is limited to the duration of the course. The requirement here is an instrument that bridges the scientific and training environments.

4. THE AMBITION

The services D4Science offers and the VRE-based approach can be applied to different scientific and industrial contexts. For example, in the BlueBRIDGE project[5], D4Science exploits its environment supporting practical uptake of scientific knowledge (in workshops and academic courses), by leveraging on its capacity offered. These specific VREs for education offer services enabling collaboration and integrated access to digital research resources, cross-disciplinary and cross-community tools, data and services. On one hand, it supply's courses and workshop attendees with tools to easily discover and access a rich set of structured and high quality data. By using this virtual and ubiquitous workspace, attendees can exploit a large variety of always up-to-date models to first analyse, compare and share resulting datasets, and then to discuss analytical processes. On the other hand, the effort spent by trainers to prepare the environment for these courses is reduced. This largely contributes to boosting training programmes, giving them a new volume and a new thematic and geographic reach. In addition, service-based VREs provide facilities for managing data, parameterizing models and providing standard access to them. Furthermore, they offer the computational capacity required to run model simulations. Enabling transparent access to remote computational facilities for executing experiments helps trainees who do not have access to powerful hardware, but have decent bandwidth, to process huge datasets. This facility is crucial especially in developing countries, where classes can be granted Internet access but not powerful hardware. VREs can be opened before and remain open after the events have taken place. In this way, course attendees can prepare themselves beforehand and, work also after the course, continuing experimentation on their own datasets, as in the case of a student working on degree theses.

In order to strengthen communication between trainees and trainers and also to facilitate collaboration among trainees, the VREs also offer social networking and sharing facilities in accordance with the Science 2.0 vision. These facilities simplify exchange of opinions, models, datasets and results about their research.

5. USE CASE

D4Science facilities can concretely help trainers to setup and manage a course. Since a VRE collects individuals that are interested in a certain domain, it can also be used as a support for a class of attendees (e.g. students, scientists etc.). The D4Science VRE creation tools allow adding a description to a new VREs and selecting the e-Infrastructure resources (applications, databases, data sets, etc.) that the VRE services and participants require.

Once the VRE has been created, D4Science offers the VRE manager (i.e. the trainer) web tools to invite attendees to subscribe to the VRE. Each subscribing attendee is automatically given a private Workspace area, an online file system on which he/she can upload private files. This is the main tool for data sharing offered by D4Science and the trainer can base course material distribution on this facility. In particular, VRE participants are automatically provided with a special folder (VRE folder) that is shared among all the participants. Data either uploaded or saved in that folder are automatically shared with the other participants. Outside of that folder, one attendee's data are private, but the D4Science Workspace gives the user the option to share them with selected colleagues.

The trainer usually creates a course structure under the VRE folder, where course material is uploaded and organized. The D4Science sharing mechanisms automatically notify the VRE participants as soon as new material is available. The VRE folder can be structured also before inviting the course participants to join the VRE.

Through the VRE creation mechanisms which D4Science services and Web applications attendees can use is defined. Among these, the trainer can select data mining algorithms [9], domain specific models (e.g. stock assessment, ecological models etc.) and data preparation tools [10] that are compliant with the topics of the course. Every D4Science service is have mechanisms and interfaces to save data in the Workspace. For example, the result of an

[5] BlueBRIDGE (bluebridge-vres.eu) receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 675680.

experiment, the input data and parameters can be saved as files on the shared file system. This allows attendees to communicate, collaborate and also to produce results for the trainers, as a demonstration of their learning achievements. In this phase, having D4Science services running the computations, reduces the time to configure attendees' experimental setup and facilitate the production of results in a timely manner. Attendees may have folders on the Workspace shared with the trainer, either prepared by the trainer before the course or created by the attendees. When attendees save data on that folder, the trainer is automatically alerted. This interaction modality allows, for example, attendees to do homework after the course and eventually send the results of the homework/exam to the trainer.

Finally, a set of social networking facilities is integrated with the D4Science VREs. These allow users, for example, to post information or to subscribe to other posts. All the posts are automatically sent via email to the VRE participants. This allows using posts as mailing lists, leveraging trainers' duty to create specific communication means for courses.

The here described use case is still valid if a course requires organizing attendees in working groups. The scalability of the D4Science solution allows trainers to share information with groups of people as well as with one single user. In this view, D4Science implements a multi-granular approach to teaching. Up to now, this method has been applied to several courses managed by the D4Science team in European Universities [11] and is going to be applied also to courses for scientists in the domain of marine science.

6. CONLUSIONS

The exploitation of Space Data through European Research e-Infrastructures for Societal innovation includes the need to create the optimal workspaces to create and consume academic courses. Education Needs must be given the right importance and effort to continuously improve the experience for both trainer and trainee.

D4Science, in-line with Science 2.0, makes state-of-the-art technologies and infrastructures accessible to education, ultimately increasing information-sharing and collaboration, also by giving easy access to data and models. Furthermore, D4Science promotes the reuse of science (new models, new data) in different contexts. It facilitates new models/data to be used not only to address scientific challenges, but to educate and train new scientists.

5. REFERENCES

[1] Juncker, J.-C. (2014) A New Start for Europe: My Agenda for Jobs, Growth, Fairness and Democratic Change. Political Guidelines for the next European Commission. http://ec.europa.eu/priorities/docs/pg

[2] Candela, L., Castelli, D., & Pagano, P. (2009, July). D4Science: an e-Infrastructure for Supporting Virtual Research Environments. In IRCDL (pp. 166-169).

[3] Candela, L.; Castelli, D., Pagano, P. (2013). Virtual Research Environments: An Overview and a Research Agenda. Data Science Journal, Volume 12, July 2013, pp GRDI75-GRDI81 doi: 10.2481/dsj.GRDI-013

[4] Allan, R. (2009) Virtual Research Environments: From Portals to Science Gateways. Oxford, UK: Chandos Publishing.

[5] Carusi, A., & Reimer, T. (2010) Virtual Research Environment Collaborative Landscape Study. JISC.

[6] Candela, L.; Castelli, D.; Manzi, A.; Pagano, P. (2014). Realising Virtual Research Environments by Hybrid Data Infrastructures: the D4Science Experience. International Symposium on Grids and Clouds (ISGC) 2014, Proceedings of Science PoS(ISGC2014)022

[7] Assante, M., Candela, L., Castelli, D., Mangiacrapa, F. & Pagano, P. (2014). "A Social Networking Research Environment for Scientific Data Sharing: The D4Science Offering," The Grey Journal, Vol. 10, Number 2, 2014

[8] Assante, M.; Candela, L.; Castelli, D.; Manghi, P.; Pagano, P. (2015). Science 2.0 Repositories: Time for a Change in Scholarly Communication. D-Lib Magazine, Vol. 21, Number 1/2 doi: 10.1045/january2015-assante

[9] Coro, G., Candela, L., Pagano, P., Italiano, A., & Liccardo, L. (2015). Parallelizing the execution of native data mining algorithms for computational biology. Concurrency and Computation: Practice and Experience, 27(17), 4630-4644.

[10] Candela, L., Castelli, D., Coro, G., Lelii, L., Mangiacrapa, F., Marioli, V., & Pagano, P. (2015). An infrastructure-oriented approach for supporting biodiversity research. Ecological Informatics, 26, 162-172.

[11] D4Science-based courses at University of Pisa: http://phd.dii.unipi.it/formazione/item/281-gianpaolo-coro,-istituto-di-scienza-e-tecnologie-

HIGH-PERFORMANCE COMPUTING FOR SOIL MOISTURE ESTIMATION

S. Elefante[1], W. Wagner[1], C. Briese[2], S. Cao[1], V. Naeimi[1]

[1] Department of Geodesy and Geoinformation, Vienna University of Technology, Vienna, Austria
[2] EODC Earth Observation Data Centre for Water Resources Monitoring, Vienna, c/o University of Technology, Vienna, Austria

ABSTRACT

The aim of this work is to study the performance of a Synthetic Aperture Radar (SAR) processing chain for soil moisture retrieval using a supercomputer (Vienna Scientific Cluster 3) and the capabilities of the Earth Observation Data Centre (EODC). Different processing tests were designed and performed using three benchmark data sets: the whole ENVISAT ASAR Global Mode archive, part of the ASAR Wide Swath Mode archive and one large Sentinel-1 data set. The experiments showed the feasibility of processing large amount of earth observation data using the services offered by the EODC supercomputing environment.

Index Terms— Big data, parallel computing, HPC, high performance computing, soil moisture

1. INTRODUCTION

The availability of large amounts of Synthetic Aperture Radar (SAR) data stemming from the completed ENVISAT satellite mission and from the Sentinel-1 satellites constellation represents a major opportunity for the worldwide scientific community. Earth observation satellite data are widely employed for many operational and scientific applications. One of the chief applications of Sentinel-1 over land is to monitor surface soil moisture that is regarded as a key parameter for flood forecast and numerical weather prediction systems. [1] [2]. However, the processing of such a massive data amount can be executed only by following the well-established paradigm of distributed computing. The processing of the whole ENVISAT archive and of the Sentinel-1 data in a time frame that allows researchers to make scientific conclusions within a reasonable time span can occur only by using computing environments that fully support parallel processing. The high performance computing (HPC) system represented by the Vienna Scientific Cluster 3 (VSC-3) [3] suits to this aim well. The VSC-3 is a large distributed platform installed in summer 2014 in Vienna. Due to the massive data volume, the accessing and downloading of the raw data is also not a trivial issue. To effectively deal with these challenges, the paradigm is to bring the processing algorithms close to the data as well as to the computing resources. The Earth Observation Data Center (EODC) [4] provides an effective answer to this paradigm as it is a private-public partnership aiming at delivering a collaborative cloud infrastructure for archiving, processing, and distributing EO data. Through multi-national partnerships from science, the public and private sectors, users can get direct access to Sentinel data storage and running data-intensive geoscientific models. One of the services offered by EODC is the capability of performing data processing through VSC-3. EODC environment is, therefore, a comprehensive infrastructure in which users have the possibility to directly access EO data, process them with their own algorithms and retrieve the final products.
The aim of this work is to assess the possibility of processing Big Data for remote sensing applications with a particular attention to the TU Wien soil moisture retrieval processing chain. The results of five case studies, which were performed by using the EODC infrastructure and the VSC-3 high performance computing platform, are shown.

3. TU WIEN SAR TOOLBOX AND PRODUCTS

The TU Wien SAR Geophysical Toolbox (SGRT) is a software package developed to process Envisat and Sentinel-1 Synthetic Aperture Radar (SAR) data. The SGRT is written in the Python programming language and includes some external software modules, in particular ESA's Sentinel-1 toolbox (S1TBX) for SAR data geocoding, radiometric corrections and calibration [5]. The SGRT program is split into three main processing components to ease workflows traceability and reproducibility: pre-processing, model parameter extraction, and data production. During the pre-processing step, the SAR data are first calibrated and georeferenced using S1TBX. After some quality checks and data conversions and corrections, the geocoded SAR scenes in orbital image format are resampled to the TU Wien predefined fixed, planar grid called *Equi7 Grid* that consists of seven continental grids in Azimuthal Equidistant projection. The TU Wien Equi7 Grid is designed to minimize the oversampling rate of the high resolution satellite data globally, while keeping its structure simple [6]. The output of the pre-processing step are time series of the

terrain corrected and georeferenced SAR data resampled to the respective continental Equi7 subgrid.
The SGRT is aimed to be easily adapted for new products and workflows. As a result, a wide range of algorithms can be implemented within the SGRT framework for generating various products. One of the main SAR derived products developed by TU Wien is the Surface Soil Moisture (SSM) product.

3. VIENNA SCIENTIFIC CLUSTER

The Vienna Scientific Cluster-3 (VSC-3) is an advanced HPC system and consists of 2020 nodes, each equipped with 2 processors (Intel® Xeon® Processor E5-2650 v2 from the Ivy Bridge-EP family) and internally connected with an Intel QDR-80 dual-link high-speed InfiniBand fabric. The Scientific Linux release 6.6 is installed on each node as operative system. The Simple Linux Utility for Resource Management (SLURM), that is an open source, fault-tolerant, and highly scalable cluster management and job scheduling system for large and small Linux clusters, is installed as middleware [7]. SLURM organizes the access to the computing nodes through the management of a queue of pending work and provides a framework for executing and monitoring jobs. The parallel cluster file system BeeGFS (formerly FhGFS) [8] is installed on the VSC-3 distributed volume which is constituted by 360 spinning disks connected through around 160 Gb/sec bandwidth (evaluated experimentally). The satellite data are stored in 93 NFS disks (83 x 4T, 10 x 8T) of the EODC data archive. They are connected through high-speed InfiniBand (40 Gb/sec) to the VSC-3 system.

4. CASE STUDIES

Five case studies were executed to evaluate the VSC-3 performance in processing SAR data sets of different size and spatial resolution of the images. Table 1 indicates specifications and processing results of the five case studies. In all case studies, the SGRT pre-processing workflow was selected for the evaluation which is the most time-consuming part of the SAR level-1 data processing. The workflow comprises calibration, orbit corrections, geo-referencing including terrain correction using a digital elevation model (DEM), image-edges' noise removal, data conversion and encoding, and resampling to Equi7 Grid. Three representative sample data were considered as benchmark: Envisat ASAR Global Mode (ASAR GM), ASAR Wide Swath (ASAR WS) and Sentinel-1 IW GRDH. The SGRT preprocessing is utilizing S1TBX (version 1.1.1) for georeferencing and calibration which has multi-core parallel processing capability. The SAR data are acquired in form of image slices and can be independently processed as the elaboration of one image does not require any information from the other image slices processing. This simplifies parallelization of the processing that can be implemented through the run of many independent jobs. However, the large number of data files and considerable different data file sizes from a few MBs in case of ASAR GM (in case-1) to 1-2 GBs of Sentinel-1 IW GRDH (in case-3 and case-4) give rise to different challenges when they are processed on any HPC system.
The full archive of the Envisat ASAR GM data, 189,621 relatively small files (1 to 73 MB) with a total size of about 1.6 TB and a part of the ASAR WS archive including images acquired over mid-latitudes of the Earth), 31,199 files (12 – 692 MB) with a total size of about 5.4 TB were successfully processed during case-1, case-2 and case-3 (Table 1). In all these three cases, the large number of relatively small files (more than a million in the case of ASA GM) made it burdensome to save the output directly in the BeeGFS distributed volume. The storage device was not capable to simultaneously perform the I/O operations that were needed to fulfill all the computing nodes requests. All attempts for storing the output data directly on distributed volume failed due to fatal stucks of the nodes. Thus, the images were divided in small groups of 8 and 2 images per job for ASAR GM and WS, respectively. Such a job file was, afterwards, sent to a single computing node, the output was temporary cached in the memory available at the node and only at the end of the processing the output was copied on the persistent BeeGFS storage volume. The selected number of images, that was assigned to each group (8 and 2), was a constraint imposed by the available RAM on each node. In total, an array of 23,703 and 15,600 independent jobs were submitted to the SLURM middleware respectively for ASAR GM and ASAR WS datasets processing.
As shown in Table 1, the same dataset was processed in case-2 and case-3 as these two tests were ad-hoc designed to further identify the processing chain bottleneck. The number of nodes that were provided by the SLURM middleware during the overall ASAR WS processing is shown in Figures 1 and 2 for case-2 and case-3, respectively. A maximum of approximately 400 and 600 nodes was assigned during the two processing. Such a number guarantees a significant data throughput between storage device and nodes. For instance, the processing of case-1 started in the night of the 4th December 2015 and ended early in the morning of the 9th December 2015. The peak of assigned nodes was reached during the night between the 8th and 9th of December. The average processing time was 5.65 and 2.39 MB/sec for case-2 and case-3 respectively. Figures 3 and 4 show the elapsed times with respect the data file size; while in case-3 the elapsed times have almost a processing linear trend, case-2 is characterized by a large dispersion meaning that the system got stuck quite a few times. The difference in the two case studies is that, although both tests cached the output, only the case-3 cached also the input images. We hypothesize that the reason for this is that the SGRT and S1TBX keep

continuously reading and writing data on disk. Such a programming strategy can cause either bandwidth saturation or storage volumes failing when many nodes are simultaneously active. Copying the input image files to the node-cash could solve the issue as during the processing, the node has a direct and fast access to the cache without incurring the risk of saturating the storage volume and bandwidth. However, further studies are needed to confirm that hypothesis and to exactly identify the reason of such a different behavior.

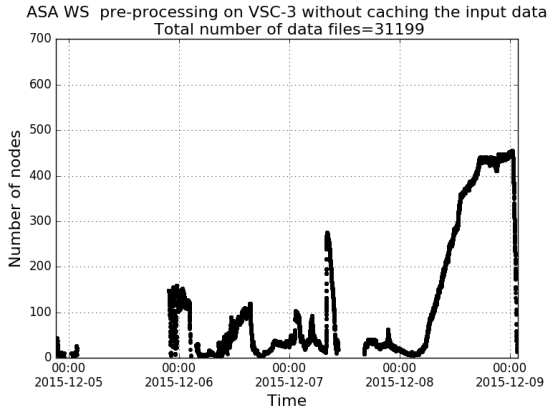

Figure 1. Nodes assignment during case study 2

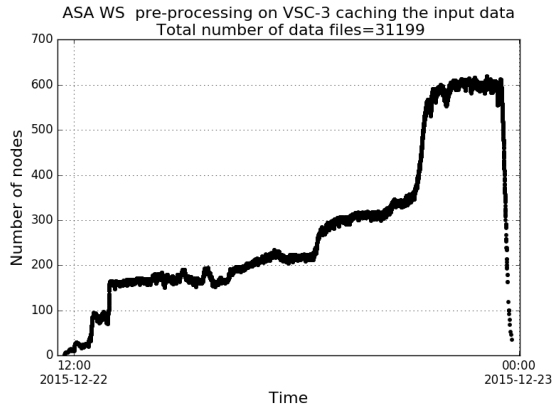

Figure 2. Nodes assignment during case study 3

The case studies 4 and 5 were designed to identify the optimal S1TBX configuration in terms of running cores within a node. The same dataset of Sentinel-1 data consisting of 1,075 data slices of Europe were considered. The input image was immediately cashed in the node to ensure a fast access to the data to SGRT and S1TBX. The output directory was also set to be the cash, and only at the end of the processing the output data were copied in the persistent parallel disk. Due to the size of Sentinel-1 data files and the limited RAM available at each node, only one image was processed at each single nodes. A total of 1,075 independent jobs, equal to the number of images, were therefore submitted. The queue system provided by the SLURM middleware has made a maximum of nearly 396 and 182 nodes simultaneously running. Table 1 shows the average processing times, for case-4 is 2.69 MB/sec in contrast of 2.83 MB/sec for case-5. Case studies 4 and 5 showed that increasing the number of cores used by S1TBX did not improve the performance.

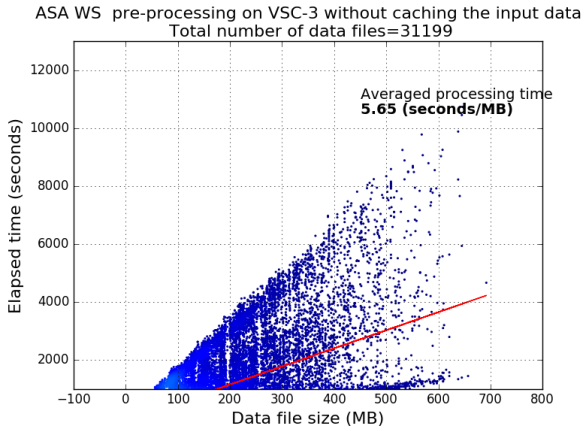

Figure 3. Average processing time for case study 2

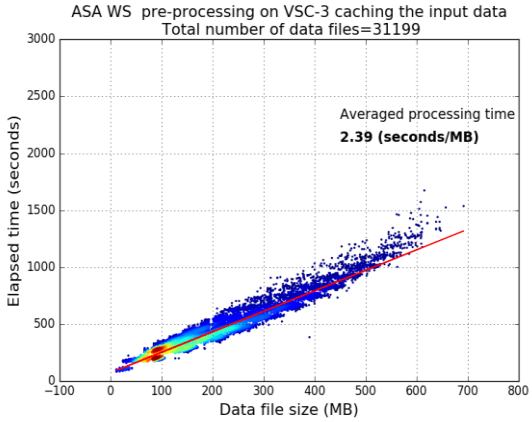

Figure 4. Averaged processing time for case study 3

5. CONCLUSIONS

This paper showed the feasibility of processing large EO data sets at the EODC high-performance platform. Different tests involving the processing of the whole ENVISAT ASAR GM archive, part of the ASAR WS archive and one large Sentinel-1 data sets were performed. The experiments showed that EODC is a comprehensive infrastructure in which users can bring their own algorithm close to both EO data and use the powerful computational resources offered by VSC. Distributed computation can be, therefore, the tool to open

novel roadmaps for addressing new research questions and challenging operational applications.

6. ACKNOLEDGEMENTS

This study was supported by the Austrian Research Promotion Agency (FFG) through “P4EODC - Preparing for the Initial Services of the EODC” (project number 36786), “WetMon - Enabling an operational Sentinel-1 wetland monitoring service for the EODC” (project number 848001) and by the EUMETSAT through “Satellite Application Facility on Support to Hydrology and Water Management (H-SAF)”. The authors would like to express their gratitude to B. Bauer-Marschallinger, A. Dostalova, S. Hahn, S. Hasenauer and to H. Thüminger for their valued and continued support.

7. REFERENCES

[1] Hornacek, M., Wagner, W., Sabel, D., Truong, H.L., Snoeij, P., Hahman, T., Diedrich, E. and Doubkova, M., “Potential for High Resolution Systematic Global Surface Soil Moisture Retrieval via Change Detection Using Sentinel-1”, *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2012, 5 (4), 1303-1311.

[2] Pathe, C., Wagner, W., Sabel, D., Doubkova, M. and Basara, J.B., “Using ENVISAT ASAR Global Mode data for surface soil moisture retrieval over Oklahoma, USA”, *IEEE Transactions on Geoscience and Remote Sensing*, 2009, 47 (2), 468-480.

[3] Vienna Scientific Cluster 3. Available from <http://vsc.ac.at/systems/vsc-3/>

[4] Briese, C., Wagner, W., Boresch, A., Federspiel, C., Aspetsberger, M., Hasenauer, S., Mücke, W., Hoffmann, C., Wotawa, G., “Challenges in the exploitation of big earth observation data”, *Proceeding of the 2014 conference on Big Data from Space*, Frascati (IT)

[5] https://sentinel.esa.int/web/sentinel/toolboxes/sentinel-1

[6] Bauer-Marschallinger, B., Sabel, D., Wagner, W., 2014, "Optimisation of Global Grids for High-Resolution Remote Sensing Data", Computers & Geosciences, ISSN 0098-3004, Available online 11 July 2014.

[7] SLURM. Available from <http://slurm.schedmd.com/>

[8] BeeGFS. Available from <http://www.beegfs.com/content/>

Case study	case-1	case-2	case-3	case-4	case-5
SAR data product mode	ASAR GM	ASAR WS	ASAR WS	S-1 IW GRDH	S-1 IW GRDH
Spatial resolution	1 km	150 m	150 m	20 m	20 m
Grid spacing	500 m	75 m	75 m	10 m	10 m
Total number of data files	189,621	31,199	31,199	1,075	1,075
Number of images for job / Total Number of jobs	8 / 23,703	2 / 15,600	2 / 15,600	1 / 1,075	1 / 1,075
Input data file size range	1 - 73 MB	12 - 692 MB	12 - 692 MB	0.8 – 1.7 GB	0.8 –1.7 GB
Total input data files size	1.579 TB	5.401 TB	5.401 TB	1.2 TB	1.2 TB
Max. number of simultaneous running nodes	417	454	612	396	182
Number of cores used by Sentinel-1 Toolbox	4	8	8	8	16
Input data cashing on node	False	False	True	True	True
Output data cashing on node	True	True	True	True	True
Processing time range (seconds/MB)	0.8 - 90	1.1 – 22.1	1.1 – 3.1	1.5 – 5.6	1.5 – 5.8
Averaged processing time (seconds/MB)	9.18	5.65	2.39	2.69	2.83
Elapsed time including SLURM queueing	≈ 3.5 days	≈ 4 days	≈ 8 hours	≈ 3.5 hours	≈ 15 hours

Table 1. Specification and results of the five different case studies carried out on VSC-3.

DISTRIBUTED THEMATIC MAPPING PERFORMANCE OPTIMIZATION IN PUBLIC CLOUDS

J. Lozano, M. Quartulli, J. Egaña, I.G. Olaizola

Vicomtech-IK4
Digital Media Department
Paseo Mikeletegi 57,
20009 Donostia, Spain

E. Zulueta

University of Basque Country
Sys. Eng. and Automation Department
Nieves Cano 12,
1006 Gasteiz, Spain

ABSTRACT

Global distributed thematic mapping in public clouds requires optimized data flows. These optimized flows can be the result of the analysis by Machine Learning (ML) of a deeply sensorized mapping system. In this sense, distributed global mapping requires a monitoring system that allows to understand the internal working of the system and enables the implementation of corrective actions to increase system performance. This work presents an implementation of a system monitoring framework and the obtained analysis results.

***Index Terms*—** System monitoring, big data, web mapping

1. INTRODUCTION

Performance-optimized tasks are required to make the best use of large scale geospatial computing infrastructures [1, 2, 3]. Improvements in the performance can result in a drop in overall computational and financial costs. This optimization can be the result of the analysis of the motorization system: contributions such as [4] describe the typical complexities and uncertainties in this kind of infrastructures, like resource contention and its uncertainty or robustness, to address real-time problems through robust big data solutions.

Recent advances in Remote Sensing technology have improved the volume of the available image data. In this contribution, we consider the data available in the Open Data Euskadi repository [1]. The ortho-imagery, acquired annually by an air-borne camera, is characterized by a spatial resolution of 25 cm in each direction. As a consequence of technological limitations in the acquisition and storage systems, the number of available channels is limited to three. The available coverage spans to the whole extension of the Basque Country, resulting in around 1500 GeoTIFF products with a size of about 33Gb each, as per figure 3. The need for an effective divide and conquer approach based on data tiling and parallel processing is evident in this context.

[1] `http://opendata.euskadi.net/`

2. SYSTEM DESCRIPTION

The prototype presented in [5] is a web map server with integrated machine learning — including classification and clustering — capabilities, able to create thematic coverage maps. The prototype is if needed able to distribute processing loads by big data frameworks like Apache Spark [6], taking full advantage of distributed in-memory computing.

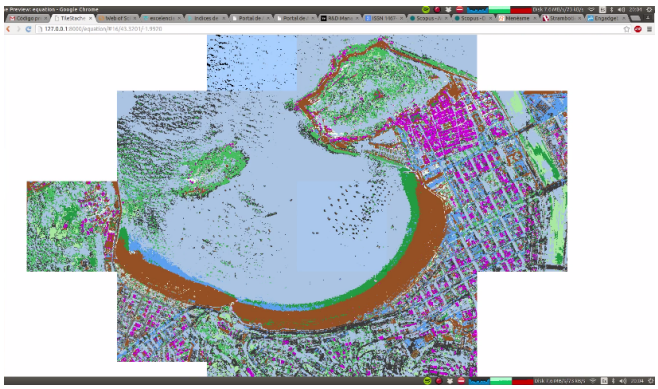

Fig. 1. Prototype UI for real time execution of supervised classification tasks across the map. The thematic coverage tiling process is similar to typical tiled map navigation data service, yet it includes components that allow causal users to define thematic classes of interest that might be unique to their interests and activities.

The operation of the prototype for supervised classification can be described as follows. A trained classification model is distributed to each worker operating in the infrastructure for the generation of the thematic map. While interactively navigating the map, each tile of the thematic coverage map is created by a worker node in a lazy mode, based on the HTTP requests put forward by an HTML-based Graphical User Interface implementing a standard web map navigator. As can be see in figure 1, the thematic coverage tiling process is similar to typical tiled map navigation interfaces, yet it includes components that allow causal users to define thematic

classes of interest that might be unique to their interests and activities. To obtain each tile, different features are computed according to the training, in order to process them by the distributed trained classification model.

The architecture of the server side (figure 2) starts conceptually with data archives in charge of storing the original data. The data can in turn be analyzed by a scalable set of processors co-located with the archive. External users can access either the original or the processed data, as well as exploit higher-level descriptions of data content also available as Services defined by the system, by accessing application servers dedicated to specific data exploitation scenarios such as those dedicated to specific applications in forestry or in the exploitation of marine resources.

Computing resource consumption measures like CPU and memory need to be monitored closely in order to analyze the behavior of the system in the face of varying user demand. Further measures corresponding to the inner operations of the training and classification system can be very useful to obtain intermediate information or to improve the performance of the prototype.

The monitoring of these metrics is the focus of the present work.

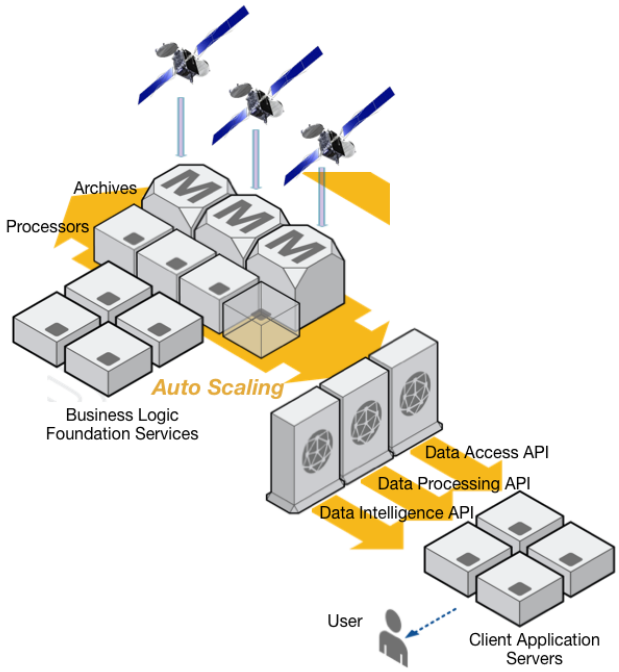

Fig. 2. System architecture. Archives in charge of storing the original data are served by a scalable set of processors co-located with the archive. External users can access either the original or the processed data, as well as exploit higher-level descriptions of data content also available as Services defined by the system.

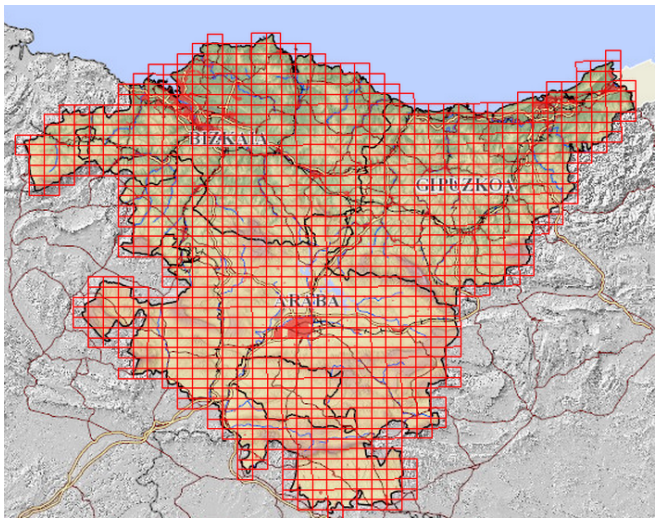

Fig. 3. Available GeoTIFF product map. The ortho-imagery, acquired annually by an air-borne camera, is characterized by a spatial resolution of 25 cm in each direction. As a consequence of technological limitations in the acquisition and storage systems, the number of available channels is limited to three. The available coverage spans to the whole extension of the Basque Country, resulting in around 1500 GeoTIFF products with a size of about 33Gb each.

2.1. Monitoring subsystem

The proposed monitoring system tries to optimize the distributed thematic mapping process by minimizing the impact on the system performance: because the user interface of this prototype is web based, processing time acquires a unusual relevance in machine learning processes where time limitations usually play a less prominent role [7].

The main task of the server is to create thematic coverage tiles based on common machine learning models, to build a coherent coverage map.

System quality as related to the classification results measured with metrics like Precision, Recall and F1 are not the object of the current contribution: they have been extensively described in [5] and [8].

The monitoring of the computing performance of the different modules in the system is here understood as an internal description of the operation of the system, rather than being related to external ground truth maps. From an architectural point of view, the analyzed performance metrics are recollected in each worker node and aggregated in a specialized node of the infrastructure destined to this task. Using a specialized web based user interface dedicated to the administration of the system, these values can be monitored in real time as per Figure 4.

Whenever required, these metrics could be represented and analyzed statistically using multi-variate analysis and representation methods, to analyze the behavior of the system

Fig. 4. Real time monitoring system UI. The UI allows to follow the operation of the prototype in real time, contributing to its understanding and allowing operators to optimize its processing mechanisms.

and to try and determine the most appropriate actions that allow the improvement of the performance of the system in any of the presented contexts.

3. MONITORING RESULTS

As stated, the generated metrics are collected to serve the internal optimization of the system.

Yet, when the system is applied to a specific dataset, their content is naturally related to the characteristics of this data.

A specific case is reported in relation to figure 5.

The model learning time is in this case reported for each visited tile across a scene covering the city of Donostia/San Sebastián. The considered machine learning model is an instance of the K-Means algorithm [9]. For investigating the observable variations if the classification costs in the prototype, the model is re-trained separately on each one of the coverage tiles — naturally resulting in incompatible assignments across the whole coverage. The measured processing times are represented in terms of a heatmap-like layer that we superimpose to a map of the processed area, as in figure 5. Variations in processing time seem to be explicable in terms of the variability of the input data: very flat areas result in slower analysis operations (a result that can easily be replicated in the case of many k-means implementations), or perhaps natural areas characterized by more inherent complexity generate data that has more variation to be accounted for in the analysis.

4. CONCLUSIONS

Optimizing the exploitation of big data infrastructures for distributed thematic mapping at global scales requires an optimal use of the resources to reduce computational and financial costs.

A monitoring system allows to control these infrastructures so that they can attain better performance by applying appropriate measures, allowing operators to develop own metrics in addition to the classics that enrich the knowledge of the system.

This understanding of system performance effects can leverage the exploitation of machine learning methodologies able for instance to identify the most important metrics or to synthesize summary measures with increased information content.

This better knowledge of the inner workings of the system can produce a better performance of the monitored system and therefore the capability of analyzing much more extended coverage maps with reduced computational costs.

5. REFERENCES

[1] Dan Wang and Jiangchuan Liu, "Optimizing big data processing performance in the public cloud: opportunities and approaches," *Network, IEEE*, vol. 29, no. 5, pp. 31–35, September 2015.

[2] Stefano Nativi, Paolo Mazzetti, Mattia Santoro, Fabrizio Papeschi, Max Craglia, and Osamu Ochiai, "Big data challenges in building the global earth observation system of systems," *Environmental Modelling & Software*, vol. 68, pp. 1–26, 2015.

[3] Yan Ma, Haiping Wu, Lizhe Wang, Bormin Huang, Rajiv Ranjan, Albert Zomaya, and Wei Jie, "Remote sensing big data computing: Challenges and opportunities," *Future Generation Computer Systems*, vol. 51, pp. 47–60, 2015.

[4] R. Ranjan, "Modeling and simulation in performance optimization of big data processing frameworks," *Cloud Computing, IEEE*, vol. 1, no. 4, pp. 14–19, Nov 2014.

[5] J. Lozano, N. Aginako, M. Quartulli, I.G. Olaizola, E. Zulueta, and P. Iriondo, "Large scale thematic mapping by supervised machine learning on 'big data' distributed cluster computing frameworks," in *Geoscience and Remote Sensing Symposium (IGARSS), 2015 IEEE International*, July 2015, pp. 1504–1507.

[6] Matei Zaharia, Mosharaf Chowdhury, Michael J. Franklin, Scott Shenker, and Ion Stoica, "Spark: Cluster computing with working sets," in *Proceedings of the*

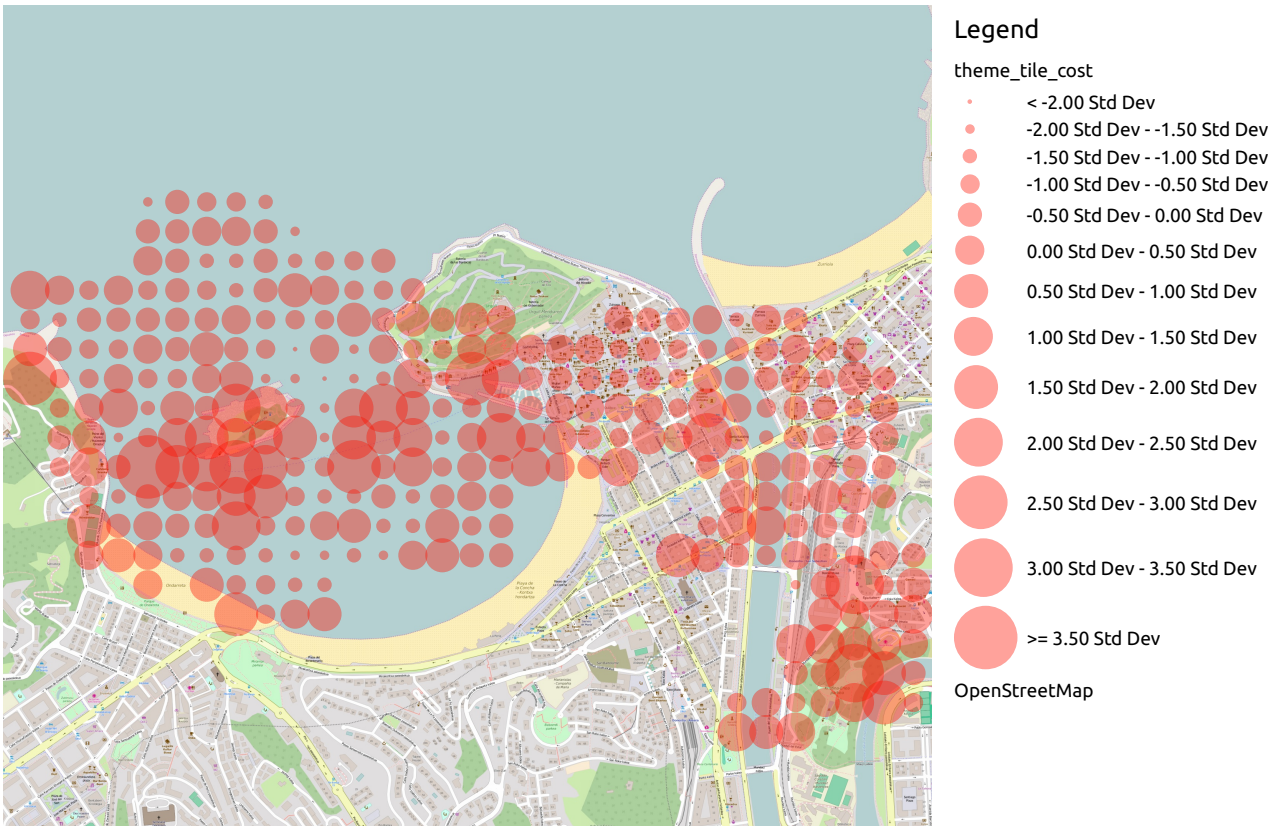

Fig. 5. K-Means clustering model learning time is represented for each visited tile across a scene covering the city of Donostia/San Sebastián. For investigating the observable variations if the classification costs in the prototype, the model is re-trained separately on each one of the coverage tiles — naturally resulting in incompatible assignments across the whole coverage. As is typical for k-means implementations, variations in processing time seem to be explicable in terms of the variability of the input data: very flat areas result in slower analysis operations (a result that can easily be replicated in the case of many k-means implementations), or perhaps natural areas characterized by more inherent complexity generate data that has more variation to be accounted for in the analysis.

2Nd USENIX Conference on Hot Topics in Cloud Computing, Berkeley, CA, USA, 2010, HotCloud'10, pp. 10–10, USENIX Association.

[7] Fiona Fui-Hoon Nah, "A study on tolerable waiting time: how long are web users willing to wait?," *BEHAVIOUR & INFORMATION TECHNOLOGY*, vol. 23, no. 3, pp. 153–163, 2004.

[8] J. Lozano Silva, N. Aginako Bengoa, M. Quartulli, I.G. Olaizola, and E. Zulueta, "Web-based supervised thematic mapping," *Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of*, vol. 8, no. 5, pp. 2165–2176, May 2015.

[9] Richard O Duda, Peter E Hart, and David G Stork, *Pattern classification*, John Wiley & Sons, 2012.

MUSCATE: multi-satellites, multi-sensors and multi-temporal THEIA Data and Services Infrastructure

Joëlle DONADIEU[1], Céline L'HELGUEN[1], Simon BAILLARIN[1],Marc LEROY[1], Philippe PACHOLCZYK[1], Olivier HAGOLLE[2], Mireille HUC[2], Jérôme GASPERI[1], Beatrice PETRUCCI[1], Olivier MELET[1], Manuel PAVY, Emmanuelle SARRAZIN[1], Robert NGO[4],Laurent LEBEGUE[1], Vincent LONJOU[1],Camille DESJARDINS[1],Anne-Claire GERARD[3], Adrien VIDOT[3], Dominique CLESSE[3]

[1] CNES, 18, avenue Edouard Belin 31401 TOULOUSE CEDEX 4, France – celine.l-helguen@cnes.fr
[2] CESBIO, 18 avenue Edouard Belin BPI 2801 31401 TOULOUSE CEDEX 9, France
[3] CAPGEMINI, 109 Avenue Eisenhower 31036 TOULOUSE Cedex, France
[4] ALTEN Sud Ouest, 9 rue Alain Fournier 31300 TOULOUSE , France

ABSTRACT

MUSCATE is the part of the THEIA data land centre dedicated to process optical images from SPOT, LANDSAT and SENTINEL-2 satellites. These images are intended for both scientific community and institutional actors.
Currently in qualification phase, MUSCATE will automatically generate and distribute either ortho-rectified images expressed in reflectance or temporal syntheses. These processes have high computing resource needs which require a performing computing centre. In order to reduce development cost and development risks, MUSCATE design is based on CNES computational shared facility and on relevant CNES image algorithms and supervision framework.
After an overall presentation of MUSCATE architecture, then implemented image processing, the paper focuses on the MUSCATE workflow dynamics and the advantage of using PHOEBUS[5] framework coupled with a dedicated sequencing module in order to regulate the processing flow sent to the HPC centre and to fulfil the MUSCATE high demanding timeliness requirements.
This particular architecture will enable MUSCATE to process a large volume of satellite images within the allotted time.

Index Terms— *THEIA, MUSCATE, High Performance Computing, SPOT World Heritage, LANDSAT, SENTINEL-2, multi-temporal, data valorization*

INTRODUCTION

THEIA is a French national multi-agency organization which promotes the use of satellite data by scientific community and institutional actors. This consortium aims at helping monitoring human and climate impacts on ecosystems and territories by delivering a large panel of products and mutualised services offering the user community to get the largest benefit of data and products from space missions.
As Data and Services Infrastructures of THEIA, MUSCATE is designed to acquire, process and distribute automatically high resolution satellite images from SPOT 1 to 5, LANDSAT 5-7 and 8, and SENTINEL-2 which cover France territories and worldwide areas of interest.
CNES has developed with CAP GEMINI the MUSCATE framework and has deployed it on CNES High Performance Computing (HPC) centre based in Toulouse (France).
MUSCATE is directly linked to the French Sentinel products exploitation platform (called PEPS [6]) to download Sentinel-2 data and generate value-added products. Developed by CNES in the context of collaborative ground segments of Copernicus program, PEPS platform broadcasts the Sentinel products at a national level increasing performances access to the high volume of Sentinel data. In the short term (up to 2017), PEPS covers the national needs for Sentinel data and associated services, while waiting for an European solution. PEPS relies on CNES (French National Space Agency) data storage facility HPSS (High Performance storage System) based in Toulouse (France).

Currently in qualification phase, the objectives of MUSCATE are to process up to 250 SPOT Level-1C, 20 LANDSAT Level-2A, and 1600 SENTINEL-2 Level-2A products a day.
This document presents firstly MUSCATE architecture, then implemented image processing, workflow dynamics and finally the first results of qualification.

1. MUSCATE ARCHITECTURE

One of the challenges MUSCATE has to take-up is to reduce production time at a minimum especially between the date of acquisition of SENTINEL-2 and LANDSAT-8 products and the availability of corresponding value-added

products on THEIA website. This requires the use of a performing computing architecture offered by the CNES HPC centre.

This computing cluster is a shared facility at the disposal of CNES projects and partners. It has a large computing capacity which should enable MUSCATE to meet its production objectives (almost 1900 products per day) without deteriorating the level of service for all the others users of this computing centre. MUSCATE will access to the CNES HPC centre which is composed of more than 70 Red Hat Enterprise Linux 6 computation nodes, each of these nodes being equipped with 16 cores, 64Go of RAM and 500Go of drive space. The use of this HPC centre will enable to meet timeliness requirements thanks to its high capacity of parallel processing.

MUSCATE relies also on various CNES computing services such as the different types of data storage, DBMS (Data Base Management System), archive facility, web-site hosting services which are also mutualised services offered by CNES computing division.

MUSCATE is composed of two main components: an Acquisition&Production Module and a Distribution Module which relies also on shared facilities.

1.1. Acquisition and Distribution Module

The Acquisition&Production Module uses HPC centre, DBMS, HPSS, GPFS (General Parallel File System), archive means. It is the major component of MUSCATE.

It relies on a CNES orchestration and processing framework called PHOEBUS and on a repository administrator component which carries MUSCATE system intelligence. PHOEBUS and the repository administrator component are deployed on an independent server.
PHOEBUS which is in charge of the processing orchestration sends jobs on the cluster centre and manages their priority. Operators can then monitor processes in progress and act when required through its interface.
Each process is split in steps which are sequentially sent to the distributed resource manager (DRM) of the CNES HPC centre. Each step can be parallelised computing in order to optimize HPC centre occupation.

As the DRM gives priority to short and mono-processor jobs, MUSCATE should find a compromise between length of jobs and the number of allocated processors per job. MUSCATE should indeed be able to process the equivalent of more than several thousands of computing hours per day which requires a perfect sequence of computing jobs.

The PHOEBUS framework integrates all the necessary functions to automatically generate products from inputs: management of complex tasks sequence, massive parallelised computing, optimisation of computing cluster use, reprocessing campaigns.

For MUSCATE project, PHOEBUS enables to define specific workflows for each type of acquired products. This framework facilitates the integration of image algorithm in an HPC architecture which should process massive volume of data. It launches processing by step at customizable hours and also gives the possibility to regulate the number of jobs send in parallel to the CNES cluster. The use of PHOEBUS in MUSCATE is a real advantage because it regulates the flow of jobs sent to the HPC centre and also enables to optimize this flow to meet MUSCATE high processing demand.

The repository administrator component manages the products, and via its specific interface, operators can request on products state. It is responsible of the treatments sequencing according to the availability of the products. It relies on production centre database which is lodged on shared Postgres database server.

In the Acquisition&Production Module, MUSCATE should use the most relevant datastorage for its different data type; temporary data should be on quickly accessible spaces but not necessarily redounded, while all inputs and outputs have to be archived with a high quality service. All these other infrastructures are also shared services offered by CNES computing division.

For temporary data necessary to production activities, the Acquisition&Production Module uses GPFS which are quickly accessible spaces. Inputs and outputs have to be archived with a high quality service (STAF).

MUSCATE acquires SENTINEL-2 Level-1C Products available on PEPS platform, Landsat-8 Products from USGS and SPOT from dedicated exchange spaces.

During the output preparation step, downloadable output products are deposed on HPSS (High Performance Storage System) and visualisable output data (multi-scaled labels) on NAS (Network Attached Storage) which are accessible spaces for the Distribution Centre..

1.2. Distribution Module

The role of the Distribution Module is to diffuse its products to the scientist and public community.

Users will have a direct access to products in the Distribution Module in CNES system through MUSCATE research HMI.

The MUSCATE Distribution Module is composed of two web sites using a VM (Virtual Machine) based on a shared infrastructure. These two web-sites are packed in CNES Apache&Tomcat software components CNES package called WEB-NG.

1.3 Interface with Expertise space

In parallel of the automatic production, MUSCATE may export defective products towards expertise space. Defective products are products whose production failed because their quality index is poor. The expertise tools help quality image experts to investigate the origin of the issue. It integrates a panel of algorithms to calculate statistics on the defective products, visualisation tools and the possibility to run processing steps outside of the production centre.

2. IMPLEMENTED IMAGE PROCESSING

MUSCATE generates and distributes 3 levels of products based on CEOS standard;
- level 1C: orthorectified product in TOA reflectance (SPOT, Landsat5,7),
- level 2A: orthorectified product in surface reflectance (Sentinel2 and Landsat8),
- level 3A: temporal synthesis of level 2A products (Sentinel2 and Landsat8).

This processing level of acquired and distributed products varies with satellites and sensors.

Acquired SPOT products are first controlled and eventually corrected before being orthorectified and converted in TOA reflectance. MUSCATE will thus distribute level 1C SPOT World Heritage products.

LANDSAT processing varies according to the level of the acquired product and the sensors. For example LANDSAT-7 products acquired after 31st May 2003 first need correction due to a defective mirror while L1G products must be orthorectified. All LANDSAT products are then converted in TOA reflectance, tiled and mosaicked before being converted in surface reflectance. MUSCATE is designed to distribute level 2A LANDSAT products.

SENTINEL-2 inputs will be level 1C products. They might be reprojected and tilled before being converted in surface reflectance. MUSCATE is planned for distributed level 2A.

To generate these products, MUSCATE is based on CNES image algorithms for the complex and sizing computation steps: orthorectification, surface reflectance computation and temporal synthesis algorithm. The intermediate steps such as original product correction, conversion in TOA reflectance, tiling, mosaicking and geographical reprojection, are specifically developed for MUSCATE.

Orthorectification is composed of a registration and a resampling step. Both of these algorithms are implemented in a CNES component, called SIGMA (integrated, generic and autonomous mosaic system).

The conversion in surface reflectance and cloud detection are implemented in another CNES component called MACCS (Multi-mission Atmospheric and Cloud Correction Software). This algorithm has been designed by CESBIO (centre for the study of the biosphere from space) and is based on a temporal analysis. Indeed the generation of a level 2A product from a level 1C product nominally required the previously generated level 2A product acquired on the exact same location no more than a few days before the level 1C product to process. In order to initiate a new time series of level 2A products , MACCS also offers a “backward” mode witch takes as input the level 1C product to process up to level 2A as well as N subsequent level 1C products located on the exact same location. The degraded mode offers by this algorithm is the “init” mode which only takes the level 1C product to process as input but the results quality is much poorer than with the “nominal” or “backward” modes.

MUSCATE Production Centre activates automatically MACCS in the adequate conversion mode depending on the inputs products inventoried by the repository administrator component; the “nominal” mode if there are level 1C and level 2A products as inputs, the backward mode if there are N level 1C products as inputs, the backward mode if there are N level 1C products as inputs, and the “init” mode if there is only the level 1C to process as input.

3. MUSCATE WORKFLOW DYNAMICS

MUSCATE has been developed in order to automate the processing of large amount of satellite images.

In the case of SENTINEL-2, Level-1C products are processed by MUSCATE through 6 different steps; acquisition, conversion and ingestion by the repository administrator component, possibly tiling and mosaicking, conversion in surface reflectance and cloud detection (MACCS), preparation for distribution, archiving.

For example, when SENTINEL-2 Level-1C products have been acquired from PEPS platform and arrive in the MUSCATE repository administrator component database, an “available” state is assigned to these products and the repository administrator component predicts the outputs of each of these next processing steps which will be applied to the images. New objects are created in MUSCATE database at a “predicted” state: tiled and mosaicked, converted in TOA reflectance and prepared objects.
The repository administrator component creates tasks for the different steps in the production database with available inputs.
Then PHOEBUS collects the tasks in the production database and sends jobs to cluster to execute the processing steps.
At the end of a processing step, the outputs objects state is updated to “available” while the next step product state becomes “to be launched” and so on until the preparation for distribution.
All this processing is automatically chained without operator intervention.

4. QUALIFICATION-FIRST RESULTS

MUSCATE offers three HMI: Phoebus HMI to follow processes status, a repository administrator component HMI to follow in-progress products and a Web distribution HMI to research, visualize and downloading output products.

The MUSCATE distribution HMI offers a unique research service of all products organised by projects (SENTINEL-2, LANDSAT, SPOT WORLD HERITAGE). Each project is independent and has its own research, visualization and downloading web services. The access to web services of the distribution centre is dependent on each user rights.

This chapter gives an overview of the very promising results on MUSCATE operational processing and distributing activities of SPOT1-5, LANDSAT and SENTINEL2 products.
Several scenarii with nominal treatments without retreatments have been tested with a dedicated sequencing module in order to regulate production flow.
We have to be very vigilant and control resources (jobs, threads ...) as we are on shared computing facilities.
By first feedback experience, PHOEBUS should be configured to launch very regularly few workplans in order to smoothing memory and CPU consumption.
For example, a scenario with 250 Spot World Heritage, 18 Landsat 8, 210 SENTINEL-2 products a day lasts 14 hours with a minimal occupation of the shared cluster.

Others scenarii with massive retreatments (up to 1250 products) are foreseen to be tested to fulfil the MUSCATE high demanding timeliness requirements and to fit with the operational scenarii for the exploitation phase.

5. CONCLUSION AND OUTLOOKS

Thanks to this program, the data heritage of SPOT 1 to 5 satellites family will be valorised while LANDSAT 8 and SENTINEL-2 products will be processed and distributed to the scientific community and institutional actors in only a few days after their acquisition.
As MUSCATE is designed to be at least 10-years operational, a database of more than 30 years of satellite optical images should be available for the scientists to analyse and better understand how our planet changes.
The very promising results suggest scientific community to integrate in the near future new processing developed by scientists such as snow detection, soil occupation map and temporal synthesis.

5. REFERENCES

[1] S. Baillarin, P. Gigord and O. Hagolle, « Automatic Registration of Optical Images, a Stake for Future Missions: Application to Ortho-Rectification, Time Series and Mosaic Products », Geoscience and Remote Sensing Symposium, 2008,2:II-1112-II-1115.doi:10.1109/IGARSS.2008.4779194 , 2008
[2] O. Hagolle, G. Dedieu, B. Mougenot, V. Debaecker, B. Duchemin, and A Meygret, « Correction of aerosol effects on multi-temporal images acquired with constant viewing angles: Application to Formosat-2 images », Remote Sensing of Environment 112 (4) (avril 15): 1689-1701. doi:10.1016/j.rse.2007.08.016, 2008
[3] O. Hagolle, M. Huc, D. Villa Pascual and G. Dedieu, « A multi-temporal method for cloud detection, applied to FORMOSAT-2, VENµS, LANDSAT and SENTINEL-2 images », Remote Sensing of Environment 114 (8) (août 16): 1747-1755. doi:10.1016/j.rse.2010.03.002, 2010
[4] C. L’Helguen et al. “MUSCATE: Multi-satellites, multi-sensors and multi-temporal THEIA data centre”, Big Data From Space 2014, Frascati, ESRIN.
[5] J.N. Hourcastagnou. “Phoebus a generic fully integrated workflow execution framework for high performance data processing centres: overview, description and lessons learned in customizations ”, SpaceOps,2012.
[6] Plateforme d’Exploitation des Produits Sentinel https://peps.cnes.fr

DEIMOS' GS4EO OVER ENTICE: A COST-EFFECTIVE CLOUD-BASED SOLUTION TO DEPLOY AND OPERATE FLEXIBLE BIG EO DATA SYSTEMS WITH OPTIMIZED PERFORMANCE

José Julio Ramos[†] *and Jonathan Becedas*[‡]

Elecnor Deimos

[†] Deimos UK Ltd.
Building R103, Fermi Avenue – Harwell Oxford
Didcot, OX11 0QR, UK
Tel.: +44 (0) 1235 56 7173
e-mail: jose-julio.ramos@deimos-space.co.uk

[‡] Satellite Systems
Francia 9, Poligono Industrial la Nava
13500 Puertollano, Ciudad Real
Tel.: + 34 926 44 35 78
e-mail: jonathan.becedas@elecnor-deimos.es

ABSTRACT

Processing and distribution of big space data still presents a critical challenge: the treatment of massive and large-sized data obtained from Earth Observation satellite recordings. Remote sensing industries implement on-site conventional infrastructures to acquire, store, process and distribute the geo-information generated. However these solutions do not cover sudden changes in the demand of services and the access to the information presents large latencies. Cloud infrastructures are envisaged as a possible solution to improve EO systems dealing with large datasets, however the current technology still presents some limitations: i) the virtual machine images (VMIs) are not optimized being highly over-sized which impacts in the costs of using the infrastructure and in the dynamic resources provisioning; and ii) the deployment of Virtual Machines (VM) in cloud has large duration, normally between 10 and 20 minutes which directly affects to the flexibility and dynamic scalability of the system. Deimos' research focuses in the development of Future Internet technologies in order to improve Earth Observation (EO) services and to highly reduce the costs associated with on premises deployment. Within the ENTICE H2020 project, Deimos intends to implement a flexible, cost-effective Payload Data Ground Segment (PDGS) in a cloud computing infrastructure by reducing 60% the VMI size, 30% VMI delivery time, 25% deployment time, 25% VM cost and 25% VM storage.

Index Terms— *Earth Observation, Remote Sensing, Future Internet, Cloud Computing, Big Data, ENTICE.*

1. INTRODUCTION

Although the proven importance of EO in society, the access to the information obtained from satellites follows traditional and expensive paths to cover on demand services for different potential customers: conventional data centres and conventional distribution of services. This presents several drawbacks: i) the cost of acquiring recent images of the Earth is very high, which limits the access of the general public, ii) clients cannot access directly neither fast to the information they need because this has to be processed and ad-hoc distributed, iii) the service is not flexible to be adapted to sudden changes in the demand.

Cloud computing is presented as a possible solution to improve common services and create new market opportunities because it is elastic, scalable and it works on demand through virtualization of resources [1].

Satellite and Earth Observation applications are then clear use cases for deployment on the cloud for the following reasons:

- The global nature of EO data, with ground stations and users geographically located all over the world, mandates to deploy a worldwide infrastructure connecting stakeholders.
- The massive size of EO data generated by today´s sensors, in the order of daily Terabytes, suggests the use of a cost-effective procurement of the computer infrastructure for archive and processing.
- The amount of resources needed to optimally process and distribute EO data to the global user community is large.

Earth Observation systems are ideal use cases to be deployed on a cloud infrastructure with a dedicated middleware for automatic and intelligent provisioning and deployment to run EO applications.

2. BACKGROUND

In previous work of the research group, an Earth Observation system was successfully implemented in Future Internet cloud infrastructure: the GEO-Cloud project [2].

The GEO-Cloud project was an experiment in which a complete EO system was implemented in Fed4FIRE (http://www.fed4fire.eu/) [3]. A data centre was

specifically designed to be implemented as a distributed system, and a simulator of a satellite constellation was developed and implemented. Numerous users were emulated remotely accessing to the infrastructure. Promising results were obtained from that work [4]:

- Parallelization of processes to reduce the image processing time.
- Reduction of delivery time to end users.
- Possibility of having a distributed data centre on the World to facilitate the ingestion of images from a network of antennas.
- Externalization of infrastructure costs.
- On demand cost of computing resources

Then the system was implemented in a commercial infrastructure, in SoftLayer®, with the support of our IBM and Itera Process partners [5] with a successful increase in the computing performance among other improvements:

- Reduction of data transfers by using shared storage.
- Auto scaling
- Image processing time reduction
- 24/365 support
- Replication and download to local storage.
- Reconfiguration of instances on the fly.
- Security.

However, there are two major issues that can highly improve the applicability of cloud computing to EO applications:

- Storage in public cloud infrastructures is expensive, so effort in the optimization of virtual machines has to be done to reduce the high costs of implementing a whole mission, which have several years duration, in cloud.
- A major aspect is that the deployment of VMs is of several minutes, usually about 10 minutes [5]. For real time elasticity this is a major problem since the system cannot be reconfigured in real time. For example to scale down or scale up the distribution of products in function of the users' demand. This increases the costs of using the infrastructure and the latency to deliver the final products.

Within the ENTICE H2020 project (http://www.entice-project.eu/), it is expected the creation of a middleware that will optimize the size of VMIs, will reduce the deployment of the virtual machines and will increase the performance of the whole system. In this work we introduce the use of ENTICE to increase the performance of the Earth Observation Payload Data Ground Segment which is being used to manage the data recorded by the high resolution satellite Deimos-2. It is intended that the commercial system will be included in the gs4EO suite as the Cloud Processing Station (CPS) component.

3. USE CASES

The CPS system will typically have three main trialling use cases iterating with different external actors:

- Use Case A. "Land Management" - Governments, planning offices and agricultural agencies need to characterize landscape and crops in vast rural regions. These users demand a multi-temporal product which allows analysing the land cover dynamics and detecting changes during a vast amount of time. The main objective is to optimize the archive of satellite recordings in order to facilitate reprocessing campaigns and time-series analysis.
- Use Case B. "Infrastructure monitoring" - During the construction or operations of major infrastructures in remote, hazardous areas such as high performance railway lines in the desert. These infrastructures would need constant monitoring to minimize the impact of the ambient conditions. The main objective is to optimize the processing of the satellite imagery to facilitate on-demand processing.
- Use Case C. "Disaster response" - Upon the event of a disaster, say an earthquake on a remote area, government and safety agencies would need to acquire, as soon as possible, images from the affected area before, during and after the event. Those images would be later demanded by a huge number of users from all around the world, either to manage emergency services and to assess the damages. The objective is to optimize the distribution of satellite imagery at a global scale by creating a Content Distribution Network.

4. OPTIMIZATION OBJECTIVES

Thus, to fulfil the objectives described in the Use Cases of section 3, the Cloud Processing Station (CPS) pilot case for ENTICE has to optimize a Space Data System deployed on cloud. Two main metrics are defined to validate the system: its overall data latency and system cost without any re-engineering activity.

- **Data latency** - Space Data Systems deployed on premises generates products after three generic steps which defines the overall system's data latency: processing, archiving and distribution. Each of these steps includes a data movement phase over a network (with latency depending on data quantity and server proximity) to present the data to the core functionality, respectively processing, archiving and distribution. A scalable deployment of the system on the cloud would include added activities to each general step dedicated to Virtual Machines management (provisioning, importing/exporting, spinning up and down, migration, destruction, etc.) which, while improving the system's

scalability, will impact negatively on the overall data latency. One of the ENTICE optimization objectives is to reduce the overall data latency by increasing its flexibility, minimizing the penalties on latency introduced by the VM management activities. Additionally, it also improves the core functionality performance.

- **System cost** - The main economic appeal of cloud computing consists of 'converting capital expenses to operating expenses" and paying for the use of utilities (resources in IaaS cloud model) reducing the risks of over- and under-provisioning a system [1]. Storage and computational resources like CPU architecture and memory directly affect the overall cost of a cloud deployment. ENTICE aims to minimise both needs by reducing the sizes of VMs, VM images and VM templates and the overall use of computational resources stripping down key functionality (processing, storage and distribution) of unnecessary dependencies, plus the time dedicated to VM management.

By applying ENTICE technology during the trialling of the distributed data centre proposed, we expect to reduce 60% the VMI size, 30% VMI delivery time, 25% deployment time, 25% VM storage and for instance at least a 25% VM of cost reduction in the use of cloud infrastructures.

5. ENTICE TECHNOLOGY

The development of the ENTICE middleware has the following main objectives:

- To simplify the creation of lightweight and highly optimised VM images tuned for functional descriptions of applications.
- To automatically decompose and distribute VM images based on multi-objective optimisation to meet application runtime requirements, performance, economic costs, storage size and QoS needs.
- To elastically auto-scale applications on Cloud resources based on their fluctuating load with optimised VM interoperability across Cloud infrastructures and without provider lock-in.

Those objectives are in line with the detected problems of current cloud technology detected in the GEO-Cloud project, and which are proposed to be solved with this work. The ENTICE middleware facilitates the implementation of the Cloud Processing Station system in cloud is done with the through functional descriptors. These functional descriptors define the system Service Level Agreement define at high and functional level the VMIs of the modules of the architecture and feed a Knowledge Base which highly contributes to optimize the VMIs. ENTICE analyses, synthesizes and optimizes the VMIs by automatically removing unnecessary VM content like unused libraries or old log files or merging different VM templates stored in the system. ENTICE also provides a federated cloud as a testbed for trialling the use cases. Figure 1 depicts the high level architecture of the ENTICE environment.

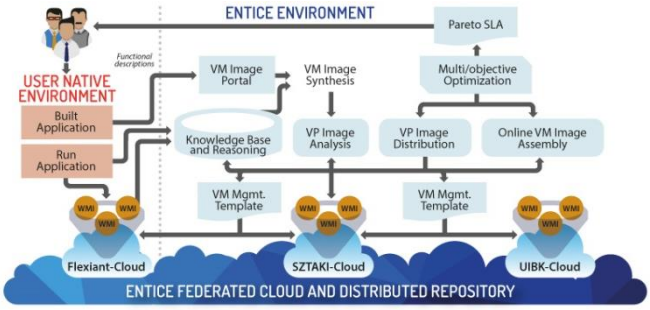

Figure 1 Architecture of ENTICE technology. ***Source:*** ***http://www.entice-project.eu/***

6. GS4EO CLOUD BASED IMPLEMENTATION WITH ENTICE

The GS4EO distributed architecture computed in cloud makes use of the ENTICE environment for its virtualization. The system is constituted by a GS Monitor which continuously pools over the different ground stations subscribed to the system, a Product Processors module which processes the row data obtained from the satellite, an Archive and Catalogue module which stores and classifies the processed images and a module providing User Services. Figure 2 shows a scheme of the implementation concept with ENTICE.

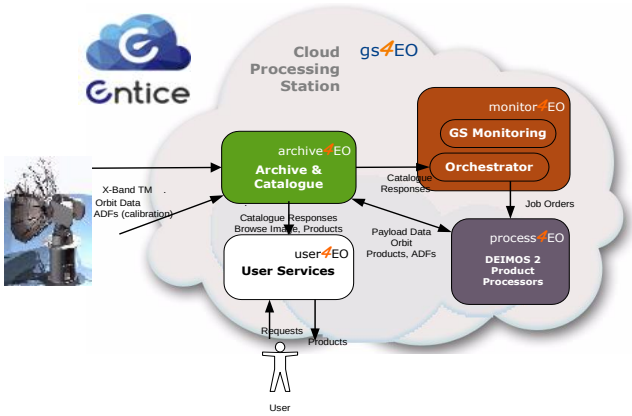

Figure 2 Implementation concept of the gs4EO Cloud Processing Station with ENTICE

doi: 10.2788/854791

7. FURTHER DISCUSSIONS

One added effect of applying the ENTICE technology on a system is obtaining a process for automatic generation of micro-services [6]. The nature of the ENTICE optimization phase allows a "top-down" approach on system engineering, i.e., by gradually removing all non-essential features not included in a functional description. If the functional description includes only one task, this optimization would generate a single task-focused micro-service without having to re-engineer the complete software from "bottom-up" into smaller modules.

One further consideration on this project regards system scalability. The Use Cases for testing the ENTICE technology, i.e. Land Management, Infrastructure Monitoring and Disaster Response are designed for scaling the three main parts of the system (respectively Storage, Processing and Distribution) and having to response to variable data demand scenarios. Scaling the data supply side is currently limited to adjusting for sporadic changes on the data reception (failures in ground station antennas or satellites) or to well-known passing-by calendars as defined by satellite flight dynamics. However, there exists the possibility to include further unpredictable (or linked to user requirements) variations on the data acquisition side of the system by using virtual, swarm-like constellations of satellites like the Proba-3 mission concept [7], federation of satellite constellations or even virtual satellites providing access to internal resources to different customers in the same fashion as the "ground" IaaS model, a sort of cloud computing in space [8].

8. CONCLUSIONS

The implementation of Earth Observation infrastructures proposed in this work contributes to optimize traditional EO data centres on premises with centralized storage and slow distribution of final products to convert the system in a flexible data centre based on scalability and flexibility of resources that optimizes the costs, improves the distribution of final products, provides on-demand services and facilitates the automation of the system while complying with a contracted level of service. The implementation approach with ENTICE facilitates the high level implementation and optimization of the resources by reducing the size of the VMIs, enhancing the overall system functionality by eliminating unnecessary processes and the deployment time of instances, which directly affects to the increase in the performance of the product processors and to the reduction of VM storage and costs in the use of cloud IaaS.

9. ACKNOWLEDGEMENTS

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 644179.

10. REFERENCES

[1] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia. "A View of Cloud Computing. Clearing the clouds away from the true potential and obstacles posed by this computing capability". *Communications of the ACM*, vol. 53, n. 4, pp. 50-58. April 2010.

[2] J. Becedas. "The GEO-Cloud Experiment: Global Earth Observation Computed in Cloud", *Eighth International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing (IMIS)*, Birmingham, pp. 9-15, 2-4 July 2014.

[3] T. Wauters, B. Vermeulen, W. Vandenberghe, P. Demeester, S. Taylor, et al. "Federation of Internet experimentation facilities: architecture and implementation. European". *Conference on Networks and Communications (EuCNC 2014)*, Bologne, June 2014.

[4] J. Becedas, R. Pérez and G. González, "Testing and validation of cloud infrastructures for Earth observation services with satellite constellations", *International Journal of Remote Sensing*, vol. 36, n. 19-20, DOI: 10.1080/01431161.2015.1070324, 2015.

[5] J. Becedas et al. "Evaluation of Future Internet Technologies for Processing and distribution of Satellite Imagery", *The 36th International Symposium on Remote Sensing of Environment, ISRE36*, Berlin, 11-15 May 2015.

[6] S. Newman, "Building Microservices. Designing Fine-Grained Systems", *O'Reilly Media.* ISBN 978-1-4919-5035-7, 2015.

[7] L. Tarabini Castellani, S. Llorente, J.M. Fernandez, M. Ruiz, A. Mestreau-Garreau, A. Cropp, A. Santovincenzo. "Proba-3 mission". *5th International Conference on Spacecraft Formation Flying Missions and Technologies*, Munich, May 2013.

[8] K. Kanev and N. Mirenkov, "Satellite Cloud Computing," *2011 IEEE Workshops of International Conference on Advanced Information Networking and Applications (WAINA)*, Biopolis, pp. 147-152, 22-25 March, 2011.

TASK ALLOCATION IN HIGH PERFORMANCE PROCESSING OF GEOSPATIAL DATA

V. Syrris, D. Rodriguez, and P. Soille

European Commission, Joint Research Centre (JRC)
Institute for the Protection and Security of the Citizens, Global Security and Crisis Management Unit

ABSTRACT

In a sandbox framework for scientific computing, we deal with the task allocation problem when processing a high volume of geospatial data. A predefined meta-information catalogue guides the data selection and a profiling procedure based on data sub-sampling estimates the memory and processing requirements. The task allocation is formulated as an optimization problem conditioned by time dependencies and job execution priorities. The procedure can work as add-on to any task scheduler that provides configuration options for computational resources allocation. Experiments demonstrate that the SLURM fine-tuning with the proposed method leads to better resource management and shorter schedules.

***Index Terms*—** Scheduling, Resource allocation, Sandbox, High Performance Computing, SLURM, Geospatial data

1. INTRODUCTION

The context of this work can be determined by a sandbox environment where a number of users are developing code prototypes and testing their processing chains over Earth Observation imageries and in general on geospatial data. The sandbox is supported by a cluster manager equipped with a job allocation and scheduling controller. Our starting point for this study is that the average user is not knowledgeable on how to optimally configure the jobs distribution, especially when the input data are different in size and the computational resources are heterogeneous in terms of processing units (CPU/GPU, number of cores, processing power, etc.) and memory cells. Intuitively, the user tends to overestimate the processing requirements a fact that leads to high latency and ultimately to suboptimal exploitation of the cluster platform. We note here that a task is often a part of a job or a job itself; in this work these two terms are used interchangeably.

As cluster manager, we consider SLURM [3, 8] that is an open source Resource and Job Management System designed for clusters of all sizes. SLURM uses a general purpose plug-in mechanism to select various features such as scheduling policies, process tracking or node allocation mechanisms.

The paper is organized as follows: Section 2 describes the automation of the task allocation set up by means of: i) a catalogue that comprises meta-information of the input data, ii) an estimation of the computational requirements based on task execution profiling and iii) the confrontation of task allocation as an optimization problem. Section 3 reports on the experimental results. Conclusions are given in Section 4.

2. METHOD

First, indicators from a geospatial meta-information catalogue are collected about the data volume to be processed (especially the minimal bounding box containing all the data values), overlapped geographical areas, and coordinate reference system and projection. This information sets the constraints that bound the amount of data to be processed and avoid repetitions [7].

Then, through profiling [9] on a carefully chosen subset of images, the computational requirements like processing speed and memory cells per task, network bandwidth [4], efficient number of concurrent jobs and others are estimated. The selection of the images is automatic: based on criteria defined with the aid of the meta-information catalogue as described previously, a systematic sampling takes place with the purpose of representing adequately the entire data universe. By surface fitting we extrapolate the computational requirements to the whole dataset.

The problem of task allocation on the different nodes of an heterogeneous computer cluster takes the form of the following constraint problem: Let I be the number of tasks, N the number of available nodes, R_n the number of cores in node n, M_n the number of memory cells in node n, m_i the lower memory threshold that is necessary for the task i to be executed, s_i the variable designating the order of sequential tasks and p_i the task priority of independent tasks. The execution time T_i of the task i is a function of two variables: r_{in} number of cores in node n and b_{in} number of memory cells in node n that are both necessary for the execution of i task: $T_i\left(r_{in}, b_{in}\right)$. That is, the optimal task allocation can be expressed as follows:

$$minimize \sum_{i \in I} \sum_{n \in N} x_{in}(1-p_i) s_i T_i\left(r_{in}, b_{in}\right) \tag{1}$$

subject to:

$$x_{in} \in \{0,1\}, \forall i \in I, n \in N \tag{2}$$

$$\sum_{n \in N} x_{in} = 1, \forall i \in I \tag{3}$$

$$\sum_{i \in I} \sum_{n \in N} x_{in} \leq \sum_{n \in N} R_n \tag{4}$$

$$0 \leq p_i < 1 \tag{5}$$

$$s_i - s_j > 0, i \neq j, i, j \in I \tag{6}$$

$$1 \leq r_{in} \leq R_n, n \in N, i \in I \tag{7}$$

$$m_i \leq b_{in} \leq M_n, n \in N, i \in I \tag{8}$$

The decision variable x_{in} represents the assignment of task i to node n. Condition (2) bounds the x_{in} values to 0 and 1. Condition (3) states that a task i can be assigned to only one node n, taking value 1 if the assignment is done and 0 otherwise. Condition (4) signifies that the number of allocated tasks cannot surpass the total number of cores over all the nodes. Since there are limited resources (expressed by constraints (4), (7) and (8)), in order for all the tasks to be processed, the optimization problem needs to be solved several times considering always the currently available resources.

Listing **Alg.**1 contains the algorithm's description. Steps 1–4 are executed once while steps 6–9 are running in several cycles and manipulating each task incrementally.

Algorithm 1: Task allocation algorithm

Input : Georeferenced images
Output: Sequence of tasks

1 Collect information from the catalogue;
2 $S \leftarrow$ sample, execute and profile subset of tasks;
3 Apply surface fitting and extrapolate parameter values to the entire input set;
4 p_i,$s_i \leftarrow$ Define values $\forall i$

5 **while** $i \leq I$ **do**
6 Collect available resources from each node;
7 Set bounds m_i, M_n and R_n, take the optimal decision and select one task i;
8 Allocate resources r_{in} and b_{in} and send task i to node n;
9 $I \leftarrow I \setminus \{i\}$;
10 **end**

The problem belongs to the general family of optimization problems of finding the shortest path $\{y_t\}_{t=0}^{<\infty}$ in finite time instants t. The goal is to minimize the *makespan* (1), that is the total length of the schedule (total execution time). The application of standard methods like mixed-integer linear programming is not straightforward; the component T_i of the objective function does not have a closed-form expression in terms of r_{in} and b_{in}, while the heterogeneous configuration of the cluster nodes increases the amount of constraints and the number of parameters that need fine-tuning. In the next section we show two greedy strategies that confront the specific optimization as a dynamic (online) scheduling problem.

3. EXPERIMENTS

The test case refers to the execution of an experimental cloud detector (an adjusted version of ACCA algorithm [6]) over a set of 1,000 Sentinel-2A (S2A) images. For the experiments, we used a partition of our in-house computer cluster infrastructure. The heterogeneity of the computer cluster is reflected by Table 1 which comprises the technical characteristics of the five cluster nodes.

Table 1: Specifications of the used computer cluster nodes

Name	Proc. units	Clock freq	Memory	Swap
isfhpc1	16(4;4;1)	2.1GHz	64GB	64GB
isfhpc2-3	8(2;4;1)	2.1GHz	32GB	32GB
isfgpu1-2	24(2;6;2)	2.8GHz	32GB	32GB

Notation (.;.;.) means: (socket;cores per socket;threads per core)

Figures 1a and 1b illustrate the variety and the velocity of the data under processing. They show the resource requirements, in terms of memory and execution speed, given by the profiling process.

First, nine groups were formed based on the size of the bounding box of the valid data pixels; value 0 has been assumed as *nodata* value. Figure 2 displays the 9-bin data pixels histogram of the available S2A B01 bands (until 31/1/2016). The explanation for the high size of group **c9** is that for these specific images, the entire image extent has been considered even if *nodata* may exist at some parts of the image. For instance, an image with zero *nodata* values and a lower triangular image belong both to the **c9** group. Then, the rightmost image was selected from each group as the worst-case scenario, i.e. the largest image in terms of number of pixels. Each of these images was sent to the cluster manager with 5 different resource allocation settings and the respective performance was monitoring and recorded. Numbers for the isfhpc2 and isfgpu1 nodes are not provided since they have identical configuration with isfhpc3 and isfgpu2 respectively. The cloud detector that takes as input four 10m resolution S2A bands and two of 20m resolution, exploits the fact of knowing the data domain (as this information is provided by the catalogue) and reads only the bounding box of the data area of the images. Figure 1c shows the memory profile of the cloud detector on isfgpu2, when reading the data domain of the bands of a tile that has been classified to the largest size group **c9**.

Task profiling helps to define statistically the function $T_i(r_{in}, b_{in})$ and provides output based on which the lower and upper thresholds of the optimization constraints are set up. Two greedy strategies were tested in order to incrementally schedule the tasks: i) Sort the list of tasks based on the priority and the estimated memory allocation in descending order and then, schedule the next task of the list at the first available node, ii) Apply the same ranking on the task list and then, if one of the faster nodes (isfgpu1-2) is available then select a task associated with an image from the medium to the

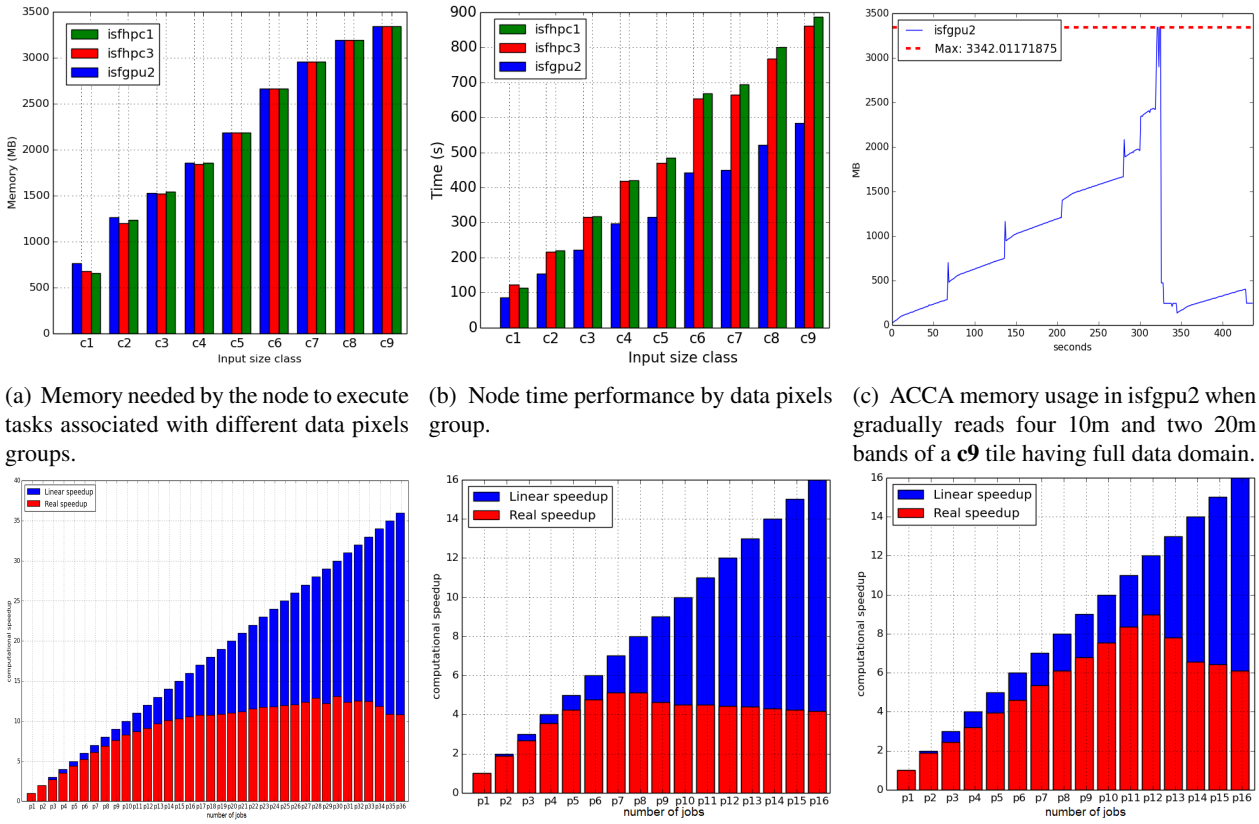

(a) Memory needed by the node to execute tasks associated with different data pixels groups.

(b) Node time performance by data pixels group.

(c) ACCA memory usage in isfgpu2 when gradually reads four 10m and two 20m bands of a **c9** tile having full data domain.

(d) isfhpc1 speedup tested on group **c9**.

(e) isfhpc3 speedup tested on group **c9**.

(f) isfgpu2 speedup tested on group **c9**.

Fig. 1: Space-Time complexity provided through profiling.

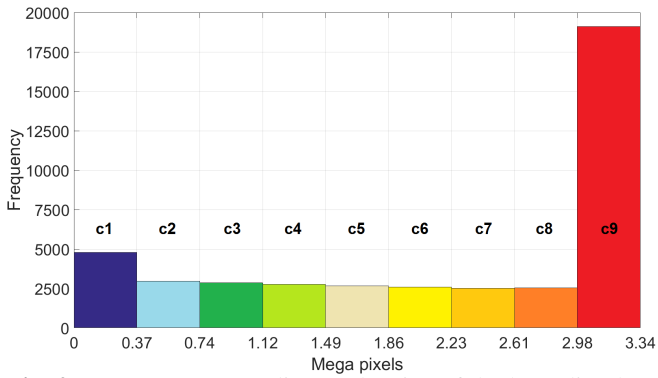

Fig. 2: Frequency according to the size of the bounding box of the data domain of the available (42,776) S2A 60m resolution bands (B01).

low size (**c1**–**c5**), whereas if one of the remaining three nodes (isfhpc1-2-3) becomes available, schedule a task associated with an image from the medium to the high size (**c6**–**c9**).

These two naive approaches together with the precise definition of the resources allocation demonstrate in practice a significant performance improvement in terms of schedule length. Figure 3 displays five makespans based on different approaches. From bottom to top: random sequence of tasks, read the full image domain and allocate 4GB for every task; the same conditions plus the bounding on the number of concurrent jobs; random sequence of tasks, read only the data domain, allocate the appropriate amount of memory according to the groups **c1**–**c9** at which the tile bands under

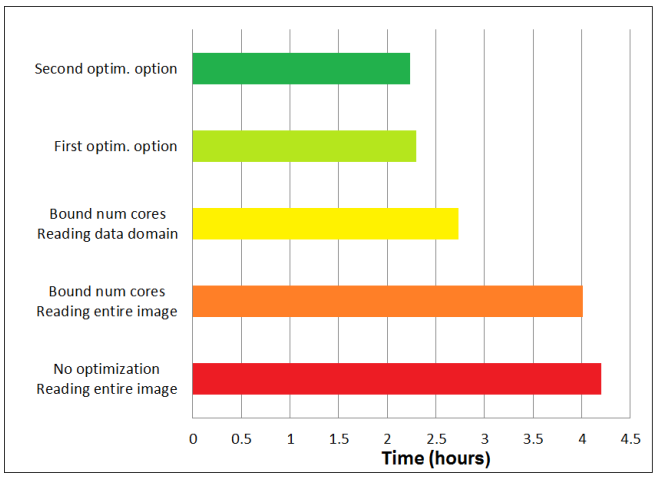

Fig. 3: Total execution time of the cloud detector (adjusted ACCA) based on the 6 bands of 1,000 S2A tiles, derived from five different schedules and managed by SLURM.

doi: 10.2788/854791

processing belong and bound the number of concurrent jobs; optimized task schedule based on the first strategy; optimized task schedule based on the second strategy. The bounding on the concurrent jobs is equivalent to the constraint imposed over the number of cores per node according to the profiling outcome (Figures 1d–1f). The explanation of the fact that a sorted list based on the size of input data domain leads to a shorter schedule length is that a homogeneous resource allocation can be achieved, permitting the fast nodes to process the larger amount of input data with the lowest waiting time. This phenomenon is amplified with the second strategy where the slower nodes are imposed to process the heavier tasks, leaving the space to the fast nodes to process evenly all the different types of tasks; in that way, they remain always occupied but in an effective mode. Table 2 shows the stability on the results when different proportions on the size of the nine groups are considered. The first column contains four different data domain distributions based on the nine bins: almost even, **c3–c6** reinforced, skewed in favour of **c1–c5**, skewed in favour of **c7–c9**. The scenarios **sc3–sc5** refer to the last three approaches respectively presented in Figure 3. The **sc1–sc2** scenarios refer to two ideal cases where the available memory is unlimited and not considered as constraint; the only difference is that in **sc1** a random task list is used as opposed to the sorted list in **sc2** case. The results derived from simulations executed 10 times for each scenario. One conclusion by observing these results is that both optimization strategies give stable schedules independently of the task list length or the input data distribution. In addition, the second strategy starts to lead to better performance in cases where the input data are of medium to low size; if the task list is dominated by input data of high size then both approaches give very similar performance.

The results validate the need for an informed preparation of the task list, and subsequently, for an adaptive decision making during the tasks execution. Apart from this finding, this work shows how the whole process can be automatized and relief the non specialized user from the burden of the task list formation.

Table 2: Simulation results when considering different data domain distributions and five scenarios (Time in hours).

Ratio by group	sc1	sc2	sc3	sc4	sc5
.1,.09,.08,.08,.08, .09,.12,.2,.16	1.68±.05	1.54±.00	2.97±.06	2.40±.01	2.35±.01
.02,.02,.2,.2,.2, .22,.03,.05,.04	1.49±.02	1.38±.00	2.62±.02	1.94±.01	1.83±.00
.2,.18,.16,.16,.16, .02,.02,.04,.03	1.11±.03	0.98±.00	1.92±.03	1.28±.01	1.16±.01
.02,.08,.01,.01, .01,.02,.23,.37,.3	2.15±.01	2.03±.00	3.85±.02	3.43±.00	3.43±.02

4. CONCLUSION AND OUTLOOK

This work demonstrates an automated method for fine-tuning a workload manager (SLURM in this specific application) in order to optimally distribute tasks in an heterogeneous computer cluster and do large-scale experiments with geospatial data. Job scheduling and resource allocation in high performance distributed environments [2, 5, 1] is now a wide area application, including also cloud computing resource scheduling [10]. Future work involves: i) the extension of the optimization problem in order to take into consideration factors like network bandwidth, users collision and composite jobs, and ii) the transferring and adaptation of the problem formulation to similar applications such as the dynamic allocation of virtual machines or containerised environments.

5. REFERENCES

[1] T. D. Braun, et al. Static resource allocation for heterogeneous computing environments with tasks having dependencies, priorities, deadlines, and multiple versions. *J. Parallel Distrib. Comput.*, 68(11):1504–1516, November 2008. doi: 10.1016/j.jpdc.2008.06.006.

[2] H. Hussain, et al. A survey on resource allocation in high performance distributed computing systems. *Parallel Computing*, 39(11):709 – 736, 2013. doi: 10.1016/j.parco.2013.09.009.

[3] M. Jette et al. Slurm: Simple linux utility for resource management. In *Proc. of ClusterWorld Conference and Expo (San Jose, California)*, 2003.

[4] S. Kiemle. Requirements for EO data processing farms. https://wiki.services.eoportal.org/attachments/HMA_WORKSHOP/20121012/Presentations-Thursday/PM/2012-10-11_DLR-Requirements-for-EO-Data-Processing-Farms_1.0.pptx, 2012.

[5] J. Polo, et al. Resource-aware adaptive scheduling for mapreduce clusters. In *Proc. of the 12th ACM/IFIP/USENIX International Conference on Middleware*, Middleware'11, pages 187–207, 2011. doi: 10.1007/978-3-642-25821-3_10.

[6] P. Soille. *IMAGE-2006 Mosaic: SPOT-4 HRVIR, SPOT-5 HRG, and IRS-LISS III Cloud Detection*. JRC, European Commission, 2008. doi: 10.2788/49355.

[7] P. Soille. Seamless mosaicing of very high resolution satellite data at continental scale. In P. Soille et al., editors, *Proc. of the 2014 Conference on Big Data from Space*, pages 222–223, 2014. doi: 10.2788/1823.

[8] S. Soner et al. Topologically aware job scheduling for slurm. http://www.prace-ri.eu/IMG/pdf/WP180.pdf.

[9] Q. Wu et al. Runtime task allocation in multicore packet processing systems. *IEEE Trans. Parallel Distrib. Syst.*, 23(10):1934–1943, 2012.

[10] Z. Zhan, et al. Cloud computing resource scheduling and a survey of its evolutionary approaches. *ACM Comput. Surv.*, 47(4):63:1–63:33, July 2015. doi: 10.1145/2788397.

STANDARDS FOR BIG GEO DATA: (HOW) CAN THEY HELP?

Peter Baumann

Jacobs University | rasdaman GmbH
Bremen, Germany
baumann@rasdaman.com

ABSTRACT

Flexible, scalable services on massive geo data receive much attention today. OGC has taken a lead in defining interoperable spatio-temporal datacubes (the coverage paradigm) together with flexible scalable Web services (the Web Coverage Service / WCS standards suite). This OGC Big Geo Data standard has attracted all major implementers of both open-source and proprietary servers and clients. We present status, uptake, and recent progress of the coverage datacube standards in OGC, ISO, and INSPIRE.

***Index Terms*—** OGC, Big Data, datacubes, coverage, WCS, rasdaman

1. INTRODUCTION

Sensor, image, image timeseries, simulation, and statistics data contribute significantly to the Big Data challenge in the Earth sciences. Serving them at a high service quality and at high speed is one of the key challenges for modern technology. In particular it becomes more and more clear that a zillion of single scenes is not the appropriate granularity for user-friendly offerings. A first step has been done some time back when seamlessly mosaicked maps enabled a smooth zoom and pan experience for users. Effectively, this has opened Earth Observation data a much larger, new community than just geo experts – we all embrace Google Maps and similar services today as part of our common Web experience. The same paradigm shift is now heralded by the datacube concept, but along the time axis: organizing all data from one satellite instrument into a single x/y/t datacube has the potential of simplifying access to multi-temporal data. Similar arguments hold for the vertical spatial axis.

The intercontinental EarthServer initiative [7] has taken a lead in enabling and promoting the datacube paradigm (Figure 1). Based on modern Array Database technology, large-scale data centers are establishing 3D and 4D datacubes for "any query, anytime" retrieval. Actually, this cube paradigm enables the EarthServer common backend technology, rasdaman, to perform manifold optimizations (see later) for smooth performance on datacube analytics and multi-cube fusion. Currently, the federation comprises ESA (operated through MEEOs.r.l.), the European Centre for Medium-Range Weather Forecast, Plymouth Marine Laboratory, National Computational Infrastructure Australia, and Jacobs University. Further data centers have already expressed interest in joining the federation.

All client/server interfaces of EarthServer rely exclusively on the OGC WCS and WCPS standards, plus WMS for flat map visualization; 3D visualization is done using NASA WorldWind coupled again through WCS / WCPS.

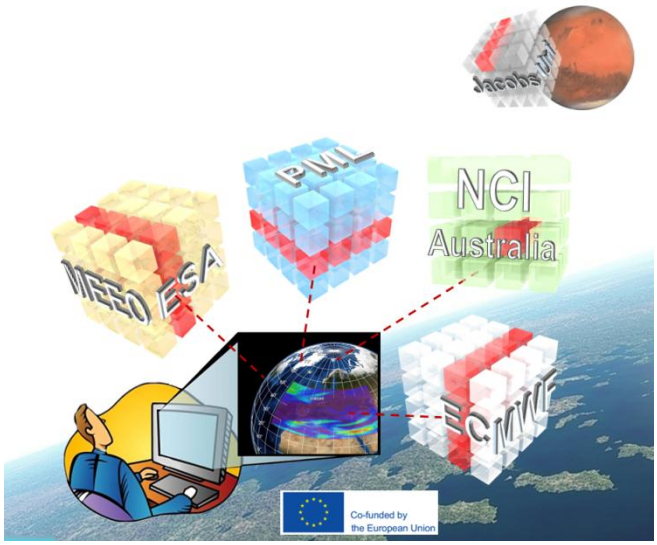

Figure 1: EarthServer datacube fusion based on OGC WCS and WCPS (source: EarthServer)

2. OGC COVERAGES

OGC and ISO together have established *coverages* as a unifying concept for spatio-temporal regular and irregular grids, point clouds, and meshes. Their abstract definition is laid down in ISO 19123 [5][10]. Being abstract implies: incompatible implementations are possible (typically visible by the fact that such services use their own homegrown clients and cannot use, e.g., standard open-source clients).

To remedy this, OGC has established a canonical coverage implementation model [2], derived from ISO 19123 and concrete, concise, and conformance testable down to pixel level. This coverage model supports both

regular and irregular spatio-temporal grids, point clouds, and meshes.

3. OGC WEB COVERAGE SERVICE (WCS)

The *Web Coverage Service* (WCS), whose mission is to provide streamlined functionality for spatio-temporal coverage, technically is subdivided into a common core and several extensions. WCS Core [2] consists of only the most fundamental functionality: extracting a coverage or a subset thereof (Figure 2). The coverage result can be delivered in any suitable encoding. In a vanilla WCS request data are guaranteed to be delivered unaltered.

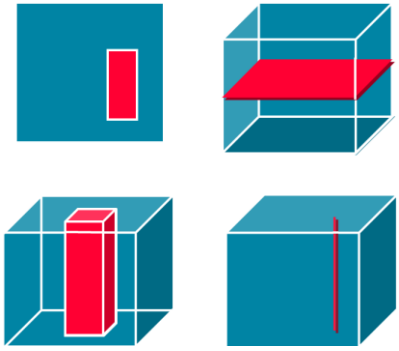

Figure 2: Subsetting of a 3-D coverage: Trimming (left) extracts a cutout which retains the number of dimensions, slicing (right) cuts out hyperplanes of reduced dimension (source: Wikipedia)

WCS extensions add further functionality facets to WCS, which an implementer may freely choose to support. Web Coverage Processing Service (WCPS) [1], for example, adds a declarative n-D geo raster query language into WCS.

4. IMPLEMENTING BIG DATACUBES

WCS Core Reference Implementation is *rasdaman*, the pioneer Array Database [8]. It offers an n-D array query language which on server-side is supported by a massively scalable array engine based on the "tile streaming" paradigm. Incoming queries undergo effective optimization, parallelization across node and heterogeneous hardware. Databases exceed 100 TB volumes [11], and single queries have been split over more than 1,000 cloud nodes [12].

5. STANDARDIZATION OUTLOOK

Since its adoption in 2010, coverages and WCS have experienced manifold, continuously growing take-up from both open-source and proprietary implementers, and recently from further standardization bodies. INSPIRE currently is establishing WCS as a Coverage Download Service.

In OGC, version 1.1 of the coverage model has entered the adoption process, adding partitioned coverages, an efficient "geometry/value pair" representation, irregular grids generalizing and encompassing GML 3.3 and SensorML, and interpolation. ISO TC211 has endeavored into adoption of CIS 1.1, which is to become ISO 19123-2 and will be followed by WCS. Also the World-Wide Web Consortium (W3C) is discussing coverages.

6. CONCLUSION

OGC coverages represent a concrete, interoperable data model which unifies n-D regular and irregular grids, point clouds, and meshes – obviously main contributors to today's Big Geo Data. The OGC WCS standards suite offers a wide range of functionality, from simple extraction and download up to complex ad-hoc analytics. Flexibility and scalability of the WCS suite has been demonstrated in practice through large-scale services. EarthServer is currently establishing datacubes exceeding 1 PB each.

REFERENCES

[1] P. Baumann: OGC Web Coverage Processing Service (WCPS) Language Standard, version 1.0. OGC 08-068r2

[2] P. Baumann: Web Coverage Service (WCS) Interface Standard – Core, version 2.0. OGC 09-110r4

[3] P. Baumann: OGC GML 3.2.1 Application Schema - Coverages, version 1.0. OGC 09-146r2, 2010

[4] P. Baumann: The OGC Web Coverage Processing Service (WCPS) Standard. Geoinformatica, 14(4)2010, pp. 447 - 449

[5] ISO: 19123 Geographic information – Schema for coverage geometry and functions. IS 19123:2004

[6] N.n.: Earthlook. http://earthlook.org, seen 2015-06-20

[7] N.n.: EarthServer – Datacubes at Your Fingertips. www.earthserver.eu, seen 2015-11-06

[8] N.n.: rasdaman. www.rasdaman.org, seen 2015-06-20

[9] OGC: WCS standards suite. www.opengeospatial.org/standards/wcs; seen 2015-06-20

[10] OGC: The OpenGIS Abstract Specification Topic 6: Schema for coverage geometry and functions, version 7.0. OGC 07-011

[11] P. Baumann, P. Mazzetti, J. Ungar, R. Barbera, D. Barboni, A. Beccati, L. Bigagli, E. Boldrini, R. Bruno, A. Calanducci, P. Campalani, O. Clement, A. Dumitru, M. Grant, P. Herzig, G. Kakaletris, J. Laxton, P. Koltsida, K. Lipskoch, A.R. Mahdiraji, S. Mantovani, V. Merticariu, A. Messina, D. Misev, S. Natali, S. Nativi, J. Oosthoek, J. Passmore, M. Pappalardo, A.P. Rossi, F. Rundo, M. Sen, V. Sorbera, D. Sullivan, M. Torrisi, L. Trovato, M.G. Veratelli, S. Wagner: Big Data Analytics for Earth Sciences: the EarthServer Approach. International Journal of Digital Earth 0(0)2015

[12] A. Dumitru, V. Merticariu, P. Baumann: Exploring Cloud Opportunities from an Array Database Perspective. Proc ACM SIGMOD Workshop on Data analytics in the Cloud (DanaC'2014), June 22 - 27, 2014, Snowbird, USA

DISCRETE GLOBAL GRID SYSTEMS FOR HANDLING BIG DATA FROM SPACE

Matthew B. J. Purss[1]
Perry Peterson[2]
Robert Gibb[3]
Faramarz Samavati[4]
Peter Strobl[5]

1 Geoscience Australia, Canberra, Australia
2 PYXIS, Kingston, Ontario, Canada
3 Landcare Research, Palmerston North, New Zealand
4 University of Calgary, Calgary, Alberta, Canada
5 European Commission, Joint Research Centre, Ispra, Italy

ABSTRACT

The rapid growth in global sensor networks is leading to an explosion in the volume, velocity and variety of geospatial and geoscientific data. At the same time, we are experiencing the rapid and unprecedented consequences of environmental changes. This coupled with the increasing integration of geospatial data into our everyday lives is also driving an increasing expectation of spatial information on-demand, and with minimal delay. It is hoped that the data and information describing these environmental challenges can be transformed into the knowledge and decision support systems that will help mitigate the cost to society. However, there remains a gap between these expectations and the present reality of easily accessible information products and tools that help us understand the Earth. A solution can only be achieved through the development and implementation of common analytical frameworks that enable large scale interoperability of 'Big Earth Data' across multiple data infrastructures.

***Index Terms*—** Discrete Global Grid System, Digital Earth Reference Model, Open Geospatial Consortium, Spatial Analysis, Big Earth Data

1. INTRODUCTION

The Open Geospatial Consortium (OGC) has recently introduced a new Digital Earth reference standard for Discrete Global Grid System (DGGS) infrastructures. A DGGS represents a spatial reference system that uses the Earth as its organizing structure [1]. DGGS use a hierarchical tessellation of cells to partition and address the globe [2] and are characterized by the properties of their cell structure, geo-encoding, quantization strategy and associated mathematical functions [2]. DGGS are analogous to any discrete "Digital" data structure – in contrast to a continuous "Analog" model of the Earth represented by geographic coordinates.

2. STANDARDIZING DGGS

There are many possible discrete global grids, each with their own advantages and disadvantages. Standard criterion for a discrete global grid are well developed by both Michael F. Goodchild [3] and Jon Kimerling [4]; the foremost requirements being a tessellation of cells that exhaustively cover the globe with each cell having equal area and representing a single point [5]. The points and cells of the various resolution grids which constitute the grid system form a hierarchy which displays a high degree of regularity [6].

Any tessellation of the Earth does not necessarily produce a DGGS. Single resolution computational grids are not sufficient to constitute a DGGS. Spatial data structures used to organize map tiles or optimize rapid spatial search cannot be considered to qualify as a DGGS in and of themselves; although DGGS often utilize hierarchical indices to identify a cell, the primary feature of the DGGS is the cell geometry not the optimization of a spatial query. Further, DGGS have data independent geometry [7] – their geometry is not formed to optimize a balanced search like R-Trees or maximal spacing of data as generated by Voronoi diagrams.

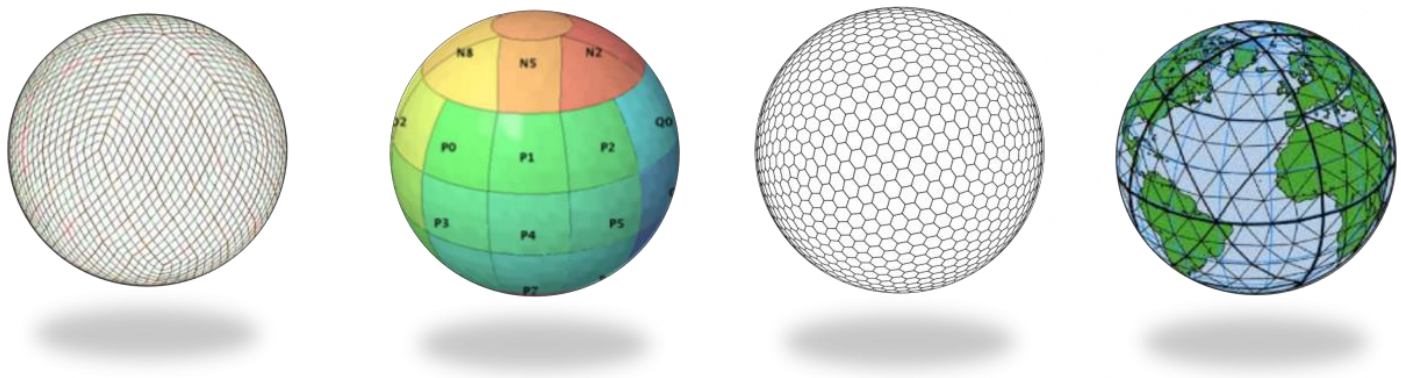

Figure 1 - Examples of conformant OGC DGGS implementations. From left-to-right: ISEA Triangular Grid; SCENZ-Grid; ISEA Hexagonal Grid; and, Quaternary Triangular Mesh (reproduced courtesy Geoff Dutton)

The only perfectly regular partitioning methods for the surface of a sphere (or ellipsoid) are based on the inverse projection of the following Platonic solids: the tetrahedron, cube, octahedron, dodecahedron, and icosahedron. This method of mapping the faces of a base polyhedron to the surface model of the Earth creates an inverse projected coordinate reference system. Traditional GIS and image analysis packages that assume flat earth geometries will need to adapt to support this new construct that more closely represents the earth. Standardizing DGGS will enable this to occur.

The new OGC DGGS Core Standard defines the DGGS core data model and the core set of requirements to which every OGC DGGS encoding must adhere. This includes specification of:

1. A concise definition of the term Discrete Global Grid System as a spatial reference system;
2. The essential characteristics of a conformant DGGS; and,
3. The core functional algorithms required to support the operation of a conformant DGGS.

3. APPLYING DGGS TO THE WORLD OF BIG EARTH DATA

The hierarchy of tessellations within a DGGS enables rapid aggregation and decomposition of data necessary for online access and transmission speeds. Geospatial data values from any source, type, format, spatial reference, spatial scale, or frequency can be held within a DGGS. With the trend to more open on-demand systems, DGGS provide a user centric approach where end-users can search, explore and integrate data from multiple content providers simultaneously. A DGGS provides a framework where the three fundamental questions of geospatial analysis can be answered as simple set theory operations (i.e. “Where is it?”, “What is here?” and “How has it changed?”). Big Earth Data that is aligned to a DGGS is easy to access, store, sort, process, transmit, integrate, visualize, analyse and model.

Visualization of complex analyses can be an effective method of influencing a multitude of policy and decision making processes which impact global issues. A DGGS represents a major advancement in our ability to understand the Earth – where data can be accurately represented and integrated with minimal distortion.

The new OGC DGGS standard provides the basis for adopting this new digital Earth approach to geospatial decision-making. Examples from a number of conformant DGGS implementations using 'Big Earth Data' demonstrate the value, application and potential of the DGGS approach.

Figure 1 shows visualizations of some examples of conformant OGC DGGS implementations. These include:

1. The ISEA Triangular mesh DGGS;
2. The SCENZ-Grid DGGS;
3. The ISEA Hexagonal mesh DGGS; and,
4. The Quaternary Triangular Mesh.

4. ACKNOWLEDGEMENTS

The development of the candidate OGC core standard for DGGS, which forms a basis for this paper, has been a collaborative effort involving contributions from many individuals and organisations involved with both the OGC and the DGGS Standards Working Group.

This paper has been published with the permission of the Chief Executive Officer, Geoscience Australia.

5. REFERENCES

[1] Goodchild, M. F., “Geographical information science”, *International Journal of Geographical Information Systems*, 6(1): pp 31–45, DOI:10.1080/02693799208901893, 1992

[2] Purss, M. B. J. (ed.), “OGC Discrete Global Grid System (DGGS) Core Standard”, OGC 15-104r3, 2016.

[3] Goodchild, M. F., “Discrete global grids for digital earth”, *International Conference on Discrete Global Grids*. Santa Barbara: University of California, Santa Barbara, http://www.ncgia.ucsb.edu/globalgrids/papers/goodchild.pdf , 2000

[4] Kimerling, A. J., Sahr, K., White, D., & Song, L., “Comparing Geometrical Properties of Global Grids“, *Cartography and Geographic Information Systems*,26(4), pp 271-88, 1999, DOI:10.1559/152304099782294186

[5] Tobler, W., & Chen, Z.-t., “A Quadtree for Global Information Storage”, Geographical Analysis, 18(4), pp 360-371, DOI:10.1111/j.1538-4632.1986.tb00108.x, 1986

[6] Clarke, K. C., "Criteria and Measures for the Comparison of Global Geocoding Systems", *International Conference on Discrete Global Grids*. Santa Barbara: University of California, Santa Barbara. http://ncgia.ucsb.edu/globalgrids-book/comparison/, 2000

[7] Peterson, P. R., “Discrete Global Grid Systems.” *In The International Encyclopedia of Geography*, edited by Douglas Richardson. Malden, Oxford: John Wiley and Sons, Ltd., (Forthcoming)

APPLYING GMLJP2 TO IMAGERY FROM SPACE

Lucio Colaiacomo, Raul Alonso Reyes

EU SatCen, Torrejón Air Base, Torrejón de Ardoz, Spain

ABSTRACT

The EU Satellite Centre delivers data resulting from satellite imagery analysis, mainly consisting of annotated imagery. The envisaged solution is to incorporate compressed imagery and annotations inside the same file being delivered to users with a jpx or jp2 extension. This can be achieved with the new GMLJP2 Version 2 format.

To fulfill these requirements, the SatCen developed a software solution based on open source software to implement the standard GMLJP2 (geographic markup language in jpeg2000), version 2. One of the main issues in dealing with satellite imagery is the size and number of files included in a typical GEOINT product. With jp2 techniques, a compression factor of 20 can be easily reached thereby facilitating transmission with minimal quality loss, and can also be used in lossless mode with a lower compression factor.

The software solution "ViewIt" is based on the GDAL library and uses openjpeg2000 and jasper libraries to enable compression decompression between geotiff and jp2. It implements the latest version of the GMLJP2 standard Version 2 and is compatible with WCS 2.0, GML 3.2.1, WFS 2.0 and KML.
With this new version of the standard, it will be possible to create compressed images with annotations and/or features. It also enables the embedding of multiple coverages in a single file and the management of different metadata formats ISO, EOP, etc. for each coverage.
The solution implemented supports multiple file encoding and decoding for publication in WCS and WFS (transactional) servers.

***Index Terms*—** satcen, gmljp2, wcs, wfs, wms, sensorML, gmlcov, CIS

1. INTRODUCTION

The SatCen manages high resolution satellite imageries, mosaics and digital elevation model in huge quantities as well as maintaining the metadata for those data. As a result, a suitable format for managing data compression (lossless or lossy) and metadata maintenance in a standardized format was required. Compressing data can result in a lossless reduction factor of 3-4 for imagery, and up to 20 for lossy compression, while preserving suitable quality for interpretation. The reduction in required storage space is evident, as well as the saving in file transmission times due to the reduced size.

2. GMLJP2 STANDARDIZATION

This can only be achieved if all the data can be processed, and most importantly, the metadata describing the data remains preserved. For this reason, a list of use cases to undertake validation has been agreed with other agencies and standardization bodies.

At the same time a standardization working group was formed at the OGC (Open geospatial Consortium) to ensure a proper process is followed, and to achieve consensus about the standard (GMLJP2 V2).
GMLJP2 version 2 bridges an interoperability gap between data and services, because it is based on GMLCOV and includes GML (3.2 equivalent to ISO 19136) and JPEG2000 and it is easily implemented inside WFS-T and WCS-T, making it easier to consume this data in OGC services. Moreover, SensorML integration has been studied and demonstrated in other Engineering Reports in OGC Testbed-11 so that it will be possible to handle a large variety of georefereanceable imagery and subsequently be able to widen the usage of the standard. The inter-standards coordination addressed by this working group demonstrates the feasibility of this format.
JPEG2000, being so widely used in many other markets (medical, film industry, etc.), could open the door to geo enabling, or geo positioning imagery of unusual sources from geospatial domain.

3. SOLUTION

Definitions

“The term "grid coverages" refers to satellite images, digital aerial photos, digital elevation data, and other phenomena represented by values at each point in a "raster" coordinate system (as opposed to "vector" geodata, in which digital map information is represented using polygons and lines). The specification describes an open interface that provides communication between software systems for purposes of requesting, viewing, and performing certain kinds of grid coverage analysis such as histogram calculation, image

doi: 10.2788/854791

covariance and other statistical measurements." (source OGC)

"Coverages represent digital geospatial information representing space/time-varying phenomena [4] – which is identical to ISO 19123 – defines an abstract model of coverages. This is concretized by the Geography Markup Language (GML) 3.2 [07-036], an XML grammar written in XML Schema for the description of application schemas as well as the transport and storage of geographic information.[2]
However, the definition contained in GML 3.2 has turned out to contain insufficient information to describe coverage instances in a flexible, interoperable, and harmonized manner. To remedy this, the document on hand defines a GML Application Schema for coverages by applying the following enhancements to the GML 3.2 Coverage data type:

- A mandatory element rangeType has been added to carry information about the range value data structure of a Coverage.
- The property CoverageFunction, which in GML 3.2 is associated with every subtype of Coverage, is moved up into AbstractCoverage in the coverage type hierarchy of the standard on hand. This semantic-preserving modification does not impact instance documents.
- A metadata hook has been added which allows definition of application specific supplementary information to be transported with a coverage.
- The grid coverage types are subtypes of AbstractCoverage rather than being subtypes of DiscreteCoverage as in GML 3.2." (source OGC)

Rectified Grid

As per ISO 19123, "Grid for which there is an affine transformation between the grid coordinates and the coordinates of an external coordinate reference system". A rectified grid is defined by an origin in an external coordinate reference system, and a set of offset vectors that specify the direction and distance between grid lines within that external CRS. "[1]

ReferenceableGrid

As per ISO 19123, "a referenceable grid is associated with a transformation that can be used to convert grid coordinate values to values of coordinates referenced to an external coordinate reference system"[1]. In general referring to a sensor capture of a portion of the earth surface and mapped into a digital "grid" and these samples equally spaced. The captured value can be elaborated (radiometric, or geometry processing) to produce derived products but are not yet related to a specific point of the earth surface. The geo-referentiation of the image on the earth surface is obtained through concatenation of coordinate reference frames (sensor, platform, earth) so to obtain the real position of the digital information on the corresponding point of the earth which belongs to. In order to be able to manage this type of data, it was required to merge inside the GMLJP2 standard the gmlcov and SensorML to proper handle such information.

Problem

The original problem was to compress decompress (codec) raw imagery coming from heterogeneous sources:

- High resolution optical imagery 8/16/24/32 bit encoding
- High resolution SAR images
- Digital elevation model
- Large files (bigger than 2GB compressed)

Additionally, georefereanceable imagery and annotations must be supported.
In particular to fulfill the last it is being produced and extension to the GMLJP2 standard Version 2.

Use Cases

Several use cases have been tested in order to check the validity of the implementation done with GDAL.

- Gray tone: Raster image with a single band is compressed.
- RGB Colour: Raster image with three bands denoting Red, Green and Blue respectively, is compressed.
- Multispectral: A multispectral image is compressed.

- Raster map: A thematic map where the pixel values denote different objects in the map.
- Ortho-rectified: An ortho-rectified image has its ortho-rectification parameters encoded in GML.
- Embed metadata: The image's metadata is embedded in the image file.
- WMS or WCS client-server query with GMLJP2 file as response:
  A WMS client sends a GetMap request to a WMS server or a GetCoverage request to a WCS-server, and gets a GMLJP2 file as a response. Usage of WCS/WFS transactional.
- Some tests have been done with the JPIP framework:

Metadata query to a JPIP-server; JPIP client-server metadata request returning JPIP data-bin containing only metadata (tested partially).

SatCen approach

SatCen started to implement the standard using a software called ViewIt (see section 4) and also using the OGC Testbed 11 in which a prototype of client server application using GMLJP2 compression has been implemented. The results of this activity are briefly reported later on.
In the picture below an UML diagram representing the GMLJP2CoverageCollection feature type.
Particular attention has been given to the compression of floating point imagery (radar). The method of compressing using jpeg2000 with the transformation in sign, exponent, mantissa has been applied.

The activity developed at SatCen included an implementation of this standard into the GDAL library. This has been achieved by enabling a definition file in which is possible to define a gmljp2 box that typically contains a grid coverage (rectified or not) with all the SRS information associated and the geo - transformation matrix. In this box it is also possible to add metadata, vector features (GML feature collection), annotation (KML) or any other XML content as an extension, thus metadata can be dynamically generated from a template file (with an xml structure) and an xml source file.
(Some open source examples are available at http://www.gdal.org/frmt_jp2openjpeg.html)
The result is the transformation of metadata into gmljp2:eopMetadata (earth observation profile), and also other types of metadata are supported.
The work done in this framework together with the important activity made at DGIWG level (Ulf Tennfors) in implementing a parser able to transform GMLJP2 V1 into GMLJP2 v2, and the GDAL library implementation, facilitate the adoption of this standard.

In order to be fully compatible with all the data types, it has been required to test the floating point data compression.
Floating point imagery codec.
The procedure followed to obtain this jp2 image is with this function: (following IEEE754):

```
void IEEE754(float f, unsigned *mantisa,int
*exponent,int *sign){
        Fl u x;
        unsigned wk = 0;
        x.f = f;
        wk = x.ui;
        if (wk & 0x80000000)
                *sign = 1;
        else *sign = 0;
        unsigned bit = wk & 0x07FFFFF;
        *mantisa = bit;
        int exp = (wk >> 23) & 0x0FF;
        *exponent = exp;
}
```

The mantissa, exponent and sign of each float32 pixel were obtained.
For each pixel a raw formatted file was written with this sequence:
Sign, exponent, mantissa1, mantissa2, mantissa3

After this, the Luratech framework was used to obtain the final jp2 file with the following command:

c -i C:\TEMP\srtm\19TBG505575_mantisse_raw_5b.bil -raw 1500 1500 -5 8 0 -o C:\TEMP\srtm\19TBG505575_lrw_mantisse_raw_5b.jp2 -lic 3002570894 3833631664

After several tests we finally managed to compress into a 3 bands of custom datatype, results from compressing with a factor 1:3 a file with 1,8,23 bit components and the results were quite promising.

Performance:

On Intel Xeon CPU E5-1620 0 @ 3.6 GHz
16MB GB memory, 64 bit operating system
Openjpeg library
1 GB imagery lossy compression factor 25 -> 3 minutes
1 GB imagery lossless -> 1.5 minutes
3 GB imagery lossless -> 6 minutes
3 GB imagery lossless compression factor 25 -> 10 minutes

Test results

The GDAL implementation, which includes OpenJPEG and Luratech libraries produced codestreams that can also be read by ArcGis and Erdas Imagine.
The following modifications to GDAL have been made:

- Support for creation (conversion to) GMLJP2 V2 format is built into a SatCen version of GDAL by extending the JPEG2000 driver. The library is built into a 64-bit environment because JasPer needs a large amount of RAM to operate.
- Compiling for 64 bits allows addressing beyond the 2Gb limit set for the WIN32 application heap allocations (or 4Gb for WIN32 applications compiled with /LARGEADDRESSAWARE switch and running on a 64-bit system.)

The library can use OpenJPEG, JasPer and Luratech seamlessly.
During the OGC Testbed the GDAL library has been compiled for Android and iOS with OpenJPEG, facilitating the use of handheld devices. This extends the range of

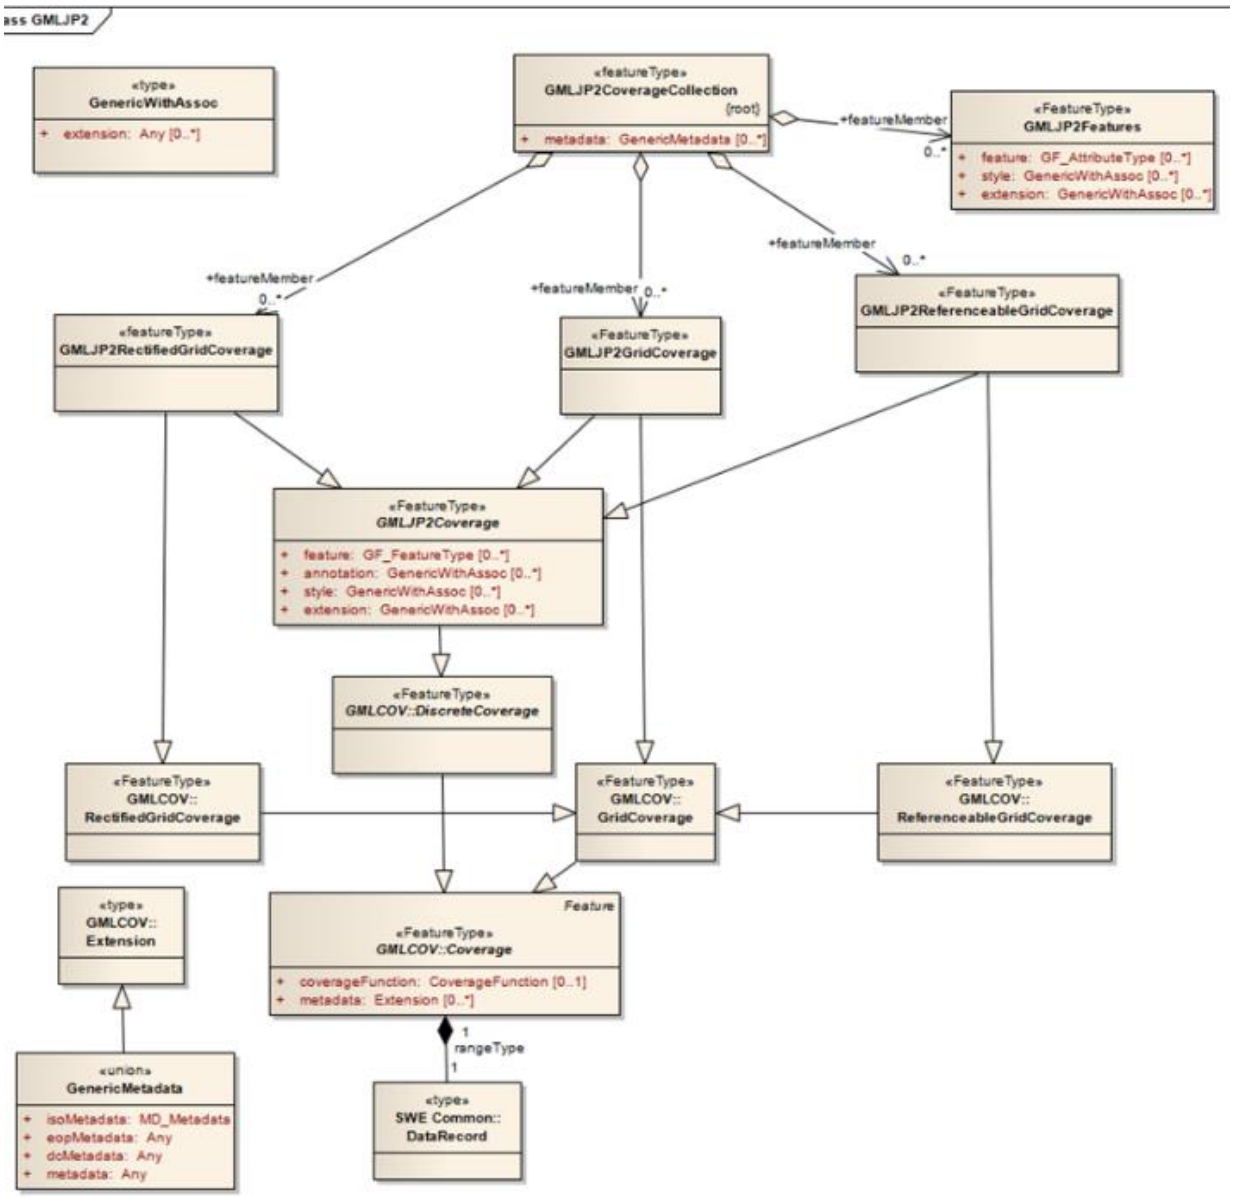

imagery production possibilities to include imagery from drones, etc. which should ease the introduction to the market of this standard. The integration inside the standard of all the elements corresponding to the GMLCOV specification make the communication with WFS-T and WCS-T server straightforward. The testbed implementation involved creating a client for android and iOS supporting communication with WFS-T (Geomatys) and WCS-T (Rasdaman).

Future activities for improvements can include:

- Imagery with associated Sensor Model (for Georeferenceable imagery test)
- Terrain Elevation data: GMLJP2,
- Data access with WCS implementing GMLJP2 parsing (GMLCOV for GMLJP2)
- Data integration for map composition creation (equivalence with NITF datasets)
- It will be necessary to study the compatibility of Luratech[3] for different datatypes for each band. Then the final file could be
    - 3 bands (1bit, 8 bits, 23 bits)

Further potential improvements would be the 23 bits compression of the mantissa component with the final result of 3 bands (1bit, 8 bits, 23bits)

doi: 10.2788/854791

4. VIEW-IT

The SatCen developed a software application based on open source code to implement the specifications previously described. The software is available from the OGC portal and implements functionality available from GDAL, in addition to geo referencing tools and metadata editing. The codec enables compression/decompression of a Region Of Interest (ROI), specific metadata (i.e. processing algorithm), decompression in geotiff, nitf, etc...
With the ViewIt application it is possible to produce annotated images /see example below).

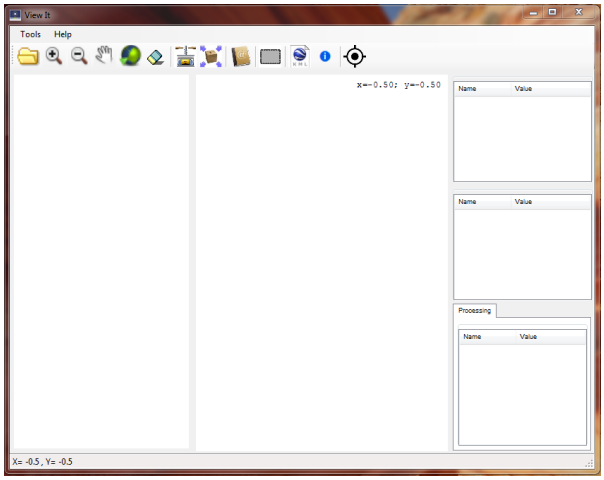

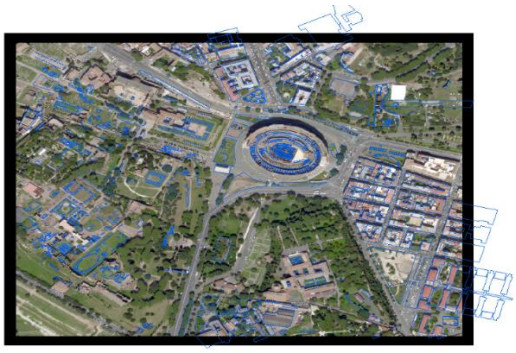

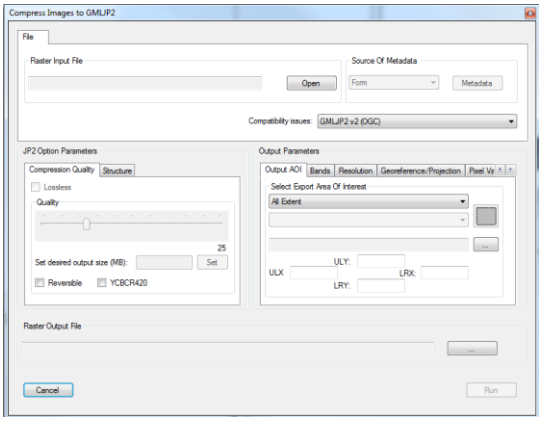

5. CONCLUSIONS

The implementation done at SatCen of the the new version of the standard GMLJP2 V2 demonstrate its feasibility and it is beneficial in terms of metadata transport in coding and decoding imagery coming from different sources (satellite, aerial, etc..) and in various input formats. The work done in managing rectified and/or refereanceable imagery with or without annotation give this format enough flexibility for its adoption. An interesting development could be the implementation using GPUs.

5. REFERENCES

[1] E.Hirschorn, P.Baumann, “15-065r1 Testbed 11 Refereamceable Grid Harmonization Engineering Report,” *OGC*, 2015.

[2] L.Colaiacomo, J.Maso, E.Devys, “GML in JPEG2000 GMLJP2 Encoding Standard Part 1 Core”, OGC , 2015.

[3]Luratech https://www.luratech.com/en/products/download-pdf-compressor-desktop/luratech-jp2-sdk/

[4] OGC Abstract Topic 6 [OGC 07-011]

6. ACKNOWLEDGMENTS

Special thanks to Diarmaid Foley.
Even Rouault for the implementation into the GDAL library.
Galdos Inc for the GML libraries and the JasPer implementation.
Alberto San Narciso for implementation of the mobile application (iOS, android) and for the testing environment.

OGC TESTBEDS AND BIG DATA FROM SPACE: EXPERIMENTAL BIG DATA DESIGN AND EXPLORATION

Ingo Simonis

Open Geospatial Consortium Europe
London, UK

ABSTRACT

The amount of data produced by earth observing missions such as ESA's Sentinels or NASA's more than a hundred currently active missions is growing at a tremendous rate. These datasets have themselves become extremely valuable assets, and their strategic importance in research and science is huge. Though the handling of the enormous data masses within the space administrations' internal processing and storage systems represents a number of research topics on its own, several others arise in the context of making the data available, exploring the data, and processing the data. These aspects are traditionally covered by the Open Geospatial Consortium (OGC). The OGC develops a number of data description, access, and handling standards for these purposes. The OGC faces new challenges, such as those addressed herein and commonly referred to as Big Data, with its Interoperability Program (IP). This paper discusses the OGC interoperability program initiative 'Testbed', a highly experimental environment that allows better understanding of the potential of Big Data, describes Big Data application scenarios, and experiments with potential solutions for Big Data induced challenges that are eventually transferred into the consensus-driven OGC standardization program for further standardization. The paper concludes with a short discussion on the value of Big Data from space once it's integrated with ground-based, crowd-sourced citizen science data.

Index Terms— Spatial Data Infrastructure, Big Data, Standards, Citizen Science

1. INTRODUCTION

Big data has been a focus of research for a number of years, but has become even more exciting lately as the amount of data produced by earth observing missions - such as ESA's Sentinels or NASA's more than a hundred currently active missions - is growing at an unprecedented rate. The availability of the data has resulted in a proposed paradigm shift from model-driven research to data-driven research [1]. These datasets have themselves become extremely valuable assets, and their strategic importance in research and science is huge. Almost all ground-based research benefits from the enormous data produced by earth observing satellites, and Big Data has influenced recent research in a wide range of fields, including ecology, city planning [2, 3], land cover analysis, and bio-geography [4, 5].

The data sets produced by the various missions have become so large that traditional data management systems and analysis tools are overwhelmed. Similarly, traditional Service Oriented Architecture approaches, implemented in Spatial Data Infrastructures like INSPIRE, NSDI, or GEOSS, are pushed to their limits when it comes to Big Data in the magnitude provided by ESA, NASA, JAXA and others. The handling of the enormous data masses within the space administrations' internal processing and storage systems represents a number of research topics on its own, and several others arise in the context of making the data available, data exploration, and data processing. These aspects add value to the originally meaningless data and are traditionally covered by the Open Geospatial Consortium (OGC). The OGC, together with ISO TC211, develops a number of data description, access, and handling standards for these purposes.

The OGC faces new challenges in its Interoperability Program (IP). This paper discusses the OGC interoperability program initiative 'Testbed' - a highly experimental environment that allows better understanding of the potential of Big Data. The Testbed describes Big Data application scenarios and experiments with potential solutions for Big Data induced challenges that are eventually transferred into the consensus-driven OGC standardization program for further standardization. Testbeds have been executed almost yearly since 1999. The current Testbed-12 has a budget of several million USD and tackles a number of Big Data challenges. These include aspects such as: cataloging Big Data; handling of self-descriptions of Web services that provide access to, or processing of, Big Data; quality matrices of multi-tile images; offline usage of big data; transfer of big data into tile stores; and compression and generalization techniques to handle "low" bandwidth situations.

The paper concludes with a short discussion on the value of interoperable solutions for other domains that are interacting with Big Data from space. As an example, we use the

This research has been supported by the European Commission through the Citizen Observatory Web (OBWEB) project, grant agreement no. 308513

citizen science domain, where a constantly growing community of volunteers contributes to scientific projects by providing real-world observations on a voluntary basis. Citizen science data holds a great potential to provide 'ground truthing' for earth observation data coming from satellites, or for verification or calibration of ground devices by fusing human-sourced data with information obtained by sensors.

2. OGC TESTBED-12

OGC Testbeds are an essential component of the OGC Interoperability Program. The OGC Interoperability Program, in general, provides global, hands-on, collaborative prototyping for rapid development and delivery of proven candidate specifications to the OGC Standards Program, where these candidates can then be considered for further action. In Interoperability Program initiatives, participants team together to solve specific geo-processing interoperability problems posed by the initiatives sponsors. These initiatives are designed to encourage rapid development, testing, validation and adoption of open, consensus-based standards. The OGC Interoperability Program policies and procedures, including descriptions of other types of initiatives, can be found at [6]. Testbeds are the experimental laboratories among the interoperability program activities. New challenges are tackled, and new solutions elaborated and tested in a prototype environment. Testbed results are documented in engineering reports that serve as direct input into the consensus-driven OGC standardization program. Testbeds analyze and explore the potential of new technologies as they seek solutions for already identified problems.

The current Testbed-12, whose first part started in January 2016 with the majority occurring in early March, is investigating and developing standards to solve Big Data integration and usage scenarios. These will, in turn, help explore the full potential of Big Data in various contexts, from general earth observation to highly specialized citizen science and crowd sourcing projects. The Testbed is scheduled to end with a demonstration meeting at the end of 2016. All experiences and results will be made available publicly. Up-to-date information as well as opportunities to join the ongoing testbed, are documented online [7].

The following sections will introduce the work items addressed in Testbed-12. They cover a wide range of Big Data aspects, starting with the discovery phase, where catalogs advertise the available data and processing options (section 'Catalogs') and self-descriptions allow exploring the potential of data access and processing Web services (section 'Capabilities'). Testbed-12 includes work items that help referencing Big Data subsets (section 'OWS Context') and others that support availability of Big Data subsets in offline situations (sections 'GeoPackage' and 'Tile Store'). The section 'Quality Models' introduces new concepts to define quality parameters of multi-tile images where individual tiles are provided by different missions or satellite overflights.

3. CATALOGS

Current Catalog solutions are highly dependent upon the metadata model employed for the service and data descriptions. The catalogs themselves have been designed to operate in a rather static environment, often with humans adding new entries following time consuming processes. Cross-catalog harvesting has been implemented, but dynamic updating of catalogs based on incoming new data and derived signals is still rare. The metadata models have been optimized to describe individual data sets, but do poor jobs for relationship management across data sets. This is particularly true for current catalogs in the context of standardized spatial data infrastructures that make use of the ISO metadata model defined in ISO 19115 [8]. The W3C has released the DCAT recommendation in early 2014 [9], an RDF vocabulary designed to facilitate interoperability between data catalogs published on the Web. In the context of bid data, it is necessary to evaluate what role DCAT can play in the integration of different catalogs. At the same time, catalog auto-entries need to be enabled to allow registering signals detected in big data, references for ground truthing using big data and citizen science data, and the various quality assurance levels need to be supported.

4. CAPABILITIES

Capabilities documents are used to describe services and their data or processing capacities in all OGC Web service standards. With the introduction of Big Data, extremely large amounts of data holdings may be available on a service or in a Cloud environment that may be scattered across multiple servers. With each server offering a high number of data sets that need to be advertised in the Capabilities document, capabilities documents tend to grow in size beyond useful limits. Paging mechanisms can only mitigate some issues, but new mechanisms and approaches are required to make service self-descriptions fit for Big Data. Testbed-12 will explore the issues induced by Big Data and explore possible mitigation solutions that scale with the ever growing amount of data.

5. OWS CONTEXT

The OGC Web Services Context Document (OWS Context) was created to allow a set of configured information resources (service set) to be passed between applications primarily as a collection of services. These services include resources such as OGC Web Feature Service (WFS), Web Map Service (WMS), Web Map Tile Service (WMTS), Web Coverage Service (WCS) and others. Referenced in a OWS Context document, clients can integrate any number of services into into a 'common operating picture', as one client can make its view

available to other clients by defining the actual services and rendering information.

To address Big Data challenges, OWS Context needs to support use cases that include distributed search results (i.e. subsets of Big Data) to be integrated into service description and processing chains. Instead of referencing a number of data services only, OWS Context for Big Data needs to reference Big Data subsets that are then further processed by additional services. At the same time, the often dynamic nature of Big Data needs to be taken into consideration, i.e. a client watching the latest data might need to reference a historic data sets to establish the same view on another client at another point in time. Additionally OWS Context can deliver a set of configured processing services (Web Processing Service (WPS)) parameters to allow the processing to be reproduced on different nodes.

6. GEOPACKAGE

The OGC GeoPackage standard supports the ability for the military (and others) to load tiled map/image data and vector data onto a hand-held device and utilize that data in a disconnected or limited connectivity environment. The standard has been incrementally improved through the efforts of the Testbed activities each year. Current Testbed-11 efforts have enabled in-the-field data collections to update the database with the most current information. Testbed-11 also developed a peer-to-peer in-the-field update via Bluetooth connection.

Testbed-12 shall utilize the GeoPackage as the implementation of Vector Tiling requirements. This allows to overlay raster-based satellite data with extracted or externally generated vector data sets in a common format and container. For that reason, Testbed-12 needs to develop a consistent tiling scheme to support both raster and vector tiling. This new scheme needs to be evaluated under real world conditions. The Testbed-12 scenario identifies a routing situation where a car gets abandoned and the further routing calculation takes satellite-based coverage data into account in order to calculate the best route between two locations.

Further on, Testbed-12 will evaluate commercial solutions such as the Common Map API (CMAPI) to further explore the combination of Big Data sets with commercial solutions from the mesh-up domain. The Common Map API provides a standardized approach for web applications from multiple organizations to visualize and interact with data on a common view without compromising application encapsulation and portability (cmapi.org). These research activities address the growing demand for agile integration of commercial large scale systems such as ESA's Copernicus platform or NASA's Atmospheric Science Data Center (ASDC) or Global Imagery Browse Services (GIBS) with state of the art tools from the mesh-up and rapid deployment communities.

7. TILE STORES

Though Big Data is commonly considered as data that requires processing as close to the data as possible, data exchange of Big Data subsets remains an important requirement for many applications and usage scenarios, in particular in situations with limited bandwidth. Testbed-12 will assess a file format (database format) for exchange of a large (global/regional) tile stores. This work item starts with the evaluation of solutions such as simple reuse of what is defined for raster tiling in GeoPackage and applying it to a PostgreSQL database instead of the GeoPackage standard container Sqlite DB. Additional solutions depend on the first evaluation results.

8. GEOSPATIAL IMAGERY QUALITY

The last decade has seen a proliferation of sensors on various platforms (satellites, aerial, UAVs) that collect imagery at multiple scales/resolutions. It is estimated that several hundred small satellites and UAVs will be added to this asset base in the next 5-10 years. Currently, there is no consistent quality framework that will allow end users to compare imagery from multiple sources across different quality attributes, in order to have a holistic view of imagery value as it may apply to a particular set of requirements.

Testbed-12 shall develop a quality framework called A3C (Accuracy, Currency, Completeness, and Consistency) that can be used to compare imagery from multiple sources. Accuracy refers to spatial accuracy of a location derived from the pixel in X,Y dimensions and potentially in the Z dimension. Currency refers to the temporal extent of the imagery products used to cover the associated area, since multiple dates of collects are typically required to cover a large area. Completeness of imagery products refers to quality metrics including cloud cover, sensor specs on collection geometry, temporal range of the data, other spectral bands if any, radiometric depth of the pixels, etc. Finally, the Consistency metric describes the consistency of colors, relative accuracy over time and over different sensors, spectral and spatial error propagation from collection to production, etc.

9. CITIZEN SCIENCE INITIATIVES

Citizen science initiatives have a great potential to serve as ground truth for big data sets. Several initiatives are under way, e.g. NASA's 'cloud of the day' project where volunteers take a picture of the clouds from the ground, while satellites take pictures from the same clouds from above [10]. With such simultaneous satellite and ground observations, scientists can compare the two perspectives to determine if satellites are missing any important details [11]. Other initiatives include land cover mapping campaigns or the application of satellite image processing for epidemiological

doi: 10.2788/854791

research, where various analyses have shown that fully automatically generated land cover maps or land characterizations exclusively based on satellite data may have huge offsets in some areas [12, 13].

Calibration processes based on citizen science initiatives make use of the high motivation of volunteers to contribute to the scientific process. Though results and quality may vary, recent publications have shown promising results that highlight the value of citizen science campaigns for anchoring satellite measurements [14].

10. REFERENCES

[1] Huadong Guo, Lizhe Wang, Fang Chen, and Dong Liang, "Scientific big data and Digital Earth," *Chinese Science Bulletin*, vol. 59, no. 35, pp. 5066–5073, 2014.

[2] Stephanie E Hampton, Carly A Strasser, Joshua J Tewksbury, Wendy K Gram, Amber E Budden, Archer L Batcheller, Clifford S Duke, and John H Porter, "Big data and the future of ecology," *Frontiers in Ecology and the Environment*, vol. 11, pp. 156–162, 2013.

[3] Rhiannon Crain, Caren Cooper, and Janis L Dickinson, "Citizen Science: A Tool for Integrating Studies of Human and Natural Systems," *Annual Review of Environment and Resources*, vol. 39, no. 1, pp. 641–665, oct 2014.

[4] Jin Xing and Renee E Sieber, "A land use/land cover change geospatial cyberinfrastructure to integrate big data and temporal topology," *International Journal of Geographical Information Science*, pp. 1–21, nov 2015.

[5] Richard J Ladle, Ana C M Malhado, Ricardo A Correia, Jhonatan Guedes dos Santos, and Ana M C Santos, "Research trends in biogeography," *Journal of Biogeography*, pp. n/a–n/a, aug 2015.

[6] OGC, "OGC Interoperability Program Initiatives," http://www.opengeospatial.org/ogc/policies/ippp, 2015.

[7] OGC, "OGC Testbed-12," http://www.opengeospatial.org/testbed12, 2016.

[8] ISO, "ISO 19115-1:2014," Tech. Rep., ISO/TC 211, 2014.

[9] W3C, "Data Catalog Vocabulary (DCAT)," http://www.w3.org/TR/vocab-dcat/, 2014.

[10] Hideo Shiogama and Tomoo Ogura, "Climate science: Clouds of uncertainty," *Nature*, vol. 505, no. 7481, pp. 34–35, jan 2014.

[11] NASA, "Calling All Cloud-Loving Citizen Scientists," http://earthobservatory.nasa.gov/IOTD/view.php?id=84662, 2016.

[12] Vincent Herbreteau, Gérard Salem, Marc Souris, Jean-Pierre Hugot, and Jean-Paul Gonzalez, "Thirty years of use and improvement of remote sensing, applied to epidemiology: from early promises to lasting frustration.," *Health & place*, vol. 13, no. 2, pp. 400–3, jun 2007.

[13] Steve Kelling, Daniel Fink, Frank A. La Sorte, Alison Johnston, Nicholas E. Bruns, and Wesley M. Hochachka, "Taking a Big Data approach to data quality in a citizen science project," *Ambio*, vol. 44, no. S4, pp. 601–611, 2015.

[14] C J Hodges, D G Leibovici, S F Ties, and Leaf Watch, "Building confidence in crowdsourced data for biological monitoring and policy making as part of FP7 Project COBWEB," , no. May, pp. 7–8, 2014.

Structural classifiers for contextual semantic labeling of aerial images

Hicham Randrianarivo*, Bertrand Le Saux*, Nicolas Audebert*, Michel Crucianu† and Marin Ferecatu†
*Onera The French Aerospace Lab
F-91761 Palaiseau, France
† Conservatoire National des Arts et Métiers, Laboratoire CEDRIC
292 Rue St Martin FR-75141 Paris Cedex 03

***Abstract*—In this work we propose a novel approach to classify aerial images with structured predictors. They get high performances by encoding the contextual information that exists around a region of interest. More precisely, we learn a graph model that takes into account the structure existing between adjacent regions (or superpixels) belonging to various categories with a Structural Support Vector Machine. The whole image classification task is broken down in multiple local subtasks so that we can deal with large amounts of regions (more than 100k).**

I. Introduction and previous works

A. Introduction

This work focuses on semantic labeling of aerial images. Nowadays more and more images captured by aerial or satellite systems become available. For example the Sentinel satellites collect more than 2To of images a day. The spatial resolution of these images vary between 300 meters to few centimeters. Manually analyzing this amount of data would take a lot of time and manpower: this is why we want to automatically analyze these images. Our goal is to automatically classify all the pixels of an image into several predefined categories. The task of segmenting and labeling Very High Resolution (VHR) imagery is one of the most active research axes in the remote sensing community [1], [2], [3], [4], [5], [6] as it is essential for a wide range of activities like deforestation analysis or urban modeling. The standard approach consists in extracting a feature descriptor on several regions of the images then learning a model over the features in order to predict the category of the region of interest. Our approach is based on the idea that we can enrich the description of a region of interest with more information from the image. By incorporating in the model information about the neighborhood of the regions of interest we can improve the recognition abilitiy of a model and then improve its predictions. In this paper we show that even novel and state-of-the-art approaches like deep neural networks (AlexNet [7]) can be improved if we extract useful contextual information of the neighborhood of the region of interest. To this aim we learn the pattern which exists in the local neighborhood structure using a graph representation and a Structural Support Vector Machine (SSVM) framework.

B. Overview

Recently a new state of the art in semantic labeling has been established with deep neural networks. Such networks (for example AlexNet) are known to be very effective when it comes to extract a powerful representation of the content of an image, even when trained on general purpose datasets [5]. We use two kinds of setup based on AlexNet as our baseline algorithms. The first one compute features on small patches centered on the superpixels while the second one uses quite larger patches which include more context. In both cases feature descriptors are compute on the patches using AlexNet and used to train a multiclass Linear Support Vector Machine (SVM).

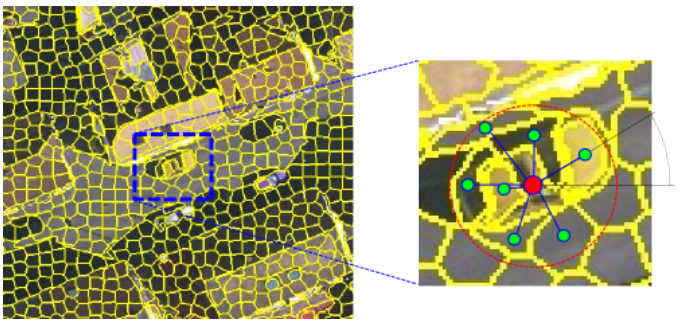

Fig. 1: Overview our method to model local relations between supepixels. For a superpixel of interest we extract a local graph of relations that will be used to make a prediction. The node of interest is the node in red and the neighbor nodes are in green. We model interactions between a node of interest and the neighbor nodes to improve the prediction.

Our approach builds on the small patch baseline but moreover encodes the context using higher order information. In an image, some categories of objects are more likely to appear next to other categories of objects (cars are more likely to appear on a road than on a tree). Our method proposes to model the structure that exists between differents regions in the image and to learn a contextual model that take into account the local interactions between the regions to predict a category. We model the local interactions between regions using a graph structure. The nodes of the graph are the feature vectors of the region of interest. Using a graph structure allows to add extra information on the edges which will be useful to model the interactions between the nodes. To this effect we define a contextual feature which captures the relative positions of the

doi: 10.2788/854791

regions that lie below a given radius. Locally for each region we regularize the local graph computed from single classifier using SSVM which is an extension of the popular SVM. The SSVM framework then predicts new classes for all regions of the graph.

II. Contextual modeling

A. Baseline classifier

Our baseline classifiers predict a label for each superpixel. The visual information associated with a superpixel is taken from the patch centered on it. It will be processed by the deep network. We use the following pipeline for semantic segmentation:

1) Divide the image into small regions of interest (superpixels) using the Simple Linear Iterative Clustering (SLIC) algorithm [8]
2) For a region of interested extract a patch of size $(N \times N), N \in \{32, 64\}$ centered on the superpixel.
3) Resize the patch to 228×228 and process them through AlexNet.

We use theses features in order to train a linear multiclass SVM. As a groundtruth we have label maps where all the pixels are assigned to a category. One issue with using superpixels as training samples is that superpixels can incorporate pixels from two differents categories. To deal with this problem we perform a majority vote according to the groundtruth to label a superpixel. The resulting semantic map is obtained thanks to the prediction of the SVM.

B. Structural context

Usually the region of interest (for instance a patch) we want to classify are very small parts of a semantic object. Looking only at patch level shadows or unusual appearance may lead to misclassifications. Introducing visual information from the neighboorhood allows to reduce uncertainty about the patch.

1) Graph model: The second contextual information we use in this work is the structure of the local interactions of regions of the image. The underlying assumption is that superpixels surrounding a region of interest have a structure that can be learnt. One natural datastructure to model structural relationship between several elements is a graph model. Using a graph to model interactions between superpixels allows us to incorporate extra contextual information in the graph. We construct the graphs that capture interactions between regions of an image in the following manner:

1) Divide the image into small regions using the SLIC algorithm [8]
2) For a region of interested extract a patch of size (32×32) centered on the superpixel.
3) Create the nodes of the graph using the features
4) Create edges between nodes where the distance between the centroids is less than a radius r
5) For each edges in the graph compute a feature of context composed of the distance ρ between the centroid and the angle θ which is the relative orientation between the superpixel centroids.

The result for each image is a set of nodes $\mathcal{V}$ and a set of undirected edges $\mathcal{E}$. A graph $G = (\mathcal{V}, \mathcal{E})$ is then associated with each image of the training set.

2) Structural model for context: Our model is composed by unary features (describing the nodes) and pairwise features (contextual feature between two nodes) which jointly describe interactions between input and output variables. For training, we use a set of N images associated with their label maps. From the images, we extract graph models as explained in section II-B1. We then extract a set of local relation graphs that we use as training samples $X = \{x^n\}_{n=1}^N$ with the corresponding groundtruth annotations $Y = \{y^n\}_{n=1}^N$ generated from the label maps. The specificity of the SSVM framework is that the output labels Y are structured, which means they are graph of classes and not single class values. A target y^n is a set a labels y_i where each label corresponds to a node x_i. The labeling of a region of interest is found by minimizing the following energy function:

$$E(x, y, ; w) = \sum_{i \in \mathcal{V}} \phi_i(y_i; w^\phi) + \sum_{(i,j) \in \mathcal{E}} \varphi_{i,j}(y_i, y_j; w^\varphi) \quad (1)$$

$$= \langle w, \psi(x, y) \rangle \quad (2)$$

With $\phi_i(y_i)$ as the unary term and $\varphi_{(i,j)}(y_i, y_j)$ as the pairwise term. The parameters to learn are w^ϕ and w^φ.

3) Max-margin structured learning: The max-margin structured learning framework optimizes discriminatively the weights of the energy function described in Equation (1). Learning the weight parameters of Equation (1) does not scale well because the computational cost is quadratic. The authors of [9] propose an efficient method to solve this issue using Block-Coordinate algorithm which allows to break down the optimization problem into simpler linear subproblems. The SSVM framework finds model weights that maximize the energy of any labeling y with respect to the one of the groundtruth y^n by the largest margin $\Delta(y, y^n)$

$$w* = \arg\min \frac{1}{2}||w||^2 + \frac{C}{N} \sum_{n=1}^{n} l(x^n, y^n, w) \quad (3)$$

with

$$l(x^n, y^n, w) = \max_{y \in \mathcal{Y}} \Delta(y^n, y) - g(x^n, y^n, w) + g(x^n, y, w) \quad (4)$$

$$g(x, y, w) = \langle w, \varphi(x, y) \rangle \quad (5)$$

where C is a penalized hyperparameter and $\Delta(y, y^n)$ is a loss function that measures the error of predicting y knowing the real configuration is y^n. Several loss functions have been defined for the max-margin structured problem. In this work we use the common Hamming loss which aims at penalizing wrong labeled nodes equally and is defined by $\Delta(y, y^n) = \frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} [y_i \neq y_i^n]$

doi: 10.2788/854791

4) Predicting a label: For each superpixel of an unknown image we consider the local graph of neighbors inside a given radius. The SSVM framework predicts the labels of all the nodes in the graph of local interactions. One approach could be to keep only the label of the region of interest and using it to predict the label of the superpixels. In our approach we also predict the labels of surrounding superpixels, so for the whole image we get several predicted labels for each superpixel. We exploit this by setting up a voting procedure where for each superpixel we accumulate the votes of all the neighbors.

III. Experiments

A. Setup

The methods are tested on the ISPRS 2D Semantic Labeling Dataset [10]. We use part of the Vaihingen data, consisting of 16 *IR-R-G* orthoimages with pixel-level ground truth. We asset the quality of the classification with the ground truth.

We split this dataset as follows : tiles 1, 5, 7, 11, 17, 23,26, 28, 34 and 37 form the training set, while tiles 13, 21 and 30 form the validation set and tiles 3, 15 and 32 form the testing set. Note that the "clutter" class is not represented in the testing set. This is justified by the fact that the ISPRS evaluation procedure does not take this class into account. We evaluate the performances of the various algorithms using f_1-score for each category in the dataset.

B. Results

Table I shows a quantitative evaluation of the algorithms for semantic labeling. The structural context method gets the best overall classification rate and outperforms the baseline on 3 categories. This method is efficient to model the interactions of the superpixels of large areas like impervious surfaces or buildings. Objects with many superpixels are more likely to vote for the right label than objects with few supepixels because the error is divided between the neighbors. The poor performances of the model on cars is a consequence of the voting method. The superpixels of roads are more likely to consider a neighbor as a road than a car. The error is then propagated by the neighbors which lead the model to make a great number of bad predictions.

Figures 2 and 4 show the classification maps produced by the algorithms on tile 3 and 32. Figure 3 shows a zoom on a particular area of tile 3: the baseline model often confuses solar panels with cars while our structural model is able to correct this type of errors. Table II show how to read the classification maps: the colors in the table correspond to the colors in the classification maps. These results show that using structural context produces semantic maps with less noise caused by misclassifications of superpixels: it allows smart regularization of object borders.

IV. Conclusion

In the paper we described a context model for semantic labeling in aerial images. We use local graphs of interactions between superpixels of an image to model contextual relations. We use an efficient SSVM framework to learn a model of context with more than 100k training samples. We have shown it increases the classification performances on the challenging ISPRS dataset for urban semantic labeling.

TABLE I: f_1-score for each category in the dataset. The last column is the multiclass accuracy score.

Model	Imperv.	Build.	Veget.	Tree	Cars	Overall Acc.
	f_1-score					
Baseline 32	81.26	81.58	**62.71**	77.88	40.10	76.33
Baseline 64	81.13	82.36	62.46	76.13	**41.03**	75.98
Structural Context	**82.00**	**82.40**	58.18	**78.38**	32.46	**78.36**

V. Acknowledgment

The Vaihingen dataset was provided by the German Society for Photogrammetry, Remote Sensing and Geoinformation (DGPF) (http://www2.isprs.org/commissions/comm3/wg4/semantic-labeling.html).

References

[1] D. J. Marceau, P. J. Howarth, J.-M. M. Dubois, and D. J. Gratton, "Evaluation Of The Grey-level Co-occurrence Matrix Method For Land-cover Classification Using Spot Imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 28, no. 4, pp. 513–519, 1990.

[2] S. Xu, T. Fang, D. Li, and S. Wang, "Object classification of aerial images with bag-of-visual words," *IEEE Geoscience and Remote Sensing Letters*, vol. 7, no. 2, pp. 366–370, 2010.

[3] A. A. Popescu, I. Gavat, and M. Datcu, "Contextual Descriptors for Scene Classes in Very High Resolution SAR Images," *IEEE Geoscience and Remote Sensing Letters*, vol. 9, pp. 80–84, jan 2012.

[4] G. Moser, S. B. Serpico, and J. A. Benediktsson, "Land-Cover Mapping by Markov Modeling of Spatial-Contextual Information in Very-High-Resolution Remote Sensing Images," *Proceedings of the IEEE*, vol. 101, pp. 631–651, mar 2013.

[5] A. Lagrange, B. Le Saux, A. Beaupere, A. Boulch, A. Chan-Hon-Tong, S. Herbin, H. Randrianarivo, and M. Ferecatu, "Benchmarking classification of earth-observation data: from learning explicit features to convolutional networks," in *IEEE International Geoscience and Remote Sensing Symposium (Invited talk in the special session on Data Fusion)*, 2015.

[6] H. Randrianarivo, B. Le Saux, and M. Ferecatu, "Urban structure detection with deformable part-based models," in *International Geoscience and Remote Sensing Symposium*, 2013.

[7] A. Krizhevsky, I. Sulskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *Advances in Neural Information and Processing Systems (NIPS)*, pp. 1–9, 2012.

[8] R. Achanta, A. Shaji, and K. Smith, "SLIC Superpixels Compared to State-of-the-Art Superpixel Methods," *Pattern Analysis and . . .*, vol. 34, no. 11, pp. 2274–2281, 2012.

[9] S. Lacoste-julien, M. Jaggi, M. Schmidt, and P. Pletscher, "Block-Coordinate Frank-Wolfe Optimization for Structural SVMs," *International conference on Machine learning*, vol. 28, 2013.

[10] M. Gerke, J. Jung, C. Baillard, S. Benitez, G. Sohn, F. Rottensteiner, and U. Breitkopf, "the Isprs Benchmark on Urban Object Classification and 3D Building Reconstruction," *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. I-3, no. September, pp. 293–298, 2012.

TABLE II: Ground truth classes for semantic labeling for the ISPRS dataset described in section III-A.

Impervious surface	Building	Low vegetation	Tree	Car	Clutter

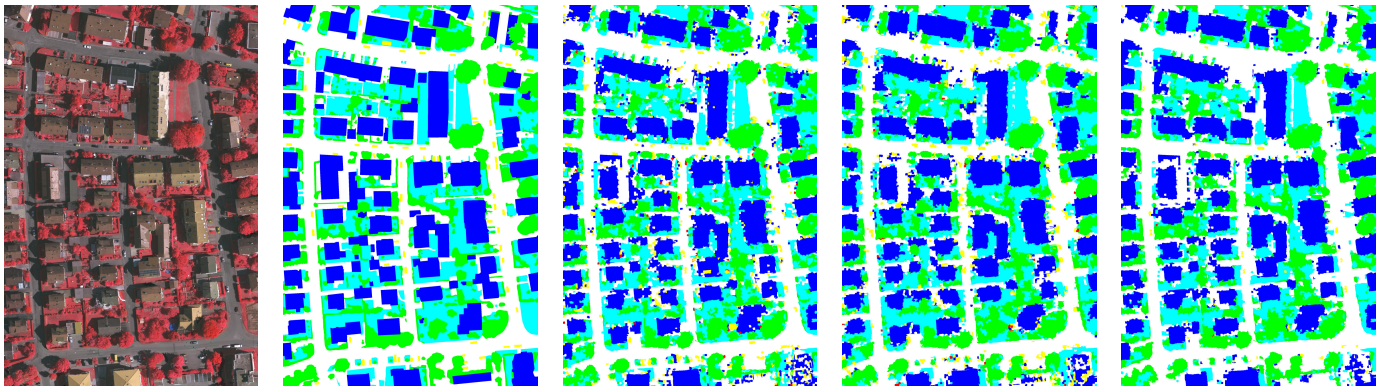

(a) Original image (b) Groundtruth (c) Baseline 32×32 (d) Baseline 64×64 (e) Structural context

Fig. 2: Semantic maps predicted for the tile 3

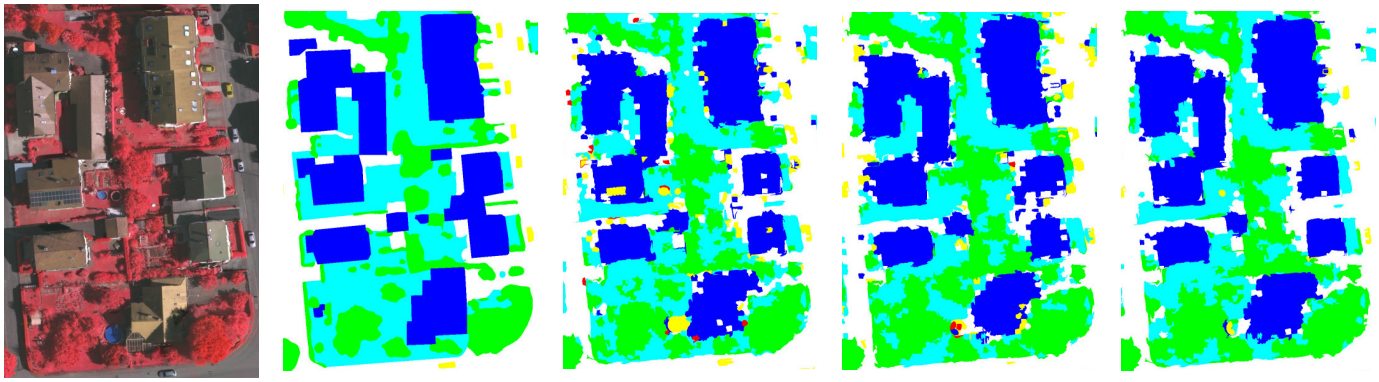

(a) Groundtruth (b) Groundtruth (c) Baseline 32×32 (d) Baseline 64×64 (e) Structural context

Fig. 3: Zoom on the tile 3. We observe that our method produces semantic maps with less noise than the baseline.

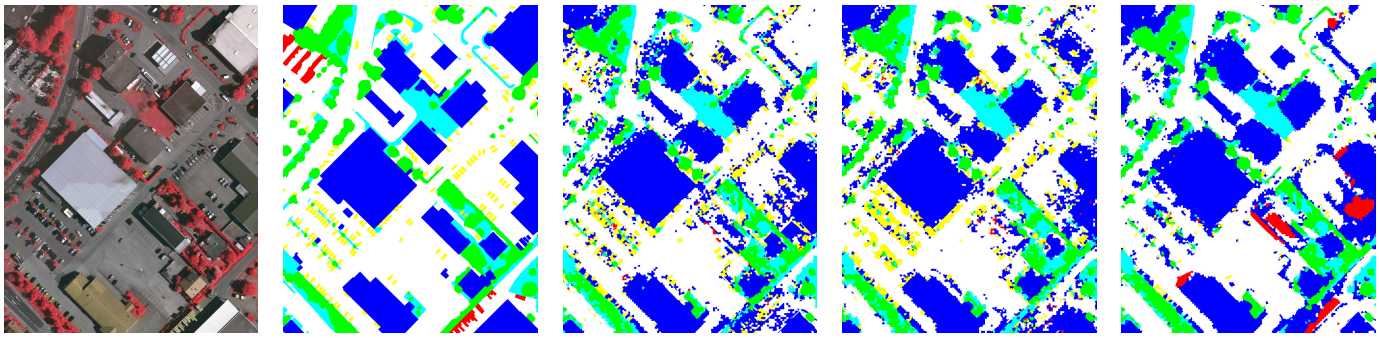

(a) Groundtruth (b) Groundtruth (c) Baseline 32×32 (d) Baseline 64×64 (e) Structural context

Fig. 4: Semantic maps predicted for the tile 32

doi: 10.2788/854791

DEEP LEARNING APPROACH FOR REMOTE SENSING IMAGE ANALYSIS

Amina Ben Hamida **, Alexandre Benoit*, Patrick Lambert*, Chokri Ben Amar***

*LISTIC, Université Savoie Mont Blanc
74940 Annecy le Vieux, France
(amina.ben-hamida,alexandre.benoit, patrick.lambert)@usmb.fr
** REGIM, Ecole Nationale d'Ingénieurs de Sfax
Route de Soukra, B. P. W, 3038 Sfax, Tunisie
chokri.benamar@ieee.org

ABSTRACT

The paper explores how multimedia approaches used in image understanding tasks could be adapted for remote sensing image analysis. In a first step, we show on 3 channels color images through the UC Merced Land Use Dataset how Deep Learning approach provides a significant performance increase compared to Bag of Visual Words approach. In a second step, we propose an extension of deep learning scheme to deal with hyperspectral data. The proposed scheme is based on a 3D architecture which jointly processes spectral and spatial information.

Index Terms— Deep Learning, Remote sensing image, classification

1. INTRODUCTION

The world is speculating a data explosion phase, which demands efficient solutions for its storing, classification and retrieval. Such phenomena has managed to tackle all aspects of today technologies, namely, Remote Sensing (RS) where data flow has witnessed remarkable advances on both levels, spatial and spectral resolutions, but therefore hardening the classification and retrieval tasks. This situation deflected the interest to the classical multimedia approaches in order to inspect them in search for possibly adaptable solutions for RS. In most of the studied cases, the main concept is to resort to images descriptors at different levels reaching for the highly semantic methods, namely, Bag of Visual Words (BoVW). As first inspired from the text recognition tasks, the BoVW approach has managed to introduce a baseline for RS images classification. However, such applications are unable to rise the challenge introduced with the emergence of hyperspectral data. In fact, at that point, such approaches suffer then from some difficulties to efficiently investe the abundant amount of information included at each scene. On the other hand, new models that are able to learn the representations from the data itself have been created. The ILSVRC2012 challenge was first to witness the highly ranked performances of the Deep Learning along with its convolutional neural networks incorporated by Hinton and al [1] in the multimedia field. Since then, most of the attention has been dedicated for the neural network approaches which seem to provide an innovative tool for data mining such as: Deep Learning (DL) [1] and more specifically Convolutional Neural Networks [2].
In early trials, we have resorted to the "UC Merced Land Use Dataset" [3] as a baseline in order to compare the BoVW and DL approaches. As detailed in [4] ; first, a BoVW is tested in its both sparse and dense sampling form reaching for respective classification rates of 86% and 91%. Then we compare with DL following the classical fine tuning procedure applied to the AlexNet network [5] by retaining its initial weights and only retraining the last layer leading for a 94% accuracy. Finally, a comparison study is established between the two approaches proofing that despite its robustness, the BoVW has lost the lead for the sake of DL. However, the challenge of the "UC Merced Land Use Dataset" was resolved with a 100% performance rate as depicted in [6] by fusing the scores obtained by 2 DL models.

In the aim of evaluating DL on harder problems, we focus in this paper on hyperspectral images. We emphasize the need of deep models that handle a small number of parameters using 3D convolutional layers. This paper is organized as follows. We recall in Sec. 2 the most influential works developed for hyperspectral data classification, from which we have derived the concept of our approach as presented in Sec. 3. The interest of the proposed method is illustrated through experimental results in Sec. 4 before concluding and suggesting directions for future work.

2. DEEP LEARNING FOR HYPERSPECTRAL DATA CLASSIFICATION

The reshaping of the previously used neural networks gave birth to the Deep Learning approaches. Such concept has managed to outperform the already existing methods and take over the field of remote sensing image classification. That

goes without saying that other approaches like the one developed in [7] or the BoVW as in [4] can result in very high performances rates. However, due to its impressive abilities to auto adapt itself to different contexts, the Deep Learning has managed so far to draw most of the attention. Multiple networks designs have been proposed. One of the simplest yet performing architectures is AlexNet [5] which is originally dedicated for multimedia RGB images. Nowadays, regarding the abundance of the hyperspectral data, taking into account both the spectral and spatial information has so far been one interesting topic that catalyzes more research. As first introduced in [8], a joint spectral-spatial deep neural network is established. In such case, the spectral information is first processed apart from the spatial component which is lately extracted to be both jointed for feature extraction based on deep architectures like stacked autoencoders (SAE). Neural network classifiers are finally implemented in the final layer. Another approach which relies on auto encoders was detailed in [9], presenting a spatio-spectral framework that merges spectral information from adjacent pixels to add spatial information to the processed pixel. Next, hidden layers are inserted in order to learn the spectral features and a supervised learning is ensured by an output softmax layer. Basically the incorporation of both spatial and spectral information improves the classification performances as mentioned in all methods above. That goes without saying that sometimes, even if the spatial information is disregarded such in [10] , excellent performances are reached up. However, the use of SAE or more generally using large layers of fully connected neurons explode the number of parameters to train and inquire a large amount of training samples. In our case, such condition is very hard to achieve since we suffer from the lack of rich annotated databases. In fact, a variety of approaches have been progressively enriching the hyperspectral images classification domain. But the Deep Learning based ones have opened wide doors into taking huge laps. In search for an accurate and computationally efficient solution, the 3D CNNs have been introduced. As presented in [11], a 3D like approach starts the process with a randomized PCA applied on the spectral dimension of the image. Therefore, the inputs for the first convolutional (Conv) layer (C1) are 3D patches of size $s \times s \times Cr$ where s is the width and height of the spatial patch while Cr is the number of the retrained principle components. A second Conv layer is applied to the output of C1. Finally, a C2 elements vector is produced and fed as input to a Multi-Layer Perceptron (MLP) classifier. However, such approach does only consider the spatial dimension when applying convolutional filters. Another approach is developed in [12], where literally a 3D convolution is deployed by the first layer followed by two 1D Convs and ended with two Fully Connected Layers (FC). This approach is close to the one presented in [10] but adds spatial information. However it still involves large amont of parameters. Our approach follows similar ideas but strongly reduces parameter amount and computational costs. The main concern in most of the cases presented above is then how to deal with the high computation costs since too many parameters need to be trained with few samples. Therefore, we propose to investigate lighter solutions using 3D convolutions for hyperspectral data processing.

3. OUR APPROACH: A LIGHTWEIGHT CNN FOR HYPERSPECTRAL DATA CLASSIFICATION

To achieve the best understanding of the hyperspectral image content, an examination of both its spectral and spatial components is obviously mandatory. The inspiration behind our newly developed approach is then derived from the concept of combining the spatio-spectral aspects along the data categorization process. In this paper, we introduce a new lightweight CNN architecture. Both of the two components are fused at early phases of the process and they are joint in a non separable way thanks to a 3D like version of the usually deployed CNNs. Actually, each pixel is classified accordingly to its $n \times n$ neighborhood and f spectral bands thus representing a $n \times n \times f$ volume. As shown in fig.1, the main purpose is then to establish a CNN that gathers a sequence of both 3D and 1D Conv layers that ends with fully connected ones. Our models mostly relies on convolutional layers. In order to get a better idea of the architecture, lets remind the output size (SizeOut)of one Convolutional (Conv) layer along one dimension with respect to layer parameters (stride, kernel size and padding), where the padding is a zero extension on the signal boundaries that can be handled according to the experiments purposes.

$$SizeOut = \frac{(SizeIn - Kernelsize + 2 \times pad)}{stride}$$

Starting from scratch, a series of *N3D* 3D convolutional layers encompassing a number of *ki* filters is utilized at first place where *ki* may be different from one layer *i* to the other. The Kernels of the filters are of size : $(mi \times mi \times fli)$ where *mi<=n* and *fli<=f*. Each Conv layer uses a stride equal to 1. Each conv layer is followed by a pooling layer to down sample the signal. Such pooling is also based on a convolution layer but with a larger stride as suggested in [13]. Such Conv+pool sequence gradually reduces spatial and spectral dimension and we tune it to get a 1D signal in the end. Afterwards, a number of *N1D* 1D Conv layers is introduced to the network involving *pi* filters each. Finally, the network relies on a set of *NFC* Fully Connected (FC) layers with a softmax. Since the final layer's size is chosen to be equal to the inputs number of classes, such FC guarantees a probabilistic representation for the different classes.

Regarding system complexity and behaviors, we consider both spectral and spatial information in a non-separable way. This enables a lower number of parameters to be trained which can be made even lower than the one used in [10] that is restricted to spectral information processing. First, we have tried to recreate a 3D like version of the architecture

doi: 10.2788/854791

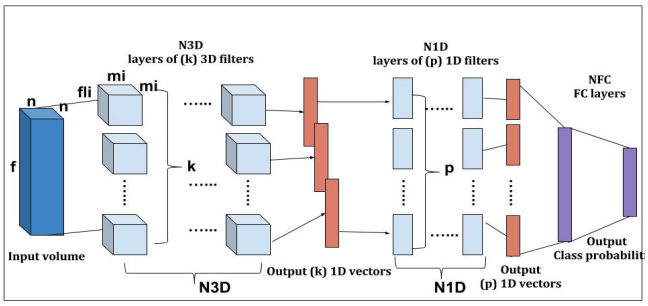

Fig. 1: Proposed 3D CNN architecture

presented in [10]. Then, we investigate deeper models that prove to involve much fewer parameters and provide higher performances.

3.1. From 1D CNNs to their 3D like version

The recreation of the architecture as introduced in [10], has lead to accurate classification rates for the hyperspectral Data. Based on the concept presented above, a multilayer-architecture was recreated. Therefore, a first single 3D based convolutional layer is established and lately followed by one pooling and two FCs. Note here that in our case, we didn't deploy max pooling layers, while instead we did resort to Convs with a stride equal to 2 as suggested in [13].

3.2. Delving into the details of the 3D CNNs

The hardest task in the parametrizing process of Deep Learning networks is then to select the most suitable combination of its specifications. When it comes to the processing on the spatial level, for each pixel to classify, we consider its neighborhood of size $n \times n$. When fed into the DL model, each layer manages its own neighborhood $mi \times mi$ and spectral band width *fli*. Actually, the *"n"* of the input voxels has to be fixed in a way that it is neither too small to incorporate meaningful details, nor too large to the point of degrading local spatial information. Regarding *mi* and *fli*, we follow DL recommendations that favor small sized kernels typically $3 \times 3 \times 3$ allowing then a slow dimension reduction along layers and deeper models.

4. EXPERIMENTS

All the developed codes rely on the Caffe framework which provides a complete Deep Learning toolkit for training and testing models. Besides, Caffe is at that point one of the fastest available implementation for such algorithms.

4.1. Our approach VS the spectral CNNs

- **The Dataset**:

We have chosen to process on the University of Pavia dataset. The employed data was collected by the AVIRIS sensor, capturing an area over Pavia, northern Italy, with a spatial resolution of 1.3 m. The image comprises 610×340 pixels with 103 bands. It mainly contains urban features, soil and vegetation.

- **Configuration of our 3D approach**:

First, a Caffe implementation of the algorithm presented in [10] is tested. This model is based on 1 Conv layer along spectral dimension, 1 max pooling layer and 2 FC layers. The same experimental environment is held during the process. Therefore, 200 pixels from each class were randomly selected from the dataset and dedicated for the training phase while the rest of the pixels was kept for testing. Then, we rely on our 3D version of the previously mentioned architecture. The tests were executed along with a variation of the spatial neighborhood size $n \times n$: 3×3, 5×5 and 7×7. As noticed in Table 1, surprisingly the Caffe application of the architecture presented in [10] didn't rise up to the performance level as declared in the original paper. The only difference is the use of Conv layers for pooling instead of max pooling. However, the introduction of its 3D version, draws paths to enhance the accuracy rate and outperform [10]. Thanks to the adjustment of the spatial neighborhood size to 5×5 pixels, we are able to outperform the [10] presented network with the bonus of training less parameters.

Table 1: 3 layers CNN models: 1D spectral vs 3D spatio-spectral approaches

Approach	Classification rate	Number of parameters
Results of [10]	92.5%	61249
Caffe version of [10]	75.9%	63329
Our 3D version: 3×3	84.0%	16626
Our 3D version: 5×5	93.8%	18269
Our 3D version: 7×7	85.9	20669

4.2. Configuration of our Deep 6 layers approach

Relying on $5 \times 5 \times 103$ input volumes, our lightweight deep approach encompasses 6 Conv layers. The first layer is a *(3,3,3)* kernel sized Conv characterized by a stride equal to *1* and number of neurons equal to *20*. Next pooling is applied using a layer identical to the previous one with the difference of a 1D kernel size *(1,1,3)* and a larger stride equal to *2* in order to reduce the spectral dimension. Outing signal of this 2 Conv layers set is then *(*$3 \times 3 \times 52$*)*. Then, a duplicate of the first and second layers is created with 35 hidden neurons per layer. Output is then 1D *(*$1 \times 1 \times 26$*)* for each neurons. Finally, the 1D spatial dimension is progressively reduced thanks to the use of two Conv layers, 35 neurons each, with respective kernel sizes of *(1,1,3)* and *(1,1,2)* and strides respectively equal to *(1,1,1)* and *(1,1,2)*. Outing signal is then of size (1,1,13) for each of the 35 neurons. The architecture ends with a fully connected layer where the number of neurons is equal to the number of input classes.

4.3. Our 6 layers approach VS the state of the art

One of the best performing approaches dealing with the University of Pavia dataset is the graph based one presented in [7]. In such method, spectral features are computed from spatially identified regions and built from hierarchical image rep-

doi: 10.2788/854791

resentation with prior knowledge. In order to compare our method in similar conditions, only 5% of the total number of pixels (2092 pixels) is randomly selected for the training phase while the rest of the pixels is dedicated for the test. As seen in Table 2, although under the same conditions, our method didn't reach up for the performances level of the approach presented in [7], it still has the benefit of auto adapting itself. And as proofed later, at specific circumstances, it can break the records of [7]. Actually, as basically recommended in most of the DL applications, the main key for a highly performing network is to have enough amount of data for the training phase. As shown in Table 3, the choice of the number of training samples plays a major role in the increase of the classification rates. When using 50% of the training data, we obtain similar performances as the one obtained at [11] trained on 80% of the data and much fewer parameters.

Table 2: Classification rates for our approach VS the [7]'s

Approach	Classification rate	Number of parameters
[7] approach	98.1%	$\times$
OurNetwork: 1D	86.5%	6359
Our Deep Network: 3D, 3×3	95.6%	4419
Our Deep Network: 3D,5×5	93.9%	6074

Table 3: Classification rates according to the amount of training samples

% of training samples	4,4%	5%	50%	95%
Our 6 layers Network 3×3	92.3%	95.6%	98.1%	99,2%
Our 6 layers Network 5×5	93.8%	93.9%	99.6%	99.9%

5. CONCLUSION

The manipulation of hyperspectral data as 3D volumes, ensures a more accurate classification providing that a sufficient amount of training samples is available. The advantage of our developed deep architecture is then to provide a less costing solution while ensuring an accurate classification of the hyperspectral data with few training samples. The Deep Learning can then gain a huge benefit from abundantly annotated data and achieve remarkable results when it comes to remote sensing images classification. One research stimulating path at this stage is then to invest more efforts in the studies of data with higher numbers of spectral bands.

6. REFERENCES

[1] A. Krizhevsky, I. Sutskever, and G. Hinton, "Imagenet classification with deep convolutional neural networks," 2012.

[2] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, "Return of the devil in the details: Delving deep into convolutional networks," August 2014.

[3] UC Merced Land Use Dataset, "University of california, merced," http://vision.ucmerced.edu/datasets/landuse.html/.

[4] A. Ben Hamida, A. Benoit, P. Lambert, and C.Ben Amar, "Could multimedia approaches help remote sensing analysis?," in *Conference Image Information Mining: Earth Observation meets Multimedia*. IIM, 2015.

[5] Y. Jia, E. Shelhamer, J. Donahue, R. Girshick S. Karayev J. Long, S. Guadarrama, T. Darrell, and Trevor, "Caffe: Convolutional architecture for fast feature embedding," in *the ACM International Conference on Multimedia*. New York, NY, USA, 2014.

[6] A. Otavio, B. Penatti, Keiller Nogueira, and A. dos Santos Jefersson, "Do deep features generalize from everyday objects to remote sensing and aerial scenes domains?," in *Journal of Sensors*, 2015.

[7] S. Lefèvre, L. Chapel, and F. Merciol, "Hyperspectral image classification from multiscale description with constrained connectivity and metric learning," in *6th International Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing*. Lausanne, Switzerland, 2014.

[8] X. Ma, J. Geng, and H. Wang, "Hyperspectral image classification via contextual deep learning," in *EURASIP Journal on Image and Video Processing*, 2015.

[9] A. Gatta, C. Camps-Valls, and G. Romero, "Unsupervised deep feature extraction for remote sensing image classification," in *in Geoscience and Remote Sensing*, 2014.

[10] W. Hu, Y. Huang, L. Wei, F. Zhang, and H. Li, "Deep convolutional neural networks for hyperspectral image classification," in *Journal of Sensors*, 2015.

[11] K. Makantasis, K. Karantzalos, D. Anastasios, and D. Nikolaos, "Deep supervised learning for hyperspectral data classification through convolutional neural networks," in *International Geoscience and Remote Sensing Symposium, Milan, Italy, year = 2015*.

[12] S.Viktor, V. Steven, D. Neve, Wesley, V. Sofie, and V. Rik, "Hyperspectral image classification with convolutional neural networks," in *the 23rd ACM International Conference on Multimedia*, 2015.

[13] S. Tobias, D. Alexey, B. Thomas, and A. Martin, "Striving for simplicity: The all convolutional net," *CoRR*, 2014.

PRIORITY SCORES BASED ON NOVELTY DETECTION TO IMPROVE THE EFFICIENCY OF GROUND-OPERATIONS

Jacopo Biancat, Chiara Brighenti, Attilio Brighenti
S.A.T.E. Systems and Advanced Technologies Engineering S.r.l., Venezia, Italy

José-Antonio Martínez-Heras, Alessandro Donati, David Evans
ESA/ESOC, HSO-OSA Human Spaceflight and Operations Department

ABSTRACT

It is well known that in space ground-operations the number of spacecraft telemetry parameters and telecommands keeps increasing while the manpower for operating the ground systems keeps being reduced [1]. This means that it is not possible for Flight Control Engineers (FCEs) to analyze manually each parameter or telecommand to verify their consistency with the overall status of the satellite. Therefore there is the need of automatic methods and/or systems that identify in these large datasets, including today over 15.000 parameters and telecommands for one satellite, which parameters should be the focus of FCEs.. Indeed FCEs should analyse first those showing the higher level of inconsistency with respect to the expectations or to the past behavior.

This paper presents some of the methods and software tools developed by the authors in two studies, under contract with the European Space Agency-European Space Operations Centre (Darmstadt-Germany) regarding the Automatic Behaviour Detection and Interpretation from Low Level Data Sets such as Telemetry and the Automatic spacecraft status characterization by data mining mission history[1], that allow the automatic extraction of the nominal behavior and the reliable detection of novel behaviors, by generating *priority scores*, which are associated to the degree of novelty computed as degree of violation of the detected knowledge, which is extracted in several formats, either for *check* definition or *pattern* analysis.

***Index Terms* —** Knowledge extraction, novelty detection, pattern analysis, checkability

1. INTRODUCTION

The detection of a novel or anomalous behaviour and its possible causes by visual analysis of the acquired raw data is an extremely demanding task for FCEs, given the large amount and variety of data and the possibly hidden cause-effect relations among subsystems parameters. Among the several possible approaches, the most promising ones are those based on data analysis, relying mainly on data mining techniques, because they allow the interpretation of large amounts of heterogeneous data using no or very little a priori knowledge. Usually these approaches elaborate acquired data, automatically characterizing the *nominal* behavior (knowledge extraction) against which, then, any possible *new* phenomenon or unknown relationship among subsystems or processes is identified. Indeed this automatic extraction of the *nominal* behavior from data allows avoiding the step, usually required in a diagnostic system, of knowledge extrapolation and representation by field experts, e.g. for the definition and implementation of a model (*model-based* approach) [2]. It is beyond doubt that this step requires resources and time that are often not available or not compatible with the constraints defined for the system to which diagnostics is applied, especially in industrial fields such as the space sector. In alternative, *threshold-based* approaches, which consists in checking measurable variables for upward or downward transgression of fixed limits, may be used. However, the major drawback of this technique is the need to set wide threshold limits to avoid false alarms, with the consequence that only sudden major faults or long-lasting gradually increasing faults can be detected (eventually with relevant time delays) [3].

Therefore the use of historical process data, collected during the system operation, to generate references of normal conditions without the need of a priori knowledge by field experts, appears to be the most promising solution for analyzing spacecraft telemetry parameters and telecommands.

However the generation of this *nominal* behavior either to apply *pattern* analysis or predefined *checks*, is not trivial because this must be *useful*, i.e. significant and reliable. Significant means that it provides valuable information, allowing at least novelty detection, i.e. the determination of the presence of a novel behavior in a system while reliable means that it does not generate too many false positive, represented by the identification of novel behaviors that are not such.

[1] The view expressed herein can in no way be taken to reflect the official opinion of the European Space Agency.

doi: 10.2788/854791

2. PRIORITY SCORES

Priority scores are synthetic measures proposed by the authors to quantify how different are the behaviors of parameters and telecommands in two different time periods. The greater the difference, the higher the priority score, which suggests that the parameter/telecommand should be further analyzed by the FCEs to investigate if failures occurred to a component or part of the satellite or to confirm whether the observed changes are associated to specific commands sent from the ground stations or generated on board.

Therefore these *priority scores* allow the generation of priority lists, i.e. lists of parameters ordered according to the *priority scores*, indicating which parameters should be further analyzed first.

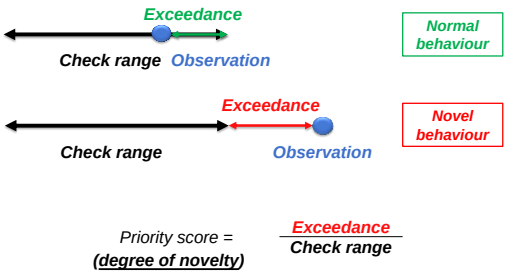

Figure 1 - Example of priority score.

Figure 1 shows an example of priority score computation for a predefined *check*. When the new *observation* falls within the *check range* a negative *exceedance* is computed which is associated to a *normal behavior* of the system (null priority score). When the new *observation* falls out of the *check range* then a positive *exceedance* is computed which is associated to a *novel behaviour* and its *degree of novelty* is quantified as the ratio between this exceedance and the check range.

3. FEATURES

One of the most common approach in data mining is the computation of *features* [4], i.e. characteristics of a set of data or of a signal, which can be measured or calculated over a specific time interval (e.g. the mean, the standard deviation and the minimum values of a signal or parameter in a given time window are features of the signal). Features are often used to reduce the volume of data, while preserving relevant information. In this application context, features can also be used to characterize the *nominal behavior* of a parameter, through the characterization either of the *single features* or of the combination of different features computed over the same time window, representing the so called *feature-based patterns*. Different checks can be defined for both types of approaches, to perform *novelty detection* and *priority scores* computation.

4. SINGLE FEATURE - FETCH

Two alternative approaches were investigated for the definition of *useful* checks (i.e. significant and reliable) of a *single feature*. The first one is based on a density-based measure called *entropy* of a feature. This measures the order of the data in the feature's space. The idea is that a low entropy value corresponds to a set of data with clusters (i.e. well separated data), whereas a high entropy value corresponds to a set of data without clusters [5]. For this reason entropy could be exploited for the selection of checks whose performance may be influenced by the presence of cluster in the data. However the results showed that the entropy does not prove to be a discriminant factor for the reliability of the check.

Indeed the use of *features* for check definition poses the issue of defining a suitable time interval, named Time Window Duration (*TWD*), for their computation. Therefore an innovative method, developed by the authors, called FEaTure CHeckability (***FETCH***), allows identifying the shortest time window beyond which a robust nominal behavior of the feature may be generated so that a *useful* check can be done.

Different kind of checks have been defined to verify that the behaviour of a feature does not change over time from different points of view, including:

1. Domain check: to verify the range or values of the feature.
2. Variability check: to verify the trend of the feature.
3. Distribution check: to verify the distribution of the feature.
4. Frequency check: to verify the frequency content of the feature.
5. Inter period check: to verify the modes of a feature.
6. Few samples check: this check applies to the parameters that do not have enough samples to compute features. It consists in verifying if the parameter values change over time with respect to the reference ones, computed from a nominal dataset.

One of the main outcome of the studies performed by the authors is that features checkability is strongly related to the *TWD* used for the computation of the feature or its characteristics (e.g. distribution, frequency).

Indeed the results showed that the same feature computed for the same parameter may have different behaviors which lead to different checks to be applied.

For example, considering the parameter shown in Figure 2 and the feature *minimum*, it is observed that using a *TWD* equal to 1 hour the feature shows a recognizable periodicity (Figure 3) which allows performing a *frequency check*, to verify that the frequency content remains the same; instead, using a TWD equal to 36 hours (Figure 4), the feature has a constant behavior which allows applying a *domain check* to verify that the values of the feature do not change over time.

doi: 10.2788/854791

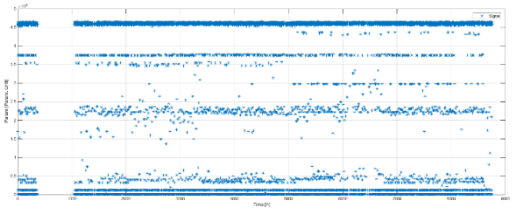

Figure 2 - Example of one satellite telemetry parameter.

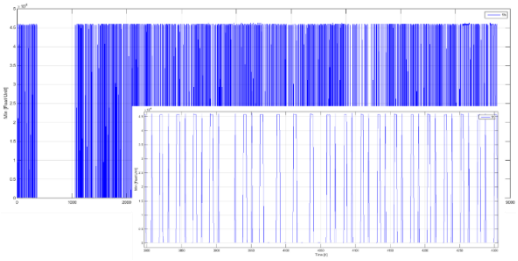

Figure 3 - Feature *minimum* computed with *TWD* = 1h.

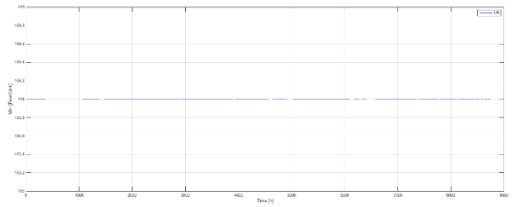

Figure 4 - Feature *minimum* computed with *TWD* = 36h.

5. FEATURE BASED PATTERN - KETTY

Features-based patterns were implemented in a software prototype, named ***KETTY*** (Knowledge ExctracTion from TelemetrY), performing a characterization of the behavior of a parameter synthesizing it into a finite number of states associated to *patterns*.

Several alternative approaches were considered for the analysis of these *patterns* including the detection of the appearance of new *patterns*, the changes in the values of the *patterns*, the analysis of the degree similarity of a new observation with respect to the identified *patterns* and the analysis on the number of occurrences of the patterns. This last approach was selected because it provided the most robust and reliable results, allowing both a static and dynamic characterization of the behaviour of a parameter in a given time period (*reference dataset*). This characterization may be used to detect novelties in the static or dynamic behaviour of the same parameter in another time period (*comparison dataset*), by comparing the occurrences in the *reference dataset* with those in the *comparison dataset*.

In the following some examples are provided using the blue colour to show the behaviour of the parameter in the *reference dataset* and the red colour to show the behaviour of the parameter in the *comparison* dataset which is detected by ***KETTY*** as being characterized by novelties.

The first example is related to a parameter having shown novelties detected by the long-term analysis which aims at the detection of novelties of long duration within a long time interval (e.g. one month). The parameter is categorical, taking the two states ON and OFF. It is clear from Figure 5 the different behaviour of the parameter in the two datasets.

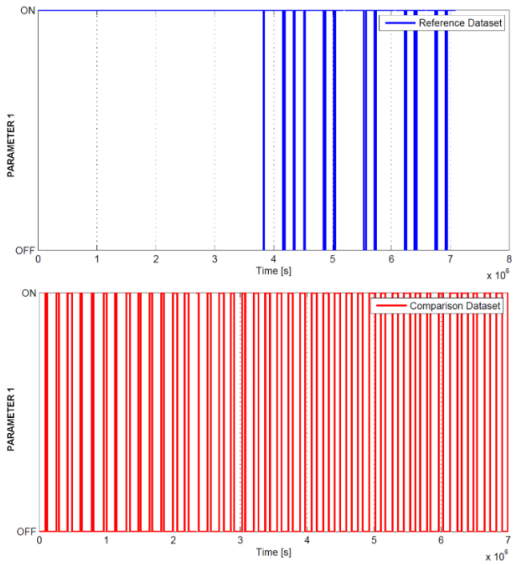

Figure 5 - Long-term novelty analysis - Categorical parameter (top blue plot for the *reference* behaviour; bottom red plot for the behaviour being characterized by novelties).

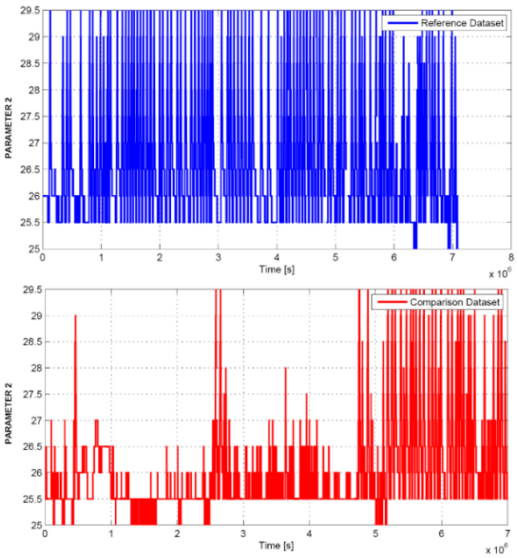

Figure 6 - Short-term novelty analysis - Numerical parameter (top blue plot for the *reference* behaviour; bottom red plot for the behaviour being characterized by novelties in the first part).

doi: 10.2788/854791

The second example is related to a parameter having shown novelties detected by the short-term analysis which aims at the detection of novelties of short duration within a long time interval. The parameter is numerical and also in this case it is clear (Figure 6) that the parameter behaviour is different (with respect to the *reference dataset*) in the first part of the *comparison dataset*, becoming similar in the last part.

For this analysis the method foresees the partition of the time interval of the *comparison dataset* into short time windows, named periods (e.g. 48 hours, see Figure 7 including 40 periods) and the computation, for each window of one priority score associated to the degree of novelty observed in the parameter behaviour in that specific window. Therefore, it is possible to rank the windows and to detect the periods in which the parameter behavior differs most from the reference one.

Table 1 shows an extract of the periods ranking for the parameter shown in Figure 7.

Period N°	Period starting date and time	Period ending date and time	Priority score
9	2009-05-01,00:00:00	2009-05-02,23:59:48	0.329
8	2009-04-29,00:00:00	2009-04-30,23:59:48	0.271
10	2009-05-03,00:00:00	2009-05-04,23:59:48	0.255
7	2009-04-27,00:00:00	2009-04-28,23:59:48	0.252
11	2009-05-05,00:00:00	2009-05-06,23:59:48	0.232
12	2009-05-07,00:00:00	2009-05-08,23:59:48	0.232
26	2009-06-04,00:00:00	2009-06-05,23:59:48	0.206
⋮	⋮	⋮	⋮
36	2009-06-24,00:00:00	2009-06-25,23:59:48	0
37	2009-06-26,00:00:00	2009-06-27,23:59:48	0
38	2009-06-28,00:00:00	2009-06-29,23:59:48	0
40	2009-07-02,00:00:00	2009-07-03,23:59:48	0
29	2009-06-10,00:00:00	2009-06-11,23:59:48	0

Table 1 - Ranking of short-time windows, named *periods*, according to the priority score computed for each period.

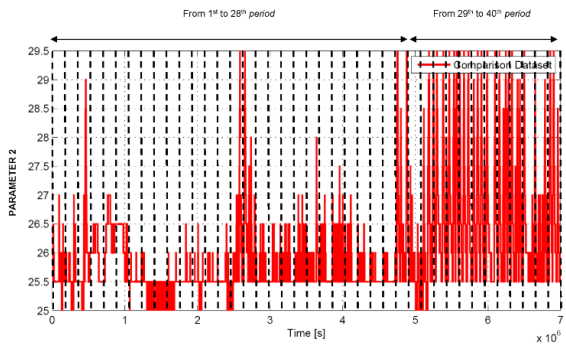

Figure 7 - Short-term novelty analysis - Numerical parameter - Windowing for the *comparison* dataset.

The periods characterized by the highest priority scores are the 8th, 9th 10th and 7th, which are all in the first part of the comparison time interval in which the parameter shows the most different behaviour with respect to the one observed in the *reference dataset*. On the other side the 36th, 37th, 38th, 40th and 29th periods are characterized by a priority score equal to 0, meaning that no novelties are detected in the behaviour of the parameter. Indeed it is clear from Figure 6 that the parameter behaviour in the last part of the *comparison dataset* is similar to the one observed in the *reference dataset*.

6. CONCLUSIONS

The results obtained by the methods included in ***FETCH*** and ***KETTY,*** applied to the telemetry data of two satellites provided by the European Space Agency including over 10.000 parameters each, were validated by visual inspection of the behaviour of the parameters on the top and bottom parts of the priority lists. These were consistent with the expectations proving that the methods are capable of identifying which parameters are characterized by the highest level of novelty and that should be addressed first by FCEs.

It must be highlighted that, although in the examples shown the methods were applied to ex-post analysis, once the diagnostic algorithms are trained they can run continuously, allowing online diagnostics, since their computational constraints are compatible with online operation.

Finally it is worth highlighting that, thanks to the general approach adopted for the definition and implementation of the knowledge extraction techniques, ***FETCH*** and ***KETTY*** are deemed easily applicable to a wide set of different application domains, which have the following characteristics:

1. Low or basic a priori knowledge about the system,
2. Large amount of heterogeneous data.

7. REFERENCES

[1] JC Gil, N. Narula, T. Lopez, "There can be only one: Heterogeneous satellite fleet automated operations with a single tool and language, the MEASAT case", SpaceOps 2014 Conference, 2014.

[2] A. Brighenti "Model-based diagnostics: a real opportunity for efficient vehicles management", SAE On-Board Diagnostics Symposium, Indianapolis Marriott Downtown, September 22-29, 2009.

[3] J.A. Martinez-Heras, A. Donati, M. G. F. Kirsch, F. Schmidt, "New telemetry Monitoring Paradigm with Novelty Detection", SpaceOps 2012 Conference, Stockholm, Sweden, June 11-15, 2012.

[4] P.-N. Tan, M. Steinbach and V. Kumar, *Introduction to Data Mining*, Pearson Education Limited, United Kingdom, 2014.

[5] M. Dash, K. Choi, P. Scheuermann, H. Liu, "Feature selection for clustering - A filter solution", 2nd IEEE Int. Conf. on Data Mining ICDM 2002, Maebashi Japan, pp. 115-122, 2002

BIG DATA STARTS ON-BOARD

R. Vitulli, P. Armbruster, D. Merodio-Codinachs

European Space Agency – ESTEC, Keplerlaan 1 – 2201AZ Noordwijk (NL)

ABSTRACT

The coming generation of multispectral and hyperspectral Earth Observation imagers will scrutinize Earth's surface in a far broader range of bands than the human eye: hundreds of wavelengths per individual pixel, acquired at a much greater spatial as well as spectral resolution than current instruments. Similarly radar instruments are subject to an increase of bandwidths. Additionally, the trend is for such instruments to observe continuously across the terrestrial land and/or sea surface on an 'always on' basis, as with the Copernicus Sentinel satellites. Some science missions face equivalent challenges related to the volume of acquired data. Gaia for instance has to detect, observe and analyze the brightness, position and motion of hundreds of stars per second in order to achieve its target of 3-D mapping the billion stars during its five-year lifetime. Mission planners are increasingly coming up against the thorny problem of how to manage such Big Data from Space – in particular, how to handle them on-board and relay them down to the ground. This paper is presenting future technology trends for on-board data handling and processing, with the focus on data selection, data compression, storage and adaptive processing.

Index Terms— *Big data, Space, On-board, payload, processing, compression, storage, adaptive processing.*

1. INTRODUCTION

The spatial and spectral resolution of space borne instruments is increasing steadily. Assuming a linear trend for the increase of instruments' resolution, the growth of produced data volume follows a quadratic curve, or even a cubic one in case of multi-band instruments. Together with the need of continuous observations, all this results in Big Data volumes to be handled on-board.

We differentiate between optical, panchromatic, and multispectral sensors and synthetic aperture radar (SAR) sensors, reporting each as a separate class. For optical systems advances occur in optical designs, radiometry, signal-to-noise ratios, dynamic range, sensor saturation, cross scene uniformity and spatial resolution among other factors; SAR systems too evolve – better physical antenna elements, improved synthesis of these, more and better echo captures, better pulse generators and the like.

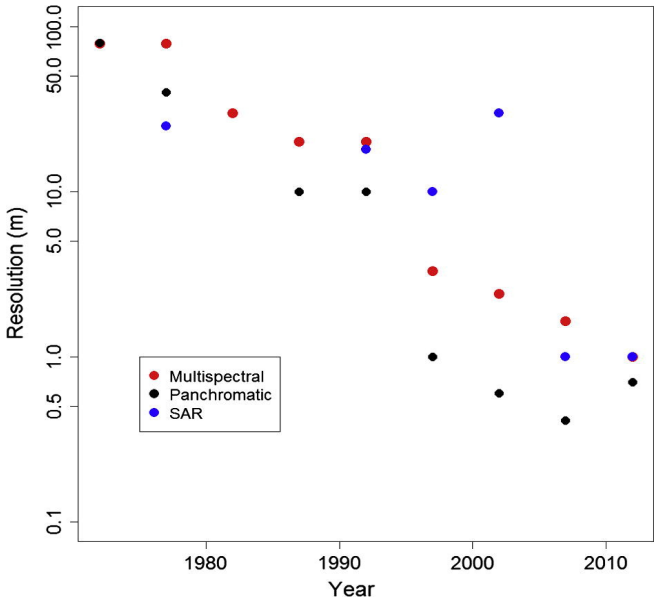

Figure 1-The finest resolution (meters) achieved from any Earth Observation sensor [1]

Fig. 1 shows how the spatial resolution of sensors on land-cover observing satellites has changed over the last four decades, plotting the highest resolution for 5-year periods for the different sensor types. In any five-year epoch panchromatic data are always acquired at a finer resolution than multispectral. However, both panchromatic and multispectral imaging now takes place at much finer spatial resolutions than in previous decades. The first high resolution SAR imagery came from Canada's Radarsat 1 in 1995; this mission provided C band imagery at various resolutions down to 8 m. Very high resolution SAR finally became available in 2007 and 2008 with Italy's first CosmoSkymed and Germany's TerraSAR, both of which provided X band imagery at 1 m resolution [1].

It is evident from the scenario depicted above, that sensors' resolution is tending to become finer and finer from one decade to the other. The volume of acquired data is becoming huge, and how to handle them on-board and relay them down to the ground is challenging.

2. DATA RELAY TO GROUND

On-board acquired data have to be transferred from space to ground for them to become useful. A mission's available

transmission capacity is bound by the combination of orbital constraints, telemetry antenna size, transmitter power and the availability of ground stations. Increasing the number of ground stations assigned per mission might not be an option due to induced cost issues.

Relaying communications via higher-orbiting satellites is one promising means of increasing connectivity, set to be carried out by ESA's European Data Relay System (EDRS). Dubbed the 'SpaceDataHighway', EDRS helps Earth-observing satellites to transmit large quantities of data down to Europe in near-real time. Its two geostationary satellites (EDRS-A, launched in January 2016 and EDRS-C, launch date 2017) use optical links to gather low-orbiting satellites' information at far higher data rates than traditional radio frequency beams. EDRS will provide an inter-satellite link capacity up to 1800 Mbit/s per optical link and a data transmission capacity of at least 50 TBytes/day.

Next interplanetary missions will also be facing limitations of sending the generated data back to ground. The "Interplanetary Data Relay System" concept (IDRS, proposed by H. Barde from ESA/ESTEC) could overcome these limitations. The idea of IDRS was inspired by the ESA BepiColombo mission, where the proximity of the spacecraft to the sun induces frequent occultations and makes some operations especially delicate, in particular the Mercury orbit insertion. A relay satellite located in Lagrangian points L4 or L5 of the Sun-Earth system would considerably reduce the criticality by allowing a permanent TM/TC contact.

In addition there is another, more direct, solution to avoid on-board data bottlenecks: shrink instruments outputs down to a manageable size before downlinking them to Earth. Powerful compression algorithms allow matching available telemetry resources. Their implementations, possibly complemented by pre-processing functions, can offer sufficiently high data throughput for the resulting imagery to be transmitted to the ground on a real-time basis.

If direct transmission is not an option, on-board generated data needs to be stored in a flexible manner until the next visibility opportunity.

As outlined here above, it can be considered that "Big Data" starts on-board. Next paragraphs will elaborate on those topics while establishing a link with related technology trends.

3. ON-BOARD DATA HANDLING OPTIONS

On-board data handling functions can mainly be divided in four categories: data compression, data selection/reduction, data storage and data processing.

Data compression. The algorithms standardized by the CCSDS [3] related working group are widely used and are quite mature. In order to improve efficiency, the algorithm provided is treated as a black box, plugged into the data handling subsystem as just one more step in the processing chain. For the near future there is a need for faster hardware to support real time operations. New technologies, like DSM (see par. 4.2), can allow filling the existing gap towards real-time operations. In case outliers and noise are heavily present in the acquired data, "ad-hoc" solutions need to be provided. Furthermore, developments have been undertaken to embed security functions in the compression ones (encryption embedded in the data stream).

Data selection/reduction. Advanced algorithms for data selection can be used for reducing the data volume. Typical use cases for Earth observation are cloud detection and ROI (Region of Interest) selection. For Science missions, data reduction is done on-board in order to transmit only the scientific useful part of the observed data (e.g. position and velocity of stars in the case of GAIA). A new trend in data reduction is Compressive Sensing (CS) [2]. Topical studies demonstrate that signal acquisition can be performed at sampling frequencies far below the minimal frequency dictated by the ideal sampling theorem. Compressive sensing can be applied to signals that don't convey the entire information amount predicted by the traditional sampling theory, regardless of the maximum frequency contained in their spectrum. Signals with this intriguing characteristic are called sparse. Nonetheless, some mathematical representation of a sparse signal must exist in which the number of non-trivial elements is less than that originated by a non-sparse signal. Such representation is an Integral Transformation (IT) of the signal itself. The sparse mathematical representation admitted by the signal can be made accessible to a sensor, provided that a dedicated subsystem performs the involved IT before its focal plane. When radiometric and spectroscopic signals are considered, an optical subsystem would be the natural choice for optically computing this transformation. Compressive Sampling always includes a multiplexing scheme that exactly implements the mathematical transformation discussed above. Compressive Sampling can bring many latent advantages to hyperspectral satellite imaging, since high spectral and spatial resolution can be attained with fewer detectors, lesser memory capacity, and narrower downlink bandwidth. The main advantage of CS is that compression takes place before the signal sampling, hence avoiding the acquisition of large volumes of data followed by standard signal compression. The possible impact of CS could be remarkable, motivating new investigations and research programs regarding this emerging technology. The principal disadvantage of CS is instead the intensive off-line data processing that leads to the desired source estimation.

On-board data storage. The management of big data volumes can be helped by a proper implementation of on-board data storage server, allowing file-based data management. Those advanced mass memories depart from the simple concept of packet stores to a more flexible handling scheme allowing to handle and transmit files, as done on ground. These files may either contain instrument data formatted according to instrument requirements (e.g. CCSDS packets) or any format depending on the type of file to be stored. To transfer these files to and from ground, a standard file transfer protocol is used – currently CFDP [4]. The Consultative Committee for Space Data Systems (CCSDS) has specified the CCSDS File Delivery Protocol (CFDP) in order to answer the increasing need of a protocol suitable for transmission of files to and from data storage mediums over a Ground-Space communication link. CFDP allows reliable files transfer between spacecraft or between spacecraft and ground (in both directions). The protocol is specifically designed to cope with a large variety of space mission needs and system constraints and to operate across interplanetary distances, despite extremely long data propagation delays and frequent, lengthy interruptions in connectivity. Moving to files for all onboard storage provides for a more modern approach in keeping with commercial ground practices. It also frees up the requirement to store all data in packets and thus reduces overhead.

On-board data processing. Future missions will require more and more demanding processing algorithms. Real time fire monitoring, natural disasters prevention, vessel recognition, sensors data fusion, among others, are examples of information extraction algorithms that could be implemented on-board. The availability of flexible hardware-based processors will make possible the implementation of those algorithms. The following chapter will present, among others, the technology trends related to on-board computing platforms.

4. TECHNOLOGY TRENDS

The evolution outlined here above is fully in line with the technological trends for near future (end of decade) and far future (next decade) as expanded in this section.

4.1. Flexible hardware-based processors

The availability of faster, multi-core microprocessors, such as the Next Generation Microprocessor for Space developed by ESA (GR740 from Cobham Gaisler) will provide enough computing power for most of the on-board required applications. Nevertheless, more demanding processing algorithms may require HW-based processors. Field Programmable Gate Arrays (FPGAs) are Integrated Circuits (ICs) that are completely customized by the designer after manufacturing. This special feature gives a phenomenal degree of flexibility compared to ICs that limit their configuration to register settings.

The anti-fuse based FPGAs (Actel/Microsemi RTSX and RTAX FPGAs) are being extensibility used in space, especially due to their radiation resilience. They brought the FPGA flexibility for electronic equipment; even though they can be programmed only once (i.e. they are one-time-programmable). Other FPGA technologies, that are increasingly being used in space, base their configuration on other technologies that increase even further this flexibility to a limited number of re-programmable cycles in the case of Flash-based FPGAs (ProASIC3 and RTG4 [5]), or even to unlimited re-programmability for SRAM-based FPGAs (Xilinx, Altera, BRAVE [6]).

The re-programmability is the enabler to have the same hardware performing different processing functions; thus to implement adaptive systems for space.

On top of the re-programmability, hardware-based processors require high performance and high capacity to implement complex processing functions. Due to the gap in performance and capacity offered by the FPGAs for Space (i.e. radiation tolerant) versus the commercial FPGAs, the later started to be used in space units where radiation requirements could be relaxed and achieved thanks to radiation mitigation techniques applied by the FPGA designers. This trend started with the Xilinx Virtex QPRO FPGA and has continued to grow. However, the latest availability of high-capacity, high-performance reprogrammable FPGAs like Virtex-5QV and RTG4, and the future European BRAVE FPGA, is paving the way to provide high radiation tolerance to these flexible hardware-based processors.

4.2. High end ASICs based on deep to very deep Sub-Micron technology nodes

The Deep Sub-Micron (DSM) ASIC rad-hard manufacturing process provides higher density and lower power demand. This technology can enable the production of a new generation of more powerful integrated circuits qualified for space, which can face the new challenges that the Big Data era will pose. In the frame of ESA and CNES supported activities, ST Microelectronics has developed the C65SPACE, a radiation hardened standard cell library for space, based on the 65 nm bulk process. The C65SPACE platform is supplied by ST directly to customers, or via intermediates such as Atmel (F) or ISD (GR). First products have been manufactured in 2015. Beyond the 65 nm, ST has a 28 nm FDSOI (Fully Depleted Silicon On Insulator) process, for which promising first radiation tests results have been obtained.

4.3. Mass Memories with file system support

Mass memories in space systems are evolving from simple stream tape-like recorders to complex intelligent (sub)-systems capable of autonomous operations. This evolution is driven by more complex functional requirements coming from multi-payload missions and by the availability of ever-higher density memory components. In commercial markets, NAND Flash memories are widespread as data storage in consumer electronics (e.g., USB flash drives) because of their compactness, low power, low cost and high data throughput.

However, High-Reliability electronics is struggling in keeping the pace with the aggressive scaling down of NAND Flash technology. As a consequence, the use of such technology in space applications is not as established as in the consumer market and still needs further research.

Nevertheless, in the past few years European industry was capable to deliver an impressive gain in performance for Mass Memories thanks to the switch from SDRAM to Flash-based memory devices. Examples are SPOT-5 and Sentinel-2 mass memories.

It is well known that the commercial availability of space-qualified memories is not an effective situation for performance and cost reasons or simply due to a lack of manufacturer's interest. Adoption of Commercial NAND Flash devices for large memory arrays is seen as the only sustainable option. Nevertheless, only a few memory chips have demonstrated the capability to survive in the demanding space environment. Screening of those components is of paramount importance. Moreover, long-term reliability combined with radiation effects over missions' lifetime including storage periods spanning on decades, has to be taken into account.

4.4. High Speed Data links and networks

Instruments require on-board high-speed data links for the handling of the data. Current architectures are based on SpaceWire networks, that are widely used in European and non-European missions. SpaceWire has an intrinsic limitation on the maximum data rate that cannot exceed a few hundreds of Mbits/s.

One alternative option for the near future is to use very high-speed links based on different encoding schemes. In this context SpaceFibre, that will be available either with fibre or copper physical interface, represents a natural evolution for instrument data links. SpaceFibre is currently in the process of being standardised under ECSS (European Cooperation for Space Standardization).

One other important feature is the compatibility with SpaceWire at packet level and higher layers. This allows making use of SpaceWire protocols and to easily integrate SpaceWire and SpaceFibre in a common network. Incidentally, DSM technologies, as introduced in chapter 4.2, are enabling such very high-speed interfaces.

Experiments have demonstrated the use of copper cables up to a distance of 5 meters. For longer connections, using an optical physical layer based on electro-optical converters is a very attractive and mass effective solution.

5. CONCLUSIONS

This paper identifies functions and technologies required for the handling of the massive data generated by space-borne instruments. The emerging interest in "Big Data from Space" can be restricted at a first glance to the mission ground segment and Data Centers. Nevertheless, the same significance should be given to the space segment. As a matter of fact, the on-board data handling sub-system is responsible for the collection, the pre-processing and the transmission of the data to the ground. Therefore, an end-to-end view is to be considered when tackling such a challenge. This is the reason why we can state that "Big Data from Space" starts on-board.

6. REFERENCES

[1] Alan S. Belward, Jon O. Skøien, "Who launched what, when and why; trends in global land-cover observation capacity from civilian earth observation satellites," *ISPRS Journal of Photogrammetry and Remote Sensing, Volume 103, May 2015, Pages 115–128*, Elsevier.

[2] Enrico Magli, Mauro Barni, Andrea Abrardo, Alessandro Barducci, Donatella Guzzi, and Ivan Pippi, ``*Technological issues in compressive sensing,*'' in 2012 On-Board Payload Data Compression Workshop (OBPDC). ESA, 2012.

[3] CCSDS website, http://public.ccsds.org/publications/default.aspx

[4] CCSDS FILE DELIVERY PROTOCOL (CFDP), http://public.ccsds.org/publications/archive/727x0b4.pdf

[5] Jih-Jong Wang, Rezzak N., Dsilva D., Jia J., Cai A., Hawley F., McCollum J., Hamdy E., "A Novel 65 nm Radiation Tolerant Flash Configuration Cell Used in RTG4 Field Programmable Gate Array", *IEEE Transactions on Nuclear Science, Volume 62, Issue 6, Pages 3072 – 3079*

[6] Space FPGA Users Workshop, 3rd Edition, SEFUW, https://indico.esa.int/indico/event/130/

PROCESSING THE GAIA DATA IN CNES: THE GREAT ADVENTURE INTO HADOOP WORLD

Laurence Chaoul[1], *Véronique Valette*[1]

[1]CNES, 18, avenue Edouard Belin 31401 TOULOUSE CEDEX 4, France

ABSTRACT

The ESA Gaia satellite launched on December 19th 2013 is scanning the sky from the Lagrange L2 point in order to build the largest, most precise three-dimensional map of our Galaxy by surveying more than a thousand million stars [1]. The Gaia data processing is handled by a European Scientific Consortium, namely DPAC and relies on six Data Processing Centres (DPC) distributed all around Europe. This paper will focus on the architecture of the CNES Data Processing Center (DPCC), based on Hadoop technologies and will present lessons learned with the processing of the first scientific chains.

***Index Terms*—**Processing Centers, Big Data, Hadoop, YARN, MapReduce

1. OVERALL GROUND SEGMENT ORGANIZATION

The Gaia spacecraft downlinks everyday an average of 30 GB of data which are transferred via the Mission Operations Centre (ESOC/Darmstadt) to the Science Operations Centre (ESAC/Madrid), also hosting the DPC that executes the two core daily systems: Initial Data Treatment and First Look. The results of both systems are then daily distributed to the other DPCs that run systems in order to monitor the payload health and to raise science alerts. At the same time, every 6 months, a new Data Reduction Cycle (DRC) begins and each DPC runs its systems on all the data acquired since the beginning of the mission and sends its results back to DPCE which starts integrating the data into a new Main DataBase (MDB) version.

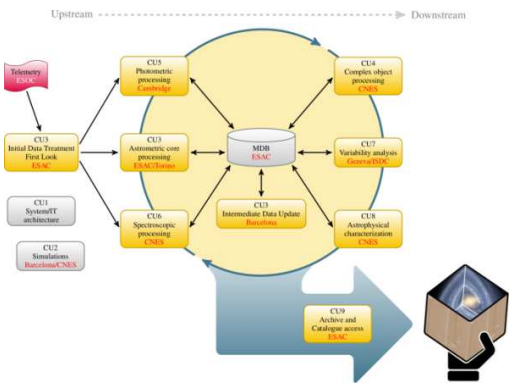

Figure 1 : The DPAC dataflow (DPACE courtesy)

2. GAIA CHALLENGES FOR DATA PROCESSING

Gaia will observe 80 times more than 1 billion stars. The ground data processing has therefore to face several challenges:

- A huge number of elements to handle with dozens of tables containing up to 80 billion rows
- Complex processing with timeliness constraints: daily systems to deal with the DRC ones
- Huge volume to handle: 3PB of data are foreseen at the end of the mission (disregarding intermediate data generated in each DPC.)

3. THE DPCC ARCHITECTURE

DPCC is in charge of Spectroscopic processing (CU6), Objects processing (CU4) (Non-Single Stars, Solar System Objects, Extended Objects) and Astrophysical Parameters determination (CU8).

3.1. Database selection

The DPCC has chosen Hadoop in 2010 as the core of the DPCC framework after a database system study that showed that Hadoop was the best solution to handle more than 10 million objects and with a scalability allowing an incremental purchase of the hardware in order to follow the growing needs in terms of volume and processing power over the 5 years of mission.

3.2. A Hadoop-powered framework

The Hadoop core Map/Reduce paradigm is a quite complex framework to develop and maintain dozens of scientific chains. The Cascading library was then chosen to solve this issue. Cascading is a Java library proposing to chain elementary operations called pipes. At runtime, Cascading translates this assembly of pipes into a chain of Maps and Reduces tasks directly executed by the Hadoop cluster. The overhead of Cascading has been measured at around 5%.

3.3. Hadoop version

The first Hadoop version of DPCC was developed with the Hadoop MR V1 version, the only available version when

Hadoop was chosen. After few months of executions of DPCC scientific chains, the limitations of Hadoop V1 clearly appear and the DPCC has thus decided to switch to the new Hadoop MR V2 version (YARN) early 2015. The main reasons for this choice were:

- MapReduce v1 only manages execution slot without taking into account the different profiles of tasks (in terms of resources utilization). YARN manages resources and can optimize the scheduling taking into account the different tasks profile in queue.
- The Hadoop MR V1 uses static definition of the number of CPU cores dedicated to Maps and Reduces. The YARN version uses the dynamic allocation of tasks on available cores resulting in a higher utilization rate and therefore, an optimization of the hardware investment.
- Finally, YARN also offers an optimised scheduling strategy that is able to better take into account the heterogeneity of the different hardware generation.

3.4. GaiaWeb, the data access Web portal

GaiaWeb is the Web server allowing scientific developers to access the data produced by their scientific chains. It basically brings two kinds of features: a statistical analysis of the data produced on the cluster, using an ElasticSearch engine that indexes the data used in the statistics, and on-demand queries launched on the cluster to retrieve data from the production Hadoop FileSystem.

3.5. Hardware overview

The Gaia DPCC is the first project running in CNES to use Hadoop. Dedicated machines have thus been purchased.
The current operational platform contains 72 calculus nodes (1152 cores), with a storage capacity of 240 TB for HDFS.
The increase of the operational platform will follow DPCC needs in terms of storage and processing power, with incremental purchases planned every year in average.
The foreseen final DPCC operational cluster is 6000 cores, available via a set of 8 racks, each of which consists of 64 servers embedded into 16 enclosures.
A validation platform is also available with 28 calculus nodes (512 cores) and a storage capacity of 110 TB for HDFS. This platform is used for the validation tests of the chains not yet executed in operations.

4. FIRST RESULTS ON THE DPCC PROCESSING CHAINS

Today, the CNES Gaia data processing centre is operational for two daily processing chains (Daily Spectrometric Calibration and the Detection of Solar System Objects).

Moreover, the cyclic chains CU4 NSS (Non Single Stars chain) and CU8 Apsis (for the computation of the astrophysical parameters) have been trained in the validation platform during the Operations rehearsals tests performed by the whole DPAC.
The DPCC Hadoop architecture is now a tried and tested solution.

Every day, on the 1152 current cores :

- about 6 hours are needed for the spectroscopic chain to process around 10 millions of observations
- 2h are needed for the Solar System Object (SSO) detection chain to process about 64 Millions of observations (with 855.000 solar objects).

The parallel execution of the two chains is very well handled by Hadoop.

The results of the validation tests on 28 days of the Eclipse Pole Scanning Law period, on CU4 NSS and CU8 Apsis chains, obtained in 160 cores of the validation platform are described in Fig; 2.

	CU4 NSS			CU8 Apsis		
	Total Duration	# Objects	Duration (ms/objects)	Total Duration	# Objects	Duration (ms/objects)
Insertion	03:08:00	4 946 346 930	0,002	01:06:35	1 366 536 558	0,003
Ingestion	14:00:32	1 193 200 015	0,042	06:40:13	1 366 536 558	0,018
Processing				00:07:37	5 013 503	0,091
	02:00:47	198 055 792	0,037	10:10:21	5 013 503	7,304

Fig. 2 : Results of validation tests on CU4 NSS and CU8 Apsis chains

The above results already underline that the performances results are closely linked with the type of executions performed:

- The insertion phase is only devoted to insert the data into the Hadoop filesystem; This part is linear with the number of objects that are processed;
- The ingestion phase transforms the input data in objects that contains all the information necessary for the processing chain. This phase can include some scientific computations, thus the performances obtained can vary according to the complexity of the algorithm.
- The processing phase can be divided into several steps. The performances of this phase is strongly dependent on the performances of the algorithms

doi: 10.2788/854791

themselves (on CU8 side, the second step is about 100 times longer that the first one)

Although these first results are quite heterogeneous, they demonstrate the capability of such a Hadoop architecture to deal with huge volume of data (more than 1 billion entries) and with promising performances.

5. FIRST LESSONS LEARNT FROM DPCC

5.1. Comparison between Hadoop MRV1 and MRV2 with the daily chains

The gain in terms of performances when changing the Hadoop version is really impressive regarding the capacity to execute in parallel the two daily chains. The Fig. 3 shows a comparison of the execution times with both Hadoop versions and indicating if the two chains were executed in parallel or not.

	MR V1 (432 cores)		MR V2 (1152 cores)	
	Alone	Both chains	Alone	Both Chains
Spectro. Chain	9h20	12h15	5h30	7h
SSO Chain	4h40	9h (and more)	1h30	1h35

Fig.3: duration executions vs Hadoop versions

With Hadoop MR V1, the spectroscopic chain was using all the resources and was preventing the SSO chain from running.

5.2. Performances closely linked to the chains design

The results obtained in DPCC also clearly demonstrate that the scalabilty and the performances of a chain are closely linked to its design.

With the example of Fig.3, we can see that the spectroscopic chain execution time was indeed improved when the DPCC adds new nodes, but not linearly with the number of cores. The SSO chain was more linear as its data policy access is more scalable.

Moreover, a performances test campaign was achieved in the CU6 Daily chain to try to characterize its execution time. The number of input data was changed at each execution to see the impact of the number of processed objects on the execution time.

The Fig. 4 shows the elapsed time spent in each part of the chain according to the number of objects in input.

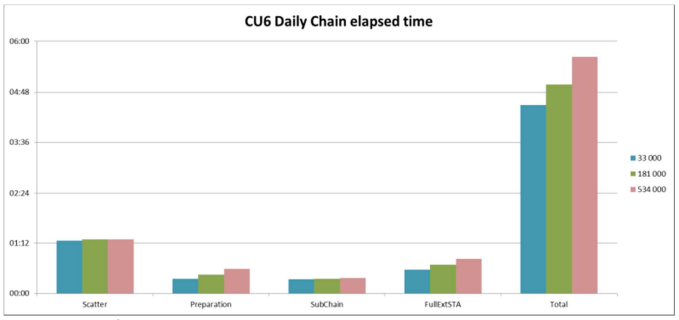

Fig.4 : performances test campaign on CU6 Daily chain

The results clearly show that the execution time of each processing part is really dependent on the design of the chain. The overall processing time increases with the amount of data, but not linearly: some parts are quite constant, while others seem to be linear with the number of processed objects.

These figures underline the importance of the design of the chains to fully benefit from Hadoop capabilities. This design shall be thought at the beginning of a project to be able to have a fully scalable system.

5.3. Management of Hadoop queues

Another important configuration is the tuning of the Hadoop queues.

Indeed, in order to deal with different executions in parallel, Hadoop allows configuring different queues to use a given fraction of the cluster capacity. This insures that a chain will not be fully blocked by another using all the resources.

These queues can be configured in different ways. In DPCC, two solutions have been tried :

- first a hierarchical tuning was used with a queue affected to each CU and with a priority given to the daily chain against the cyclic chain in a CU. This configuration leads to have the cyclic chain fully blocked by the daily one. It was then decided to abandon this option
- now a queue is defined for each chain, in order to be sure that each chain is able to have some resources to run.

Independently of the definition of the queues, some options are also important to have a cluster fully busy:

- usually a single job does not use more resources than its queue's capacity. However, if there is more than one job in the queue and there are idle resources available, then the Hadoop system may allocate the spare resources to jobs in the queue, even if that causes the queue's capacity to be exceeded. This behavior is known as *queue elasticity*. [2] This option prevents from having a

part of the cluster not used if some processing are not activated. The main drawback is that the allocation of resources is dynamic according to the cluster usage when the job starts. The extrapolation of performances is thus very difficult to obtain as it is very difficult to identify exactly how many cores have been really used by the processing.

- the pre-emption option can allow a job of the queue n°2 to delete a job of the queue n°1 that would have overflowed on its queue. This option is quite dangerous in the DPCC case, because jobs of several hours can be killed, whereas they are almost finished. So this option was abandoned.

In conclusion, it is important for each project to find the right tuning to use efficiently its cluster.

5.4. Memory Management with Hadoop

The memory consumption of each job is still a big issue;
Indeed, the switch to Hadoop MRV1 to Hadoop MR V2 has permitted to have a better usage of memory. With Hadoop MR V2, the memory of a CPU can be divided between two jobs, whereas in MR V1, the full memory slot was dedicated to a CPU. If a task required more that the memory slot size, it was affected two CPUs even if in terms of processing power, it only needed one. With Hadoop MR V2, only one CPU is affected, and the other one can be used by a job that needs less memory.
Nevertheless, the memory management in Hadoop MR V2 keeps difficult to tune finely. It is always necessary to know a priori the amount of memory that the jobs will need.

5.5. Performances monitoring

The performances follow-up in Hadoop is another issue. The Hadoop systems deliver a lot of statistics on the execution of each tasks, but it is still very complex to consolidate all this information to make an analysis at the chain level.
The DPCC has just initiated a work on systematic monitoring tools to follow the performances of each chain. The main objectives of these tools are to be sure that the existing hardware will allow to process all the data within a DPAC cycle and to anticipate and size the next purchase of hardware.

5.6. Storage management

In terms of storage capacities, the Hadoop framework is quite flexible. Indeed, it is possible to tune the replication factor in the Hadoop filesystem. The recommended value of this replication factor is 3 to be sure that all the data are well replicated.
However, the DPCC experience shows that some intermediate data can be less secured and that some data storage can be saved. The Hadoop framework allow to tune this replication factor by path in HDFS. This feature offers an interesting flexibility on the data storage management at DPCC level.

5.7. Hardware management

A last important lesson learnt is the management of heterogeneous hardware in the DPCC cluster.
The DPCC operational and the validation platforms have already machines of different generation and with different characteristics (in terms of number of nodes, memory per node, or data storage per node).
This heterogeneity is totally transparent for DPCC.
The deployment tools like Puppet and foreman allow to deploy and add new nodes in the platform with very short interruptions of services.
The management of the different generation of nodes is managed by the Hadoop Resource Manager that send jobs to each kind of nodes according to their profile.

5. PERSPECTIVES

A lot of lessons have been learned during the deployment of these first processing chains. Even if the Hadoop tuning seems better with the new YARN version, it is still necessary to monitor the behaviour of the cluster, when adding new chains and new kinds of execution.
The next step for DPCC is thus to closely monitor the performances of the cluster, to be able to anticipate possible resources overflows during a DPAC processing cycle.

6. CONCLUSION

Initiated since 2010, the challenging objective of DPCC to introduce the Big Data concepts and technology in the astrophysics computational world becomes a reality, with some first promising results.

7. REFERENCES

[1] Lindegren, L.; Babusiaux, C.; Bailer-Jones, C.; Bastian, U.; Brown, A. G. A.; Cropper, M.; Høg, E.; Jordi, C.; Katz, D.; van Leeuwen, F.; Luri, X.; Mignard, F.; de Bruijne, J. H. J.;Prusti, T., The Gaia mission: science, organization and present status, Proceedings of the International Astronomical Union / Volume 3 /Symposium S248 /October 2007
[2] T. White, Hadoop, the definitive guide, 4th edition, O'Reilly, 2015.

BIG DATA FOR SPATIAL OCEANOGRAPHY APPLICATION TO ESA/EU SENTINEL CONSTELLATION MISSION

Paul, F; Archer, O; Piollé, JF; Mouche, A ; Bouhier, N ; Chapron, B – Ifremer, France

ABSTRACT

To face new challenges coming with new satellite missions and massive reprocessing campaigns of decades of legacy archives, a multi-petabyte platform applying big-data designs to perform massive distributed processings has been built. We describe evolution steps, key benefits for end users and technical solutions implemented, starting from hardware infrastructure changes to new practices for users and projects organization when using large data archives. Two examples of concrete applications that currently benefit from this platform are presented. Reprocessing campaigns of ESA legacy archives and Sentinel-1 Cal/Val center illustrate main advantages, ongoing activities and future potential of such thematically driven data exploitation systems.

Keywords : big data infrastructure, massive data processings, data analysis platform, data management

1. INTRODUCTION

During the last decade, the ocean community witnessed the launch of over 30 new ocean-related satellite missions by 13 different contributing space agencies representing around 36 countries (http://eohandbook.com). In the last few years, Ifremer/LOS dedicated a large effort, with the support of ESA, to provide an experimental big-data platform [1][2] aiming to capitalize on recognized expertise, bringing together thematic, observation and validation of EO data (satellite, model, in situ), hardware, and information technology.

Through several reprocessing and science projects, the platform has demonstrated an effectiveness and level of flexibility that was not achievable before.

2. OCEAN REMOTE SENSING (BIG) DATA CENTER INSIGHT

Building a platform based on cutting-edge big data and cloud technologies allows to address ambitions which were not reachable few years ago. Focusing on existing infrastructure issues, this section describes the main challenges we faced and solutions we implemented to improve data center effectiveness.

2.1 Challenge 1 : benefit from a petascale online archive

Data centers hosting data from many satellite missions and sensors for decades deal with data management issues, some of them related to data volume. Historically, tape archives were extensively used to solve scalability issues as well as to ensure long term preservation of data. Yet, performing efficient data processings using such tape librairies is often not possible or complex. Nowadays, new computing concepts allow to easily build scalable and affordable storage systems which solve most of data storage issues. For several reasons including cost and data property, neither high-performance-computing (hpc) architecture nor public web infrastructures (cloud) were viable options. Relying on open-source technologies, we built a multi-petabyte online storage, using cheap hardware and network for affordability, horizontal aggregation for scalability, and data replication for reliability [3]. This infrastructure now mirrors our complete legacy archives as well as huge acquisitions from ongoing projects and missions. Moreover getting rid of all storage concerns brings flexibility in data management and offers new data mining opportunities for end users.

2.2 Challenge 2 : using efficient massive processings for continuous data quality improvements

A common use case for scientists as well as for data managers is to apply processing chains to improve data quality and generate more accurate products (eg. data munging, algorithm improvements). Processing decades of satellite missions archives or even new high resolution sensors data is usually a cumbersome process which requires dedicated projects and teams. Along this process some major issues are merely technical, which induces time wasting and tedious operations (eg. legacy processors incompatible with new hardware or operating systems, endless processing time, data volume exceeding platform capacity). We used virtual environments such as containers [4] to facilitate integration and portability of processing chains. Our distributed processing framework is designed to use automatic provisioning of containers to compute nodes, on demand, allowing convenient method to prepare and run

powerful distributed processings without headache. I/O performances are a major concern for designing big-data architectures. Inspired by new concepts from big-data frameworks like Hadoop [5], we relied on horizontal scalability for storage and processing to achieve good behaviour at heavy processing loads, using cheap hardware (discount hard drives) and widespread network links (gigabit). In addition, new generation schedulers take into consideration data-locality to limit network use (each node serving as compute and storage unit), and they offer smart features using real-time statistics to overcome infrastructure anomalies during operations. These innovative technologies allow to easily perform effective massive distributed processings. Our platform was proven through many reprocessing campaigns and thematic projects, but also for everyday use by remote partners and scientists. The benefit of distributing processings over big data archives almost naturally induces deep changes in users habits, data management and projects organization.

3. LEVERAGE OCEAN DATA ANALYSIS USING BIG DATA PRACTICES

Big-data is not only a matter of infrastructure. It also requires practical human approaches and new generation data analysis frameworks to fully benefit from the possibilities offered by such data exploitation systems. This section describes some challenges recently addressed by our teams.

3.1 Challenge 3 : foster collaborative work and gather around a shared platform

The data warehouse is a keystone of big-data architectures. Data, the raw material, is coupled with efficient analysis frameworks where share practices are fostered in order to promote user interactions and improve extraction of value and knowledge. Gathering complementary skills (scientists, data managers, engineers) or distinct thematics around a shared platform allows mutual enrichment, and each contribution instantly benefit to other users (deployed tools, new data archives, diagnostic analysis, improved algorithms). Shared platforms are proven solutions to smooth collaborative work since everything is natively shared and efficiently accessible. The advantage is even more obvious when all data archives required for a thematic are available alongside with big-data analysis platforms and on-demand cloud computing resources ([6],[7]). However, we need to be really close to end user to clearly understand every barrier they encounter, to propose suitable solutions to overcome it. Assisting users to take ownership of the platform is a major step, especially for them to understand how these new methods allow to multiply the potential to use the data.

3.2 Challenge 4 : toward a new data experience

An exciting ongoing activity is to bring new ways to use the data for researchers. Solutions from the previous challenges allowed to build a reliable multi-purposes platform which shapes a pre-requisite to a new era of data manipulation. For example, transparent use of remote map-reduce [8] frameworks for massive batch processings allows improved data analytics features, combination of multiple sensors and fast data mining and insights extractions from the data. Actually even performing simple statistics or basic machine learning operations over massive archives often raises interesting results. Using new technologies like distributed in-memory frameworks [9] also changes the way we used to interact with the data, for example by performing iterative analyses on big datasets in a few seconds only. Other interesting methods using for example scalable distributed databases [10] allow to index all data with many unstructured meta-data, providing a way to efficiently request heterogeneous data archives (search for colocated data or even extract many years regional statistics interactively for example). Many use cases prove that big-data ecosystem brings effective solutions for researchers work, but combining all these new opportunities to build a data as a service platform era will still require efforts.

4. APPLICATIONS

Several projects contributed to tune current platform practices during last years. This section illustrates some applications that currently benefit from concrete solutions detailed above.

4.1 Reprocessing campaigns of legacy ESA archives

Since 2010, many reprocessing campaigns were performed on the platform in the frame of ESA data consolidation projects (eg. [11]). Reaper (ERS-1/2 Altimetry) [12], X-PReSS (Landsat 5 TM) [13], DSI (Envisat, Landsat-4/5/7, ERS-1/2, …) [14][15][16] are some of these reprocessing activities which contributed to evolve the platform from a proof-of-concept to a scalable processing service. Data storage and processing scalability at petascale was proven by gradually increasing processing complexity level, each time with successful achievements and new effective opportunities. The flexibility gained using these new practices and the dedicated data intensive platform was a clear benefit for these ESA reprocessing campaigns. In addition, user experience and collaborative work with our

remote european partners demonstrated the advantages of this innovative process.

4.2 Sentinel-1 Cal/Val Center

As part of the Expert Support Laboratories teams of the Mission Performance Center for Sentinel-1 (MPC S-1), IFREMER/LOS is in charge of Level 2 Ocean products calibration and validation [17]. Level 2 consists of geo-located geophysical products for wind, wave and currents applications. This activity relies on massive and systematic sanity checks of data, but also on comparisons with reference data for geophysical measurements validation coming from in-situ measurements, model analysis and other satellite missions. All these activities rely on the available big-data platform for storage, indexation, and data analysis of co-located data archives. Main advantages are related to the flexibility offered by these new big-data approaches, which allow easy interactive analysis of colocated data, without worrying about scalability issues due to dataset increase of terabytes of incoming data each week. This platform allowed scientists to change their usual practices to benefit from new features offered by big-data services, which are now used every day in the Ifremer Sentinel-1 Cal/Val Center.

5. REFERENCES

[1] "Nephelae : a platform for data intensive science on oceans", Abstract book, *ESA conference : Big Data From Space*, Italy, June 2013
http://www.congrexprojects.com/docs/default-source/13c10_docs/abstractbook.pdf?sfvrsn=2
[2] Cersat big data platform
http://cersat.ifremer.fr/oceanography-from-space/our-domains-of-research/mass-data-processing-and-mining
[3] MooseFS, a file system for highly reliable petabyte storage, http://www.moosefs.org
[4] Docker, application container engine,
https://www.docker.com
[5] Hadoop: Open source implementation of MapReduce, http://hadoop.apache.org
[6] J.Shutler et al., "FluxEngine: A flexible processing system for calculating atmosphere-ocean carbon dioxide gas fluxes and climatologies", *Journal of Atmospheric and oceanic technology*, December 2015, DOI: 10.1175/JTECH-D-14-00204.1
[7] OceanHeatFlux project website,
http://oceanflux.cersat.fr
[8] Jeffrey Dean, Sanjay Ghemawat, “MapReduce: simplified data processing on large clusters”, Communications of the ACM, v.51 n.1, January 2008, DOI: 10.1145/1327452.1327492
[9] Spark, fast and general engine for large-scale data processing, https://spark.apache.org
[10] Hbase (Hadoop database), a distributed, scalable, big data store, https://hbase.apache.org
[11] "Data Consolidation of ESA's unique ERS mission archives - Keeping alive long-term mission archives : a framework for ocean data", Presentation, *ESA conference : PV 2013*, Italy, November 2013
http://www.congrexprojects.com/docs/default-source/13c17_docs/11-04_10-10_fredericpaul_ers-consolidation.pdf?sfvrsn=2
[12] Femenias et al., "Reprocessing of the ERS-1 and ERS-2 altimetry missions. The REAPER project", *Ocean Surface Topography Science Team Meeting Conference*, Germany, October 2014, DOI: 10.13140/2.1.3756.7685
[13] X-PReSS consortium, "Landsat 5 Reprocessing, a Case Study Into Reprocessing and Data Configuration", Conference presentation , *ESA conference : PV 2013*, Italy, November 2013
[14] DSI Project website, http://xpress.sp.serco.eu
[15] F. Gascon et al., "European Space Agency (ESA) Landsat MSS/TM/ETM+ Archive Bulk-Processing: Processor Improvements and Data Quality", *Conference: SPIE Optics + Photonics 2014*, San Diego, DOI: 10.1117/12.2062223
[16] F. Gascon et al. “Bulk processing of the Landsat MSS/TM/ETM+ archive of the European Space Agency”, *Conference: SPIE Optics + Photonics 2015*, Pragues, DOI: 10.1117/12.2192847
[17] Ifremer MPC S-1 ocean waves remote sensing cal/val, http://oceanwavesremotesensing.ifremer.fr

BIG DATA PROCESSING FRAMEWORKS APPLIED TO SPACE MISSIONS

Guillaume Eynard-Bontemps[1], *Alain Montmory*[1]

[1]THALES, 3 avenue de l'Europe, 31400 TOULOUSE

ABSTRACT

Big Data technologies have been at the heart of THALES strategy for already six years. In French South West division, development teams are mainly working with CNES and space industry actors and have acquired knowledge on space ground segments. THALES gets its first experience mixing Big Data and space mission by cooperating with CNES on GAIA ESA mission, using the famous Hadoop technology. We then enriched our knowledge through several projects or prototypes (L2PF processing platform for EUMETSAT, Image Processing). This Paper will present our strategy implementing Big Data processing framework and their integration for space (or other) data processing. Beginning with different experiences in processing ground segments, with each time useful things we have learned, we will finish by generalizing use cases and practices that can help processing your data.

Index Terms— Hadoop, HDFS, Cascading, Spark, YaRN, Storm, Kafka, Image Processing, Ground segment, GAIA, CNES, EUMETSAT, THALES

1. GAIA CNES DPC: HADOOP IN A PROCESSING GROUND SEGMENT

1.1. Gaia Mission and DPC Architecture

No need to present in details the Gaia mission, which aims at building an extremely precise three dimensional map of our galaxy. Its mains characteristics applied to CNES Data Processing Center are the following:

- 1 billion stars to observe,
- Several 100's billions of measures,
- About 3 Petabytes of data to handle,
- Lot of different scientific algorithms to chain,
- Complex data manipulation (join, sort).

In order to answer these needs, the main processing solution can be summarized in two words: Hadoop and Phoebus. Hadoop is probably the most famous term when people are talking about Big Data. This open source framework has benefited from its immense community and the growing needs for huge volume processing to become the most used ecosystem. Its main features are Map Reduce, HDFS and its scalability [1].

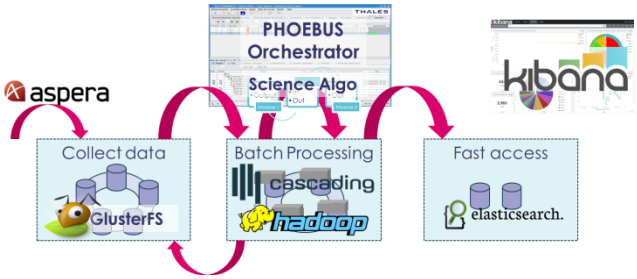

Figure 1: Gaia DPC global architecture

It is the core of GAIA CNES Data Processing Center, specifically the System of Accommodation for Gaia Algorithm (SAGA) that effectively processes the data (Figure 1). On a higher level is Phoebus, an orchestration tool that allows chaining scientific algorithms executed into Hadoop jobs.

It is also important to highlight the specificities of a Hadoop cluster (commodity hardware, meaning low costs), and also the fact that DPCC deployment is entirely automatized using DevOps tooling (Puppet and Foreman).

1.2. Gaia DPCC with Hadoop major feedback

From the design of the DPCC with Hadoop, to its operational qualification, we have learned some really useful things on Big Data processing center design that we would like to share:

1. Defining a clear interface for developing scientific algorithm and in particular the API for data access was very important: data is provided and managed by the processing platform; algorithms never do any POSIX access on any file.
2. Having a framework that allows executing integrated algorithm (i.e. Hadoop jobs) both on a local machine, on a development cluster, or fully integrated with the entire software layer and operational infrastructure is a must-have.
3. Data serialization and compression must be seriously taken care of; it can cost a great deal of computing power or storage space.
4. The data acquisition layer, how the data are inserted into the Big Data architecture, is really important.
5. Hadoop or Cascading is not magic, algorithm implementation and data request must take care of and be designed according to the processing distribution.

1.3. Gaia selection request using Spark

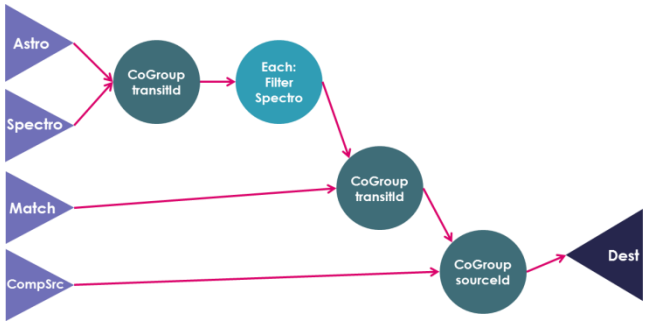

Figure 2: GAIA CU6 SPP Simplified Query

Because Spark has been one of the most promising (and now widely used in production) technology of the Big Data world for nearly two years, we did a little study with it in the Gaia context. It was quite simple, because Spark is shipped with our Hadoop Cloudera distribution, and can easily be deployed on the YaRN cluster. The principal thing to do was to implement our Cascading Query into Spark model.

We have performed the tests on a simplified CU6 data selection query (performing three distinct joins, see Figure 2), and the results are crystal clear: Spark is almost three times faster than Cascading running on MapReduce. We can also observe the linearity of performances of both Cascading and Spark when increasing input data volume.

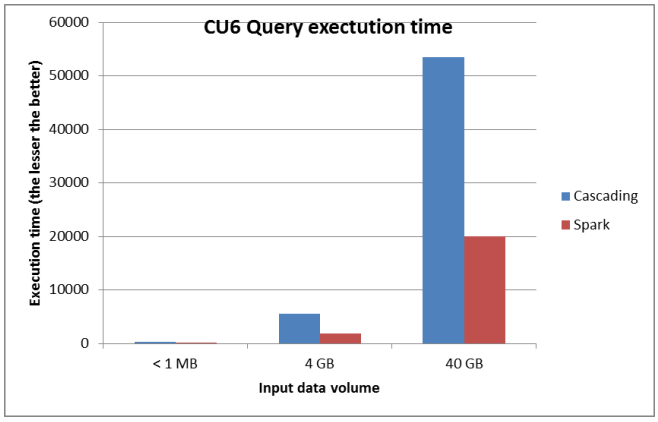

Figure 3: Query execution time: Cascading vs Spark

2. EUMETSAT MTG-L2PF: STREAMING PROCESSING WITH STORM AND KAFKA

In batch processing, we need to handle big dataset independently; Hadoop (and Spark) is very efficient to do that. But in a lot of space mission context we need to process the incoming data as soon as the telemetry is transferred on the ground segment. This is particularly true in meteorological context as we have to predict the weather as soon as possible when receiving satellite packets.

2.1. MTG-L2PF Mission and Architecture

In EUMETSAT context, data is incoming at a huge rate, and must be processed with a really low latency. Here are the principal characteristics of the mission:

- A total of 6 satellites with nominal orbit deployment configuration with 3 satellites,
- Near-real-time operations, 24x7, with stringent timeliness constraints,
- The overall downlink in excess of 600MB/s.

The answer from THALES to make complex operations on such streaming flows relies on the following technologies:

- Kafka for handling the data message layer in a distributed and very efficient way,
- Storm [2] for applying the complex transformations and manipulation on them.

We are currently implementing these solutions for EUMETSAT MTG-L2PF program.

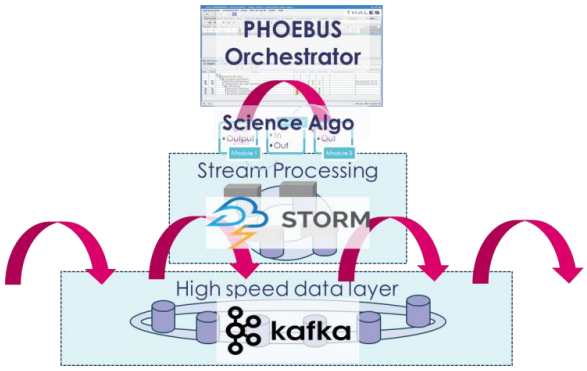

Figure 4: MTG-L2PF Global architecture

2.2. MTG-L2PF Core framework: with Gaia lessons and new challenges

The framework we are developing for MTG-L2PF mission relies on our experience for designing the CNES Gaia DPC, we've retained strong Gaia feature, and improved some less taken care of points:

- We kept the clear separation between scientific algorithm parts and data manipulation,
- But we improved the framework in order to give an access to the algorithmic power of Big Data framework to the science: scientific code can execute anywhere on a data flow or topology.

doi: 10.2788/854791

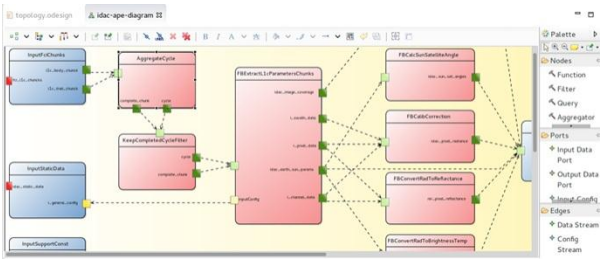

Figure 5: MTG-L2PF topology design tool

- We kept the versatile framework concept used in Gaia: with local execution for development, and distributed execution on any platform,
- We improved the way of integrating algorithm by giving a graphical tool for designing topologies (Figure 5).
- We strongly studied the best possible serialization to use (ex: FST, Avro, Kryo …).
- We also have identified our acquisition layer, and as we are working in a streaming context using Kafka, it integrates well in our architecture.
- We work in an agile way, with all the parts of the project working together on sprints.

Some challenges we have also identified:

- Streaming processing is a whole new concept, not having a coherent dataset from the start forces to see things in a different way.
- How to imply the science algorithm development more deeply and make persons developing them aware of the final execution process.

3. IMAGE PROCESSING: USING HADOOP AND SPARK

One big problematic in Toulouse is the images production ground segments for earth observation. We took a big part on PLEAIDES processing ground segment realization. Images storage volume is clearly in the scope of Big Data now; it is natural that we wanted to try to port image processing on Big Data framework. The main problem is the difficulty to parallelize image access. We have been working on this issue for already three years, with a first prototype using Hadoop Map Reduce plugged with Gluster File System. We are now experiencing on plugging Image Processing Algorithm with Spark.

3.1. Hadoop and GlusterFS in the cloud

The principal difficulty we had is that the image processing algorithm we use does not follow the first principle identified in Gaia and MTG project: they directly do POSIX access on file system to open image files.

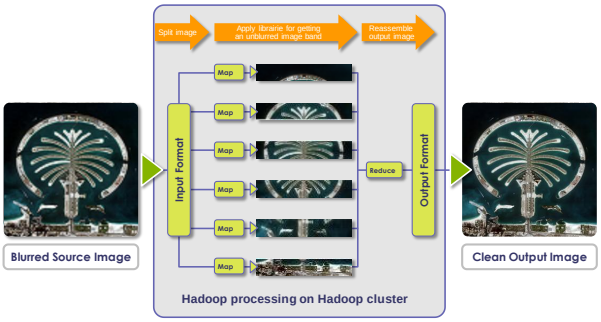

Figure 6: Image Processing on Hadoop and GlusterFS

Our first solution to overcome this was to use a distributed file system that proposes POSIX access as our storage layer, namely GlusterFS (Figure 6). Other solutions could have been implemented, with MapR NFS access to HDFS for instance. The prototype was deployed in THALES cloud to perform demonstrations.

Even is this prototype was working, it was clearly not satisfying, and was lacking a part of the parallelization power of Big data solution.

3.2. Spark, data tiling, and data interface

In order to use Big data technologies effectively in the context of image processing, and to integrate it well with the Big Data processing framework strategy we used on GAIA or MTG-L2PF, we must get rid of both the image presentation as a whole, and of the POSIX data access. So that means data must be tiled, and stored as a collection into a column format. Then algorithm must be able to work on these tiles (this is already possible with some of them), but mainly they must be able to access data from a byte buffer in memory, and not directly from the disk.

Combining these two concepts with the use of Spark, which as we have previously said can be seen as a replacement of MapReduce with an API abstraction as powerful of Cascading, we can describe really complex data manipulation, working on standard HDFS storage or other object storage or data layer.

Figure 7: image processing data flow example

It is important to note that the data manipulation part is really a piece of the algorithm, and can perform really powerful tasks. For example, with such a library as Spark with its MLib sub module, it is possible to execute Machine Learning algorithms in order for example to detect shapes on observation images, or to classify pictures.

4. DATA PROCESSING GENERALIZATION

4.1. Best practices identification

Based on our experience on Gaia, MTG-L2PF and on image processing, we've identified a set of best practices when designing a Big Data processing platform that we can sum up here:

- Define a clear interface for scientific algorithm implementation,
- Scientific algorithm access to the data through the interface, directly in memory, no POSIX access,
- Pay a great attention to your data serialization and compression, for saving computing power and storage capacity,
- Use data manipulation API as part as your algorithm, they are extremely powerful,
- But pay attention to really distribute your data into a coherent structure,
- Test your algorithms using local execution mode as soon as possible, but don't forget to try them on a real cluster.
- For first validation purpose, make technical implementations of processing chains, which use and prove all the concepts provided by the framework.

4.2. Big Data processing framework

When realizing frameworks for integrating scientific algorithm on Big Data processing facilities, we also observed that the major technologies now relies on high level API with operators that are mostly equivalent:

- Cascading: Pipe, Each, CoGroup, GroupBy;
- Spark: RDD, map or flatMap, join, groupByKey;
- Storm Trident: Stream, each, join, groupBy,
- Google Dataflow: Pipeline, ParDO, Join, GroupByKey.

This observation leads us to two conclusions:

1. These API must be good!
2. We must build a framework that can easily take into advantage one API or the other.

Our frameworks implementations, both for Gaia and for MTG-L2PF are thus based on a lot of generic component that can be plugged together using a description file or a specific GUI. Plugging those components form a data flow that chains data manipulation and scientific code execution. Then a part of code is generated according to this description file and is what we will be able to execute on our platform. The final idea is that we must be able to handle any kind of data on any kind of Big Data framework.

4.3. Big Data technology stack on lambda architecture

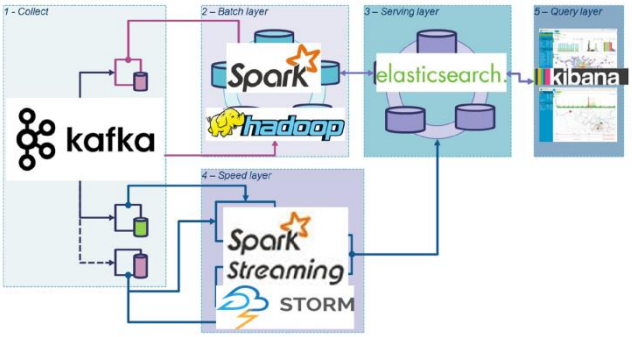

Figure 8: THALES Reference Architecture and technologies for Big Data

With all these experiences, THALES has now developed a deep knowledge in several technologies, and we have defined reference architecture for Big Data projects based on the lambda architecture [3] and our experiments. Either a sub part or the all of this architecture is used in most of our starting projects, which can be either Processing Center or Data Mining pipelines.

10. REFERENCES

[1] T. White, "Hadoop, the definitive guide", 2 nd ed, Oreilly, 2010.

[2] Ankit Jain and Anand Nayla, "Learning Storm", Packt Publishing, 2014.

[3] Michael Hausenblas and Nathan Bijnens, inspired by Nathan Marz, http://lambda-architecture.net, 2015.

ANALYZING BIG REMOTE SENSING DATA VIA SYMBOLIC MACHINE LEARNING

M. Pesaresi, V. Syrris and A. Julea

European Commission Joint Research Centre (JRC),
Institute for the Protection and Security of the Citizen, Global Security and Crisis Management Unit

ABSTRACT

This paper discusses the applicability of a recently introduced classification method for Earth Observation data and justifies its eligibility in the framework of Big Data processing and analysis. The main novelty that this classifier brings forward in the Remote Sensing domain is the automated learning from symbolic representations as opposed to the traditional quantitative approach. The method has been designed to work as a melting pot at which, data arriving from diverse sources and with heterogeneous characteristics (sensors, scales, time periods) are fused and refined to valuable information. A description of the classification method is presented briefly. In addition, results from the production of the multi-temporal medium resolution Global Human Settlement Layer that derived from Landsat satellite imagery of the last 40 years are demonstrated and assessed.

Index Terms— Big Data, Earth Observation, Symbolic Machine Learning, Classification, GHSL

1. INTRODUCTION

Several factors drove gradually the convergence of Big Data (BD) framework with Earth Observation (EO) domain: i) The establishment and support of two planetary programmes: the Sentinel missions under the European Copernicus programme [1] and the NASA Landsat mission [2]. These satellites send Terabytes of data every day and have set new standards in large-scale data management (storage, retrieval, maintenance, delivery, communication). The political decision for free and open access to these data constitutes a landmark in the history of EO and data exploration; ii) The growth of the online communities and the culture of knowledge sharing that led to a great availability of crowdsourcing data, especially geolocated ones; iii) The advancements and the broad expansion of distributed computing; and iv) The maturation of knowledge discovery and machine learning algorithms that was traced to efficient implementations of opensource libraries and tools, allowing everyone and not only the scientists or the experts to analyse automatically and quickly tons of data bits.

The continuous monitoring of Earth surface and the production of global thematic maps at decametric scale with very frequent temporal update becomes now a feasible scenario, provided that the infrastructure capacity and the proper methodologies for automated data processing and analysis fit together, are configured optimally and utilized productively. The framework of Real Big Remote Sensing Data Scenarios (RBRSDS) that we are interested in, can be defined by the following demands: high volume of satellite data, characterized by heterogeneity due to the variety of sensing devices, collected potentially at different time periods and referring possibly to dissimilar spatial domains (scales). An additional condition, mostly related with the application of crisis management our research group is involved in, is the guaranteed response within strict and specified time constraints (ranging from hours to few days).

Apart from the three classical dimensions that forge primarily the infrastructure side of the Big Data realm, that is, the size/scale of the data (Volume), the diverse sources/forms of the data (Variety) and the data flowing/streaming (Velocity), there are three additional Vs exhibiting prominent position in RBRSDS: uncertainty (bias, noise, irregularity and lack of normality) of data (Veracity), semantics/contextualization of data (Variability) and correctness (accuracy, precision) of data inputs and outputs (Validity). The last three axes are processing-oriented and challenge mostly the machine learning approaches that are based on multiple hypotheses testing, intensive pre-processing and iterative optimization techniques. The Value as principle remains invariant as can be found in the classical data analytics: to transform data to useful information. According to [3], challenges like dealing with a great amount of wide-ranging, heterogeneous and multidimensional training examples are beyond the current in-memory processing and real-time computing capability and remain still open in Geospatial Big Data Machine Learning and consequently in RBRSDS discussed here.

The standard paradigm for extracting information from Remote Sensing data that relies on physical explicit modeling of the causal relationships between target's energy absorption-reflection-emitting properties and sensor technical characteristics [4], is difficult to apply in RBRSDS due to: the high requirements of input data in terms of quality and standardization (stability, calibration), the cost for the collection of necessary ancillary data, and the cost of porting the model in different sensors, or adapting it in the same sensor in case the definition of the target to be detected is changing (eg. model transferring from water detection to built-up identification).

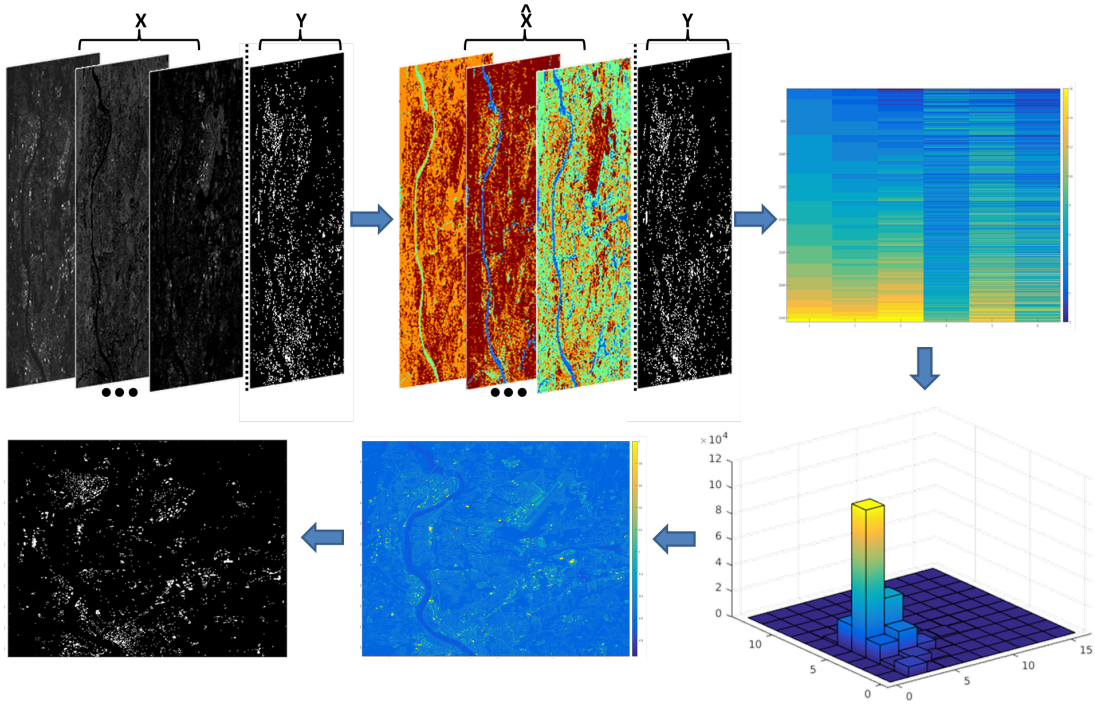

Fig. 1. The SML algorithm in successive stages. First row, from left to right: the input dataset, the dataset after the taxonomy application, the amount of unique sequences. Second row, from right to left: the multidimensional histogram, the interestingness measure and the optional step of crisp classification.

Alternative to the classical causal paradigm is the data-driven exploratory approach, where the machine learns automatically statistical relationships among features/variables based on similarities, differences, trends, covariances, or co-existence frequencies. In image classification, the machine is using either a functional, or a manifold-based, or a rule-based representation to define the class boundaries and to partition suitably the feature space. A large number of machine learning approaches [5] are currently used in state of the art image information discovery and retrieval. It is worth noting that, historically, these algorithms have been tested on small to medium size datasets (compared to RBRSDS), with a moderate to large number of features. Algorithm complexity and lack of high-quality and representative training sets are two main issues that hinder their broad applicability in these scenarios.

Recently, a new classification method named Symbolic Machine Learning (SML) [6], has been proposed in the context of geospatial BD analytics, designed to address explicitly the challenges posed by RBRSDS as stated previously. SML can be considered an evolution of the supervised learning techniques that were applied for the processing of the Global Human Settlement Layer (GHSL) from HR/VHR satellite data [7]. Compared to machine learning approaches currently used in the remote sensing community [8], it shows competitive characteristics in terms of computational cost, capacity to process large datasets and robustness against the input data quality and heterogeneity.

2. THE METHOD

SML [6] is a genuine data-driven approach encapsulating elements of Association Analysis and Machine Learning of symbolic representations. Learning is formulated as a search problem. The concept of variables coexistence (joint contribution) and the respective definition of the optimal confidence level to decide their potential linkage for the classification task at hand, guides the inspection of the hypothesis space. The term *symbolic* refers to the categorical/nominal representation of the feature dataset X, as opposed to the numerical representation of the original input signal. Another factor that affects the searching problem is the quality of the reference set Y used for the class assignment. A strong difference between SML and other classifiers which learn from a relatively small training set, assuming that this set represents at micro-scale the ideal partitioning of the complete universe of events, is the following: SML encounters directly the complete universe of events, considers approximate associations among all the available events and the reference set, and then, refines/redefines these associations to confident relationships.

In short, the input features X are abstracted to $\hat{X}$ based on a taxonomic schema. Next, associations among class labels Y and groups of features $\hat{X}$ (sequences) are defined; we characterize these initial associations *a priori*. Then, a measure is scoring these sequences according to the number of their occurrences in each class. The latter step is de-

doi: 10.2788/854791

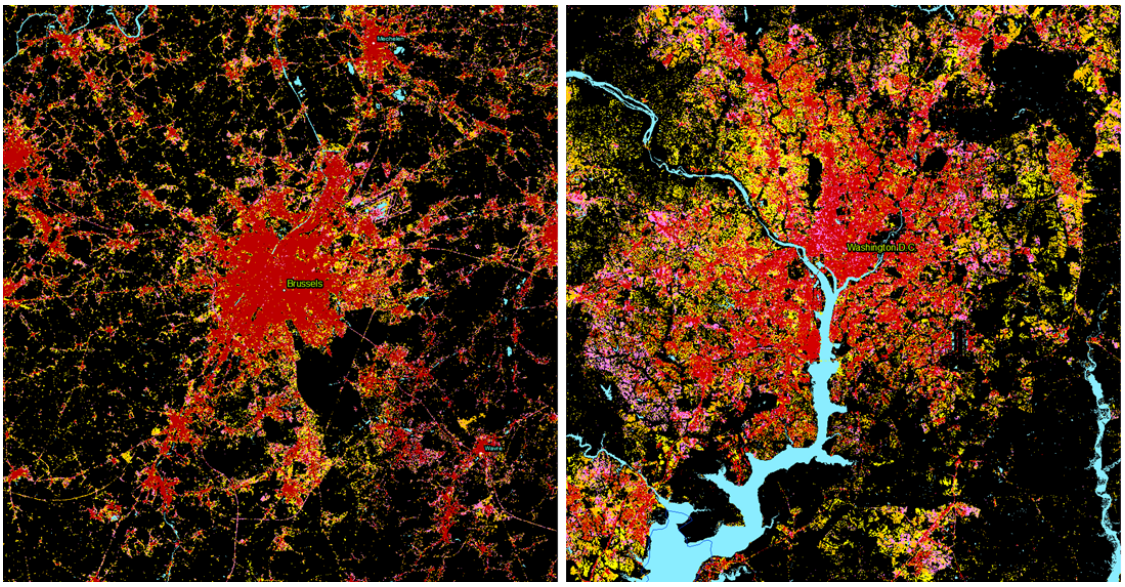

Fig. 2. Multi-temporal representation of two cities at scale 1:300,000: Brussels (left) and Washington (right). Colormap: water (cyan), land classified as non-built-up (black), red (1975), pink (1990), orange (2000) and yellow (2013-2014).

ployed with the use of the empirical multidimensional histogram. Currently, we have experimented with a measure inspired by a probability-based objective interestingness measure [9], namely the Evidence-based Normalized Differential Index. Optionally, using discriminant thresholds determined through the analysis of the cumulative distribution of the computed scores, a reclassification (refinement) takes place at the final stage.

A schematic representation of the SML algorithm is displayed in Fig. 1: 1) A matrix $[X : Y]$ containing the data instances is provided, according to the a priori associations defined among the features X and the classes Y; 2) A taxonomy is applied and determines the data granularity level $\hat{X}$; 3) The frequency of appearance of every unique sequence is estimated through 4) the empirical multidimensional histogram; 5) An interestingness measure is chosen, which scores the event that a group of feature values belongs to a specific class; 6) The last step is optional and addresses the need for crisp classification: thresholds are estimated or functions are defined that assign a sequence to a single class.

The SML functionality is adapted over a dataset through only one parameter, the setting that controls the data granularity. This is the sensitive part of the method and is under study. Low granularity (few features, small encoding alphabet) may lead to underfitting, while high data granularity (big number of features, many encoding symbols) can have the opposite effect (overfitting). Techniques like clustering, quantization and data encoding are under consideration and test.

The SML classification method has several appealing characteristics in the frame of RBRSDS. They are shortly listed below:

- Time and space complexity is independent of the number of observations. It is controlled by the number of unique sequences and the number of features. In [8], the SML complexity has been decomposed and contrasted with the complexity of other classifiers widely used in the statistical learning and data mining community.
- The fast performance allows "on the fly" learning and classification. This fact removes the need for model/bag-of-patterns storage in a knowledge-base for later classification of newcoming data.
- Model-free approach simplifying the fusion of different information sources and ancillary data.
- Histogram-based decision metrics allowing incremental learning in the presence of new data.

The SML robustness against different types of noise/distortion that may exist in a typical reference set Y used in RBRSDS has been assessed in [8].

3. APPLICATION: MULTI-TEMPORAL MEDIUM RESOLUTION GHSL

SML has been successfully applied for the processing of Landsat data records of the past 40 years and led to the production of the first medium resolution global built-up layer with a fully automatic way, permitting thus the mapping of the human settlements evolution at global scale [10].

In the specific application, the information was extracted from Landsat image records organized in four collections, corresponding to the epochs 1975, 1990, 2000 and 2013-14. Seasonality analysis was out of the scope; consequently, the images belonging to the same path and row were considered as distinct cases. Table 1 shows the type of imagery used and quantifies the volume of data and the respective processing time. Fig. 2 illustrates the resulting built-up classification output of Brussels and Washington.

doi: 10.2788/854791

Table 1. Processing of Landsat imagery in terms of Volume, Variety and Velocity

	GLS1975	GLS1990	GLS2000	Landsat-8
No of scenes	7,588	7,375	8,756	9,089
No of bands	4	6	6	9
Used bands	10-20-30-40	10-20-30-40-50-70	10-20-30-40-50-70	1-2-3-4-5-6-7-8+QA
Time* (min)	5	11	13	33

Indicative machinery: Intel(R) Xeon CPU E7420@2.13GHz, 8GB RAM
*Total processing time per image

Table 2 shows the performance metrics used for the comparison between GHSL 2014 outcome and other available global products against fine scale cartography. This reference test set consists of 1,505 raster tiles with a surface of $10 \times 10\ km^2$ each: they are derived from available vector cartographic sources at scale of approx. 1:10,000 [10].

Table 2. Performance statistics for GHSL (2014) and other datasets compared to fine scale building footprints.

	Total Accuracy	Balanced Accuracy	Informedness	Kappa	ComErr	OmErr
GHSL2014	.90±.10	.67±.11	.34±.20	.33±.20	.53±.20	.54±.30
GlobCover 2009v2.3[11]	.89±.13	.54±.06	.07±.11	.08±.12	.57±.23	.87±.21
FROM-GLC[12]	.86±.14	.55±.07	.11±.14	.11±.12	.67±.24	.75±.32
Global Land30[13]	.88±.11	.65±.11	.30±.23	.27±.20	.61±.20	.56±.34
MODIS[14]	.86±.15	.55±.07	.09±.14	.09±.12	.68±.20	.76±.32

4. CONCLUSIONS

The eligibility of the Symbolic Machine Learning classification method in the framework of Big Data analytics and Earth Observation demanding scenarios was debated. The multi-temporal medium resolution GHSL was presented shortly, demonstrating the applicability of the method in a real scenario. The same classification approach is under test in the GHSL production, using as input Sentinel 1-2 data [15].

5. REFERENCES

[1] ESA, "Copernicus programme," `http://www.esa.int/Our_Activities/Observing_the_Earth/Copernicus/`, [accessed Feb 10, 2016].

[2] NASA, "Landsat programme," `https://www.nasa.gov/mission_pages/landsat/main/index.html`, [accessed Feb 10, 2016].

[3] T. van Zyl, *Big Data: Techniques and Technologies in Geoinformatics*, chapter 7 Machine Learning on Geospatial Big Data, pp. 133–148, CRC Press, 2014.

[4] T. M. Lillesand, R. W. Kiefer, J. W. Chipman, et al., *Remote sensing and image interpretation.*, chapter 1 Concepts and Foundations of Remote Sensing, pp. 1–45, Number Ed. 5. John Wiley & Sons Ltd, 2004.

[5] D. Lu and Q. Weng, "A survey of image classification methods and techniques for improving classification performance," *International Journal of Remote Sensing*, vol. 28, no. 5, pp. 823–870, 2007.

[6] M. Pesaresi, V. Syrris, and A. Julea, "A New Method for Earth Observation Data Analytics Based on Symbolic Machine Learning," *Remote Sensing*, 2015, submitted.

[7] M. Pesaresi, Guo Huadong, X. Blaes, D. Ehrlich, S. Ferri, L. Gueguen, M. Halkia, M. Kauffmann, T. Kemper, Linlin Lu, M.A. Marin-Herrera, G.K. Ouzounis, M. Scavazzon, P. Soille, V. Syrris, and L. Zanchetta, "A Global Human Settlement Layer From Optical HR/VHR RS Data: Concept and First Results," *Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of*, vol. 6, no. 5, pp. 2102–2131, Oct 2013.

[8] M. Pesaresi, V. Syrris, and A. Julea, "Benchmarking of the Symbolic Machine Learning classifier with state of the art image classification methods - application to remote sensing imagery," JRC Technical Report EUR 27518 EN, EC, JRC, IPSC, 2015.

[9] C. Tew, C. Giraud-Carrier, K. Tanner, and S. Burton, "Behavior-based clustering and analysis of interestingness measures for association rule mining," *Data Mining and Knowledge Discovery*, vol. 28, no. 4, pp. 1004–1045, 2014.

[10] M. Pesaresi, D. Ehrlich, S. Ferri, A. J. Florczyk, S. Freire, M. Halkia, A. Julea, T. Kemper, P. Soille, and V. Syrris, "Operating procedure for the production of the Global Human Settlement Layer from Landsat data of the epochs 1975, 1990, 2000, and 2014," JRC Technical Report EUR 27741 EN, EC, JRC, IPSC, 2015.

[11] V. Kalogirou P. Defourny O. Arino, J. Ramos and F. Achard, "Globcover 2009," in *ESA Living Planet Symposium 2010*.

[12] P. Gong, J. Wang, L. Yu, Y. Zhao, Y. Zhao, L. Liang, Z. Niu, X. Huang, S. Fu, H.and Liu, et al., "Finer resolution observation and monitoring of global land cover: first mapping results with Landsat TM and ETM+ data," *International Journal of Remote Sensing*, vol. 34, no. 7, pp. 2607–2654, 2013.

[13] J. Chen, J. Chen, A. Liao, X. Cao, L. Chen, X. Chen, C. He, G. Han, S. Peng, M. Lu, W. Zhang, X. Tong, and J. Mills, "Global land cover mapping at 30m resolution: A POK-based operational approach," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 103, no. 0, pp. 7 – 27, Oct 2014.

[14] A. Schneider, M. A. Friedl, and D. Potere, "A new map of global urban extent from MODIS satellite data," *Environmental Research Letters*, vol. 4, no. 4, 2009.

[15] M. Pesaresi, C. Corbane, A.Julea, A. J. Florczyk, V. Syrris, and P. Soille, "Assessment of the added-value of Sentinel-2 for detecting built-up areas in the frame of the Global Human Settlement Layer," *Remote Sensing*, 2015, submitted.

BIG DATA MEETS LINKED DATA – WHAT ARE THE OPPORTUNITIES?

Jon Blower [1], Maik Riechert [1], Nino Pace [2], Manolis Koubarakis [3]

1. Department of Meteorology, University of Reading, United Kingdom
2. Advanced Computer Systems, Rome, Italy
3. National and Kapodistrian University of Athens, Greece

ABSTRACT

The worlds of "Big Data" and "Linked Data" are have evolved somewhat separately but both are highly relevant to the development of systems for sharing and processing large volumes of Earth Observation data. This paper focuses on Linked Data, examining this field from a user-centric point of view, based on research carried out in a number of recent European projects. A central theme is that Linked Data can help to open up Earth Observation data to new users and communities.

***Index Terms*—** Earth Observation, Data discovery, Data management, Semantics, Quality, Visualization

1. BIG DATA AND LINKED DATA

The term "Big Data" has numerous definitions but practitioners recognize many challenging aspects, including data *volume* (e.g. where data become too large to process using traditional techniques), data *velocity* (e.g. where data are captured at high rates from sensors), data *variety* (e.g. bringing together data from multiple sources) and data *veracity* (i.e. data quality).

All of these challenges are faced by users of Earth Observation (EO) data. Whereas questions of volume and velocity can be addressed through the development of high-performance "Big Data" capture and processing systems (e.g. http://www.jasmin.ac.uk), questions of variety and veracity can only be addressed by looking more carefully at how we can help users by *describing* data better and by sharing this information more effectively within and outside the EO community. Linked Data techniques provide a number of ways to achieve this. Therefore, Linked Data and Big Data are entirely complementary ideas.

There are many possible views of the concept of "Linked Data". At a high level, the promise of Linked Data is that we will be able to use the full power of the Web to break down barriers between communities that were hitherto separated. The key features of Linked Data are:

1. Information is *decentralized* (like the Web in general). Anybody can say Anything about Any topic (this is the "AAA" principle).
2. Information is *machine-readable* (unlike the wider Web).
3. Through wide use of certain *standards* (see below), diverse sources of information can be linked together unambiguously.

In this way, we do not have to collect all the possible information about a topic into a single place; we can use the Web as a highly-distributed, decentralized research environment. This allows information providers to focus on publishing the information that they "own" and in which they are expert.

The key technical steps that an information provider must follow in order to participate in the Linked Data web are [1]:

1. Create (or, preferably, reuse existing) unique and persistent identifiers for the important "things" in a community (e.g. datasets, publications, algorithms, instruments).
2. Allow users to "look up" these identifiers on the web to find out more information about them (in other words, the identifiers are essentially Web addresses, i.e. URLs).
3. Make this information machine-readable in a community-neutral format. RDF (Resource Description Framework) is the preferred choice, using terms from widely-agreed vocabularies.
4. Within this information, embed links to other things and concepts and say *how* these are related.
5. Optionally, provide web service interfaces to allow the user to perform sophisticated queries over this information. SPARQL is the preferred query language for this.

By describing data more precisely and by using high-level, community-neutral standards like RDF and SPARQL (both are standards of the World Wide Web Consortium, W3C), information should become more widely accessible and understood by users in different fields. Perhaps the main challenge is to identify the most appropriate *vocabularies* (defined terms) to describe data. Interoperability is enhanced if the same vocabularies are widely adopted; however, frequently data providers find that they need to create new terms and identifiers to describe their data accurately. Designing and publishing these new terms can be a significant undertaking. An example of a highly-relevant RDF vocabulary is GeoSPARQL [2], which describes how to encode geospatial data in RDF and query them using SPARQL.

2. USE CASES FOR LINKED DATA IN EO

The technological view of Linked Data is well-established, as is the high-level view of the overall benefits that it can bring. But the concrete benefits to the EO user community are rarely described explicitly. In this section we briefly outline some promising use cases based on experience derived from numerous European research projects, notably TELEIOS (http://www.earthobservatory.eu/), CHARMe (http://www.charme.org.uk, [3]), MELODIES (http://melodiesproject.eu, [4]) and LEO (http://linkedeodata.eu [5]).

2.1. Enabling data discovery

In the past years, considerable effort has been devoted to the development of special-purpose catalogue and search infrastructure to enable discovery of scientific data, such as the Global Earth Observation System of Systems (GEOSS).

To complement these initiatives, Linked Data can also help users to find relevant information using standard Web mechanisms. One way is simply for the user (or automated system) to traverse the links between information sources as they would on the Web. Secondly, we believe that in the coming years it will become increasingly possible to discover EO data through "mass market" search engines such as Google and Bing. These search engines are adopting Linked Data techniques for harvesting structured (RDF) information about datasets from websites. Data providers can describe their datasets using an appropriate vocabulary (e.g. http://schema.org), and search engines can harvest this information. The potential is that search engines will be able to provide richer information about datasets they find, and automatically link this with other related information. The need for special-purpose, community-specific search engines may therefore reduce and EO data will become easier to discover by new users.

2.2. Helping users to understand data.

By enabling different sources of information to be linked together in a structured fashion, Linked Data can help users to understand much more about datasets. In particular, the interrelationships between a dataset and the information surrounding it (e.g. upstream data, processing chains, publications, experts) can be described quite naturally in a Linked Data form and discovered by users (see Figure 1). User-supplied annotations (see section 2.4) can also play a role here.

2.3. Dynamically combining and integrating data

Currently the process of combining multiple data sources from different communities (e.g. in a web application) is usually lengthy and dominated by low-level technical concerns of data reformatting. The use of RDF as a universal "lingua franca" and SPARQL as a standard web service interface can greatly help developers of such systems to bring data together *at runtime* from multiple sources, without the need for further data manipulation, harmonization or conversion. Now that the essential tools are maturing (see section 3), such applications are beginning to emerge (e.g. http://almere.pilod.nl/bgtld/v2).

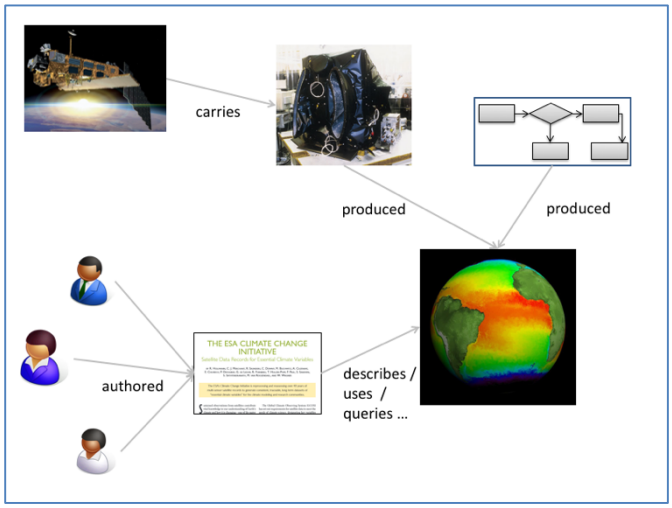

Figure 1: A simplified illustration of the links between EO platforms, instruments, algorithms, datasets, publications and scientists, represented as Linked Data.

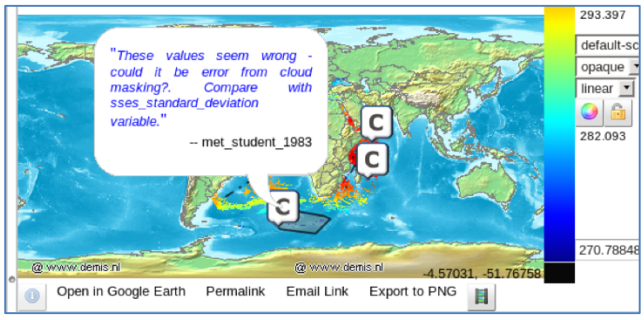

Figure 2: The "CHARMe Maps" tool, which allows user-supplied annotations to be linked with datasets and even specific subsets of datasets. Each annotation is represented by a "C" icon.

2.4. Allowing users to share information with each other

Currently users of EO data receive most of their information from the original provider of the data. However, the provider usually does not usually have complete information about important topics such as data quality and fitness for purpose; such information is often gathered or supplemented by the user community. The CHARMe project developed a system, based upon Linked Data, to allow users to create and share *annotations* about EO data [3]. These annotations include free-text comments, links to relevant publications, statements about data quality and more (Figure 2). The

system exploits the fact that data providers are increasingly assigning persistent identifiers (e.g. DOIs) to their data. These identifiers act as "anchors" for the annotations, so that users can be very clear about what they are annotating, be it a dataset, a single satellite image, a sensor or even another annotation.

3. NEW TOOLS AND APPROACHES

Data structures and services such as RDF and SPARQL are not familiar to most users of Earth Observation data. This section describes a suite of recently-developed open-source tools that help data providers and users.

Strabon [6] is an RDF data store, that is capable of storing geospatial Linked Data that changes over time. Data can be queried using both the OGC-standard GeoSPARQL dialect [2] and stSPARQL. Strabon supports a wide range of geographic functionality, such as support for spatial geometries and a wide range of coordinate reference systems. Strabon can be used to model temporal domains and concepts such as events and facts that change over time.

GeoTriples [7] is an open source tool for converting geospatial data from several common formats into RDF. It is able to generate and process mappings from source files (GML, KML, GeoJSON, ESRI Shapefile, CSV) and spatially-enabled relational databases to RDF graphs. GeoTriples employs by default the well-known ontologies like GeoSPARQL and stSPARQL, without being tightly coupled to a specific vocabulary. It offers rich support for processing geospatial data and is able to scale to large data volumes.

Ontop-spatial (https://github.com/constantB) is a geospatial extension of the system Ontop (http://ontop.inf.unibz.it). In contrast with Strabon, Ontop-spatial does not require that data be converted stored in a special RDF data store. Instead data remain in their source format (e.g. relational databases or shapefiles), and Ontop-spatial creates a "virtual" RDF graph that can be queried efficiently using SPARQL and GeoSPARQL. Therefore this system is suitable for data providers who prefer not to convert their data to RDF.

The **Silk** Link Discovery Framework (http://silkframework.org/) can be used in order to discover relationships between entities in the datasets. In particular, using Silk one can pre-process the data and create spatial, temporal and other types of relationships between data items. This can greatly increase the efficiency of complex queries.

Sextant (http://sextant.di.uoa.gr/) is a web- and mobile-based application for exploring, interacting and visualizing time-evolving linked geospatial data [8]. The functionalities of Sextant include the exploration and visualization of linked spatiotemporal data, the creation, sharing, searching and collaborative editing of maps and the production of statistical charts. While the tool heavily utilizes semantic web technologies, Sextant provides a user-friendly interface (Figure 3) to allow both domain experts and non-experts to use all features provided. Sextant can ingest heterogeneous data from files and web services that support a variety of standards, including KML, SPARQL and GeoSPARQL.

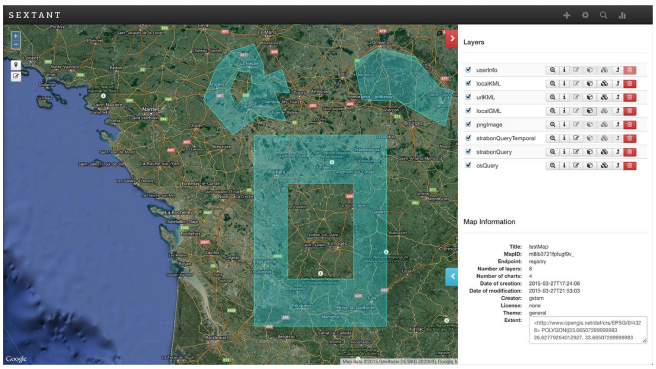

Figure 3: Screenshot of Sextant, showing visual integration of Linked Data from a number of sources.

The technologies we have described so far mainly relate to vector data, i.e. data about discrete geospatial features such as the locations of in situ observations and the boundaries of satellite images. However, satellite image data themselves usually take the form of multidimensional arrays, which can be very large in size. It would not be efficient to convert all the individual pixels in large images to RDF structures. However, the metadata about images (e.g. concerning the variables that are measured and the instrument itself) is highly amenable to being described in RDF. Therefore the EO community needs a way to handle RDF and array data seamlessly.

CoverageJSON (http://tinyurl.com/covjson) is a new (currently experimental) lightweight format for encoding scientific data (including Earth Observation data) in the JSON format. JSON (JavaScript Object Notation) is the predominant format preferred by modern web developers for consuming data feeds in web-connected applications. A "coverage" is a data structure that maps positions in space and time to data values, and can be thought of as an overarching concept that encompasses many kinds of scientific data (see the ISO19123 standard).

CoverageJSON has two main goals: (1) to integrate multidimensional array data with detailed semantic information in the same format; and (2) to make it easier for web developers to ingest scientific data into interactive web applications for in-browser visualization and processing. The key to the format is the use of JSON-LD (http://json-ld.org, an RDF variant) for encoding semantic information, whilst using "plain" JSON for data elements such as arrays. This allows the format to be easy-to-use and efficient, whilst retaining the expressive power of RDF for those applications that require it. It is therefore intended to provide a bridge between the EO and Linked Data communities.

The format has been developed and tested against use cases within the MELODIES project: e.g. in-browser reclassification of land cover datasets (Figure 4), subsetting via user-supplied polygon shapes, derivation of statistics, and intercomparison of multiple datasets.

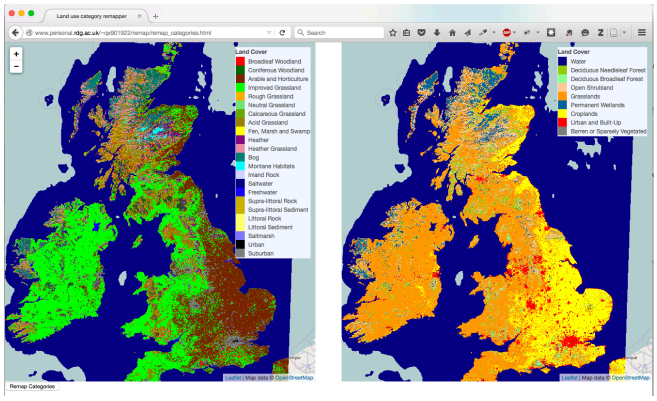

Figure 4: Reclassification of a land cover product from the MELODIES project from the MELODIES classification scheme (left) to the MODIS scheme (right). Data are loaded from a server in CoverageJSON format, then visualized and manipulated within the web browser.

4. CONCLUSIONS

In this paper we have discussed how Big Data and Linked Data are related concepts, contributing to advances in data discovery, understanding and analysis. A key current challenge is efficiency: RDF is extremely expressive but is not usually the most efficient format for data storage or analysis. Furthermore, we find that the conversion of data to RDF can be a burden on application developers, if the original data provider has not provided data in this way. Both of these issues are being addressed through the development of new tools and approaches, but it is clear that Linked Data techniques must be used *alongside* other more efficient "Big Data" techniques, depending on the user's need.

A key goal of Linked Data is to make data more widely accessible by a much wider community. We recommend that EO data providers and application developers consider the following activities to contribute towards this goal:

1. Publish "fundamental" information (such as authoritative information on datasets, instruments, satellites etc.) as Linked Data, so that the community can discover and link to it.
2. Consider publishing data and metadata in widely-used, "web-friendly" formats such as JSON(-LD), thereby making information more usable by typical developers who are not EO data specialists.
3. Publish structured metadata in forms that are understood by mass-market search engines, to enable easier data discovery.
4. Participate in joint efforts between the web community and the geospatial community in order to contribute to the standards that will bring these communities together. A current example of such an initiative is the joint OGC/W3C Spatial Data on the Web Working Group (http://tinyurl.com/sdwwg).

ACKNOWLEDGEMENTS

The TELEIOS, MELODIES, LEO and CHARMe projects have received funding from the European Union's Seventh Programme for research, technological development and demonstration under grant agreement numbers 257662, 603525, 611141 and 312641 respectively.

REFERENCES

[1] C. Bizer, T. Heath and T. Berners-Lee (2009) Linked Data - The Story So Far. International Journal on Semantic Web and Information Systems, Vol. 5(3), Pages 1-22.

[2] M. Perry and J. Harding (eds.) (2012) OGC GeoSPARQL - A Geographic Query Language for RDF Data v1.0 Open Geospatial Consortium document 11-052r4

[3] D. Clifford, R. Alegre, V. Bennett, J. Blower, C. DeLuca, P. Kershaw, C. Lynnes, C. Mattmann, R. Phipps, and I. Rozum (2015) Capturing and sharing our collective expertise on climate data: the CHARMe project. Bulletin of the American Meteorological Society. doi:10.1175/BAMS-D-14-00189.1

[4] J. Blower, D.Clifford, P. Gonçalves, and M. Koubarakis (2014) The MELODIES project: integrating diverse data using Linked Data and cloud computing. In: 2014 ESA conference on Big Data from Space (BiDS'14), Frascati, pp. 244-247.

[5] M. Koubarakis (2014) Linked Open Earth Observation Data: The LEO Project. Informal Proceedings of the Image Information Mining Conference: The Sentinels Era (ESA-EUSC-JRC 2014). Bucharest, Romania, 5-7 March.

[6] K. Kyzirakos, M. Karpathiotakis and M. Koubarakis (2012) Strabon: A Semantic Geospatial DBMS. In the 11th International Semantic Web Conference (ISWC 2012). Boston, USA.

[7] K. Kyzirakos, I. Vlachopoulos, D. Savva, S. Manegold, M. Koubarakis. GeoTriples (2014) A Tool for Publishing Geospatial Data as RDF Graphs Using R2RML Mappings. Terra Cognita 2014, 6th International Workshop on the Foundations, Technologies and Applications of the Geospatial Web, in conjuction with ISWC 2014. Riva del Garda, Trentino, Italy.

[8] K. Bereta, C. Nikolaou, M. Karpathiotakis, K. Kyzirakos, and M. Koubarakis. SexTant (2013) Visualizing Time-Evolving Linked Geospatial Data. In the 12th International Semantic Web Conference (ISWC 2013). Sydney, Australia. Demo paper.

ESA SKY: A TOOL TO DISCOVER ASTRONOMICAL DATA FROM ESA MISSIONS

Jesús Salgado , Bruno Merín* , Fabrizio Giordano* , Deborah Baines* , Elena Racero* , Belén López Martí* , María H. Sarmiento* , Raúl Gutiérrez*, Pilar de Teodoro*, Christophe Arviset**
*ESAC Science Data Centre, ESAC, Madrid

ABSTRACT

In the last years, the amount of data and resources of astrophysical data available for the scientific community at the different project science archives at ESAC is growing quite fast. The project archives offer a quite detailed and easy access to all mission data. However, the multi-wavelength discovery of data could be cumbersome and it usually requires specific knowledge of mission dependent language.

The ESAC Science Data Centre (ESDC) is working on a science-driven discovery portal, called ESA Sky, that allows the exploration of the astronomical resources (almost all the images and catalogues from ESA missions at the present stage) using a simple, intuitive and project agnostic portal.

Using techniques like visualization of multi-order all-sky mosaics based on HEALPix (HiPS), missions coverage (MOC), observational footprints, TAP services on common data models for fast and performant searches, DB geometrical indexes, internal connections between databases and wrappers around the project archives to download the final science ready data allows the handling of big amounts of data in a simplified way.

We will present the recently released first version of this tool, technologies used and future plans.

***Index Terms*—**big data exploitation, heterogeneous data sources, interoperability and standards, linked data and semantics, space science, visualization and visual analytics

1. INTRODUCTION

ESA Astronomical, Planetary and Heliophysics missions have their data archives at ESAC, Madrid. The different archives cover huge wavelength ranges and are an important resource of knowledge for the science community. In order to make this data available, the ESAC Science Data Centre (ESDC) [1] provides a set of project archives where the experts and scientists linked to a certain mission can access the data in an easy and powerful way. The different archives cover web based user interfaces to RESTful [2] interfaces, to allow the integration of the data access from different missions from the command line.

However, the different archives are focused on the discovery of data for particular missions, by using advanced query parameters based on the instruments and using project dependent jargon.

In 2014 a set of scientists and engineers at ESAC embarked on a task to design and implement a general multi-mission interface prototype that could fulfill the following basic goals:

- Allow a discovery interface, mainly based on the access to multi-wavelength data from different missions in a transparent way
- Implement an accurate discovery of observations using geometrical queries
- Use of project agnostic language

Using this paradigm, an application able to interconnect all the different archives using a discovery interface was developed as a prototype and presented internally to ESA members and projects.

There was the general agreement to convert this prototype into an operational application and, in Autumn 2015, the first beta version of the ESA Sky interface was presented to the general scientific community during the ADASS XXV conference in Sydney.

2. ARCHITECTURE

ESA Sky reuses modules and protocols defined by the IVOA (International Virtual Observatory Alliance) as a probe that the libraries and standards provided by this organization help on the design and implementation of archives modules. The application is composed of a client that makes use of the CDS Aladin Lite [3] module to visualize All-Sky maps.

The All-Sky HiPS [4] maps have been generated for all ESA missions by collaboration with the different project experts and selecting the best possible input data. HiPS are conceived as multilevel all sky maps where the data that travels from the server to the client is only the one that the client could need for visualization at the selected zoom level. That implies different resolution orders and a tesellation of the sky following *HEALPix* [5] indexes. In contrast with other kind of map projections used by other applications like, e.g. Google Sky, HEALPix makes use of

equiareal elements and good projections on the map poles, which is particularly important in astronomy.

A Table Access Protocol complaint server, TAP, [6] is used to discover observational data from ESA astronomical missions by using geometrical queries on accurate footprints. An accurate footprint of an observational based mission is considered the spatial coverage on the sky where photons have been accumulated. It is particularly important to have this kind of accurate footprint for the observations so a search for data for a certain astronomical object does not provide false positives in the results. This is why, for most of the missions, the footprints have been also generated for the ESA Sky project using only the best possible description by, e.g., ignoring sections of the sky where there is no real science data, although that spatial region could be inside the instrument coverage projection on the sky.

Most of the ESA catalogues are also provided through this TAP interface and connected to the ESA Sky graphical interface to enable the implementation of complex science use cases. Also, as all this metadata are available through a standard service, other TAP clients can make use of this service to implement, e.g. access to command line scripts.

As a last step, direct links to the full data are present to redirect the users to the different project archives. By doing that, a user can navigate from quite general all-sky maps in different wavelengths to the ultimate science data from the project archives.

The ESA Sky user interface makes use of Google Web Toolkit (GWT) technology [7] that allows a quite configurable and powerful interface for the users and, at the same time, maintainable code from the technical point of view. GWT components are written in the Java language (although, only a subset of the java classes that can be mapped to javascript are supported) and then converted to javascript after compilation. This is why a number of Integrated Development Environments (IDEs) are available to code, debug and compile GWT, allowing the development of very complex projects that can be evolved in an easy and controlled way.

3. ESA SKY GRAPHICAL USER INTERFACE

In order to provide the best possible HiPS for the different ESA missions, the generation of these maps have been coordinated between ESDC members for products generation, ESA mission members for selection of the best image products, filtering and cut-outs and CDS for providing the provides to the user HiPS generator library. ESDC has implemented a GWT wrapper on top of Aladin Lite to provide a better integration in a more complex GWT User Interface. ESA Sky makes use of the powerful CDS Aladin Lite component to display All-Sky mosaics following the HiPS format. Aladin Lite is the simplified javascript version of the well-known java application Aladin from CDS and allows an easy integration in any page/archive. In this particular case, the integration has implied the creation of a wrapper around this tool to have a GWT compliant component that can be reused, not only within ESA Sky but, also inside any of the ESA astronomical archives.

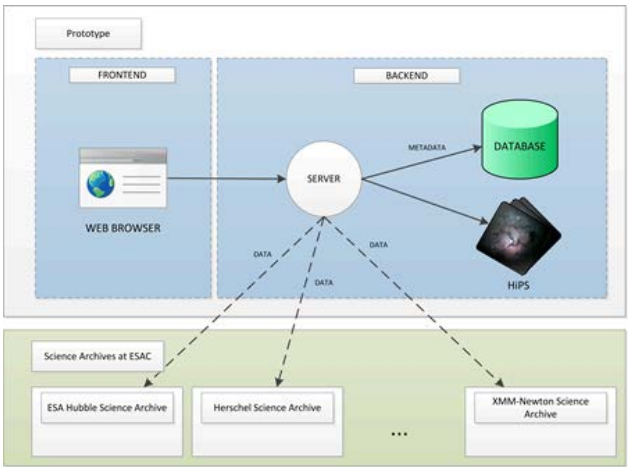

Figure 1: ESA Sky Architecture

The link to the ESA project archives is done by over-imposed detailed footprints that are shown under user requests. This is performed through the TAP service, which when invoked returns for different areas of the sky not only the detailed footprints but also tabular metadata with characterization information for the different observations. These metadata are shown in a tabular display for further access. The mechanism used for the access of observational metadata is based on the IVOA ObsCore TAP standard [8]. These metadata contain a link to a quick look image in graphical format and a link to access the data for this specific observational item. Both links are pointing to the original project archive so there is no duplication of scientific data at the ESA Sky level.

The same concept is applied to source catalogues with the representation of the positions as points in the graphical interface and tabular metadata with possible links to source related products like, e.g. finding charts.

A search on the interface provides to the user with an estimation of the number of observations per mission and sources per catalogue on the displayed area of the sky. In order to provide a fast response to this massive query, both at observational level and at source catalogues level, a *HEALPix* density calculation has been pre-calculated for the full sky so the number of observations and sources displayed is then a very fast sum of pre-computed density numbers over the visible region of the sky.
For the visualization, the HEALPix Multi-Order Coverage maps (MOCs) [9] are displayed for the observational results, when the area of the sky shown is very large. This is particularly useful when the mission spatial coverage is only a small fraction of the full sky as it helps to identify easily the regions of the sky where the mission has observations.

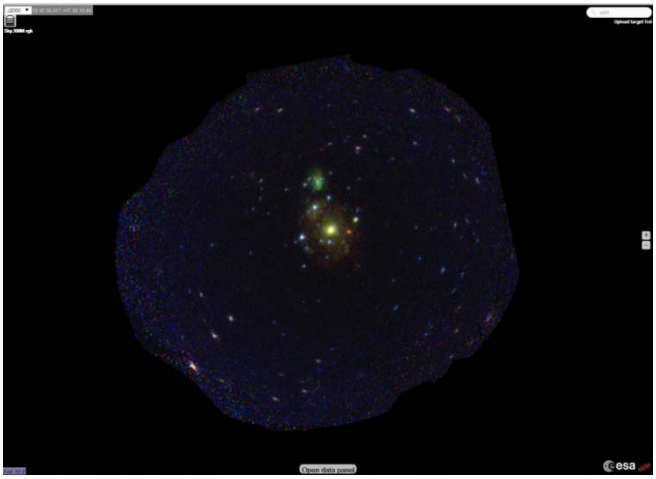

Figure 2: Fraction of the improved XMM-Newton HiPS map, created by XMM SOC and ESDC (P. Rodriguez, E. Racero et al)

The combination of MOCs, fast counts for all ESA observational data and catalogues and the presentation of results per spectral range provides to the user a quite accurate estimation of the amount of data available per wavelength range for all ESA astronomical missions in only one step. This is a rather innovative way to access big amounts of data by accessing summary histograms. This approach can be used not only in astronomy and for this particular use case, but also from a general point of view when the amount of data is so huge that a tabular representation in the first go could not provide a clear order of magnitude of the magnitude of the data found.

4. DATA IN ESA SKY

Currently, ESA Sky includes access to the data from the following missions:

- XMM-Newton images: HiPS all sky maps generated at ESAC with the support of XMM SOC. Access to observational data through detailed footprints and to data connecting to the XSA (XMM-Newton Science Archive)
- XMM-Newton catalogues: Slew catalogue, 3XMM, OM
- Herschel images: HiPS all sky maps generated at ESAC with the support of Herschel SOC. Access to observational data through detailed footprints and to data connecting to the HSA (Herschel Science Archive)
- HST images: HiPS all sky maps generated at ESAC. Access to observational data through detailed footprints and to data connecting to the eHST (European HST archive at ESAC)
- HST catalogues: Hubble Source Catalogue
- ISO images: HiPS all sky maps generated at ESAC
- Planck maps: HiPS all sky maps generated at ESAC
- Planck catalogues: PCCS2 (Second Planck Catalogue of Compact Sources), PGCC2 (Second Planck Catalogue of Galactic Cold Clumps) and PSZ2 (Second Planck Sunyaev-Zeldovich catalogue)
- Integral images: HiPS all sky maps generated at ESAC with the support of Integral SOC
- Integral catalogues: INTEGRAL General Reference Catalogue

Other major source catalogues are available for exploration and the access to the first Gaia catalogue will be also available as soon as the catalogue is made public (expected for mid 2016).

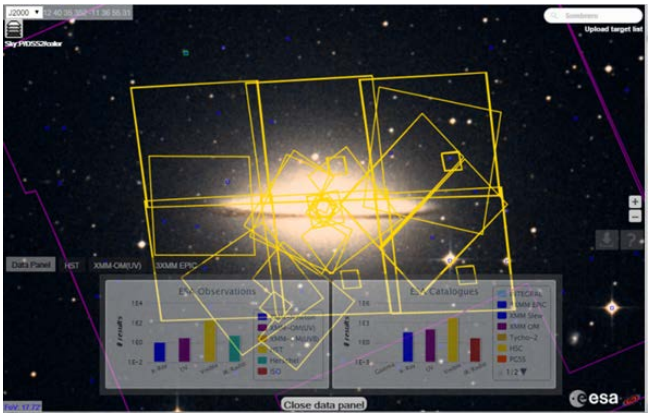

Figure 3: ESA Sky User Interface

5. INCLUSION OF NON-ESA DATA

doi: 10.2788/854791

Due to the interest from different projects and from the community, inclusion of other missions, including non-ESA ones, is also foreseen. By following some simple recipes provided by ESDC, data from other data centres can be made accessible on the ESA Sky interface. The mechanism to publish data into ESA Sky implies for the different data centres:

- Generation of HiPS maps: Taking all the images provided by the project as input, a HiPS structure is generated for exploration. This procedure is not only technical but, also, a good knowledge of the data facilitates the improvement of parameters and cuts to have the best possible all sky map. The generation of HiPS maps allows the publication of data not only inside ESA Sky, but also inside other VO applications
- Access to observational footprints: For the observational data, detailed footprints describing the spatial coverage of the observations should be generated. These footprints will be represented as an STC-S polygons [10]. Different techniques to generate these footprints are described in the ESDC how-to pages
- A table, containing general observational metadata should be generated by the data provider, fulfilling a common data model, similar to IVOA ObsCore DM
- The previous table could be exposed to the public by a TAP server by the data provider or directly provided to the ESA Sky team for ingestion inside the application data base
- The observational table should contain a link to the data located in the data provider side, so users will finally access the data from the data provider

Following these steps, data can be show in the ESA Sky interface and users can access project data directly from the data provider nodes.

Further information at:
http://www.cosmos.esa.int/web/esdc/esasky-contributing

6. FUTURE PLANS

Currently, ESA Sky has been focused on the access to images from different missions and source catalogues. A future enhancement will allow the discovery of spectral data in a similar way to the discovery of images. Spectral slits will be displayed in the interface for the different missions and a dedicated user interface module will allow the visualization and retrieval of spectral data. In the long term, there are plans to provide support for time domain data by the inclusion of time lines and time dependent HiPS maps.

At the same time, there are plans to include more missions into the ESA Sky interface in the short term and the connection to other external services.

Also, access to external VO applications will be obtained by connecting the ESA Sky interface to them using the IVOA SAMP protocol and access to other ESA archives by HTML5 interactions.

7. ACKNOWLEDGEMENTS

We acknowledge the support of the CDS team members, Pierre Fernique and Thomas Boch on the Aladin Lite component and HiPS generation. From ESA Science Operations Centers, we acknowledge Pedro Rodriguez (XMM-Newton SOC), Guillaume Belanger (Integral SOC) and Marcos Lopez-Caniego (Planck SOC) for their support on the creation of the HiPS maps for these missions.

Also we acknowledge some of the former ESA Sky team members: Iñaki Ortiz, Ignacio León, Andy Pollock and Michael Rosa.

8. REFERENCES

[1] Arviset et al, *Long-term Strategy for ESA Science Archives*, 2014ASPC..485..159A, ASP conference series, vol. 485, 2014, p.159
[2] Richardson, L., Ruby, S. *RESTful web service*, ISBN 978-0-596-52926-0 (2007)
[3] Boch, T.; Fernique, P., *Aladin Lite: Embed your Sky in the Browser*, *2014ASPC..485..277B*, ASP conference series, vol. 485, 2014, p.277
[4] Calabretta M.R., Boudewijn F.R., *Mapping on the HEALPix grid*, Mon. Not. R. Astron. Soc. 381, 865–872 (2007)
[5] Fernique, P.; Boch, T.; Pineau, F.; Oberto, *Create and Publish a Hierarchical Progressive Survey (HiPS)*, 2014ASPC..485..281F, ASP conference series, vol. 485, 2014, p.281
[6] Dowler P. et al, *IVOA Recommendation: Table Access Protocol Version 1.0*, 2011arXiv1110.0497D
[7] Dewsbury, R., *Google Web Toolkit Applications*, ISBN 978-0-321-50196-7, Prentice Hall, (2007)
[8] Louys M. et al, IVOA Recommendation: Observation Data Model Core Components and its Implementation in the Table Access Protocol Version 1.0, arXiv:1111.1758 (2011)
[9] Boch T, et al, MOC - HEALPix Multi-Order Coverage map, http://www.ivoa.net/twiki/bin/view/IVOA/IvoaApplications (2014)
[10] Rots, A., STC-S: Space-Time Coordinate (STC) Metadata Linear String Implementation, IVOA, (2007)

DATA ACCESS SERVICES AT DPCT TO SUPPORT SCIENTISTS' ONLINE AND OFFLINE DATA ANALYSIS

A. Mulone[1], *R. Morbidelli*[2], *R. Messineo*[1], *R. De March*[1], *F. Filippi*[1], *M. Vaschetto*[1], *F. Sella*[1], *S. Uzzi*[1], *C. Manetta*[1]

[1]ALTEC S.p.A Corso Marche 79, 10141 Torino, Italy http://www.altecspace.it/

[2]INAF-OATo Strada Osservatorio 20, 10025 Pino Torinese, Torino, Italy http://www.oato.inaf.it/

ABSTRACT

Received data at Gaia DPCT centre are stored in an Oracle 11G DBMS. This way, data is always online. Database size is very huge (more than 100TBs of data in 25 months) and its management is very complex. Many procedures to manage tables' partitions/subpartitions have been created and they are scheduled to run every day.

Moreover, in order to avoid wasting space, tables are compressed after data ingestion. Compression is executed after 15 days of data insertion to not affect the normal daily ingestion pipeline.

Having data online is very important duty during the operating phase, scientists can daily check and get data to execute analyses, in fact they can monitor what happens when data are received before normal daily processing phase start or scientific team can monitor data received immediately after the processing is completed if some anomalies are founded.

Thanks to online data, DPCT can perform data analyses not covered from the nominal mission pipelines. Scientists can extract data from operating database and analyze them with mission tools available (for example IDL) or a selected subset of data can be stored in a RDBMS database created on demand.

Furthermore data can be extracted and inserted into NoSQL database to lead other particular analyses. This is what has happened for TECSEL2 project.

Index Terms— Data Access, Data Management, Data Processing, RDBMS, Gaia, DPCT, Data analysis

1. INTRODUCTION

DPCT is the Italian Data Processing Centre for the Gaia Mission and its architecture has been designed considering all functional and non-functional purposes required by scientific software taking into account the requirements related to the management of a database repository. DPCT allows the Italian scientific community to execute data analyses not implemented in the software of scientific pipelines. DPCT repository will maintain a relevant part of the Gaia main database at the end of mission.

Data received are stored in two different database's groups: processing databases (LOCALDB for daily pipelines [3], [4] and GSRDB for cyclic processing [5], level 1) and the repository database (REPDB, level 0).

In LOCALDB data used and produced by daily pipeline are stored. Data ingested are pre-processed (some data are eventually discarded).

In GSRDB data used by 6-months-cycle processing are stored. Input data are taken from repository database. No direct ingestion of data coming from DPCE is foreseen.

In REPDB data are stored as they are received, so you can find all data received at DPCT from the beginning of the mission. Furthermore consolidated processing output is moved from LOCALDB to REPDB (and it is removed from LOCALDB).

2. SYSTEM ARCHITECTURE

The Data Transfer System (DTS) manages the regular data transfer between ESAC DPC (DPCE) and DPCT. It is in charge to interact with other subsystem when some data transfer occurs. This functionality allows a complete automatic pipeline.

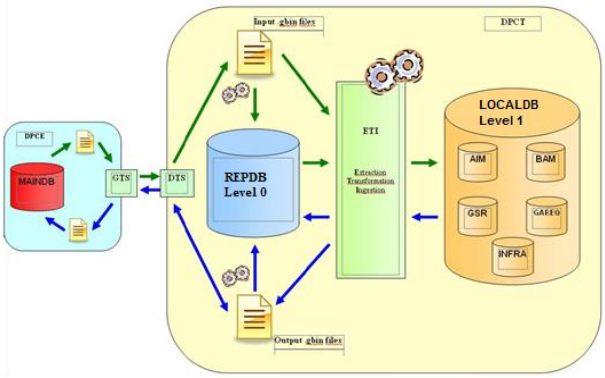

The daily transfer and the cyclic database distribution occur over the public internet connection.

When a transfer is incoming to DPCT the DTS notifies it to DPCT system that starts two or more parallel ingestions.

The component in charge to execute ingestion is the DSS (Data Storage Subsystem).

Persistent data management is the main critical point of the overall DPCT. The size of DPCT database will consist of many hundreds of terabytes which will have to be stored for a long time. In addition, data access must be efficient enough to avoid choking on the processing. The two main components for data access and data storing are the relational database and the cluster file system (full archive exceed the petabytes).

3. DATABASE ARCHITECTURE

LOCALDB (level 1) contains all data strictly tied to the processing and the infrastructure management.

The size of this database is steady because after ingested data are processed and verified by scientific and operating team, its output is moved to level 0 database and the matching input is removed.

GSRDB (level 1) contains data used by cyclic processing. The size of this database can vary since it is strictly tied to the length of the processing period. Its size is about 21TB at the moment.

In level 0 database data related to Gaia Main Database are stored. The size of this database is increased constantly; therefore this database should be managed with particular care.

Both database types (level 0 and level 1) provide advanced availability and scalability features. Oracle allows multiple computers to run Oracle DBMS software simultaneously while accessing a single database, thus providing a clustered database. The database grid uses the following Oracle products (version 11g): Oracle Server Enterprise Edition, Oracle Real Application Cluster, Oracle ASM to manage storage used by database, Oracle Compression and Oracle Partitioning.

3.1. Daily Processing Database

LOCALDB is about 15TB large and this size is steady. This database contains data used and generated by daily pipeline software and DPCT infrastructure management data.

Data are stored in different schemas: one schema for each software and one schema for DPCT infrastructure.

This database uses a lot of the partitioning features of Oracle. Tables are partitioned by range using time-based fields. Some very big tables are sub-partitioned and in this case they are partitioned by range and sub-partitioned by hash.

Hash partitioning is an Oracle feature that allows distributing uniformly rows on a fixed number of partitions/subpartitions.

Data produced by scientific software, after having been consolidated (valuated by scientific team), are moved to REPDB database and they are removed from LOCALDB. Thanks to temporal partitioning, you just need to drop partitions instead of executing many delete statements to remove data from LOCALDB. Drop partitioning is much faster than delete statements.

3.2. Cyclic Processing Database

GSRDB is about 21TB large and this size can vary depends on the cycle you must analyse. Data stored in this database are taken from repository (REPDB database) and are matched and stored in specific input tables. These specific input data are the starting point of cyclic processing (new input tables).

New input tables are partitioned by range and range is defined before processing start. You need to drop old partitions and create new partitions for each new cyclic processing.

Output cyclic processing tables are partitioned by range too, but there are exceptions: tables with just a few data are partitioned by hash or they are not partitioned at all.

After cyclic processing is terminated and validated by scientific team, its output is moved to repository and GSRDB is cleaned.

3.3. Repository Database

Repository database is the most interesting one, in fact many ‘ad hoc’ procedures have been implemented to manage tables partitioning/subpartitioning and to address database size problem.

As mentioned above, this database grows continuously and its size is more than 100TBs after 25 months. To allow ingestion, extraction and queries on stored data, the partitioning criteria have been chosen accurately. We must maintain data of different months in different datafiles[1] per each table. This is necessary to manage the backup of such huge databases and to save space.

Partitions/subpartitions are created every day as long as data are received. All tables were created only with one partition and some fixed subpartitions. To add new partitions a set of procedures have been created and they are launched every day. When you execute a procedure, table last partition is split in two so you have a partition with data

[1] Datafiles are physical files of the operating system that store the data of all logical structures in database. Datafiles are logically grouped in tablespaces.

doi: 10.2788/854791

received before procedure start and a new empty partition with data you will receive later. During this operation you can move new empty partitions on different tablespaces.

Thanks to the partitioning/subpartitioning, we can reduce the database effort in data reading, split partition operations, index reconstructions, along with the reduction of contention during ingestion. All tables are partitioned on 'solutionid' field, a numeric field assigned by DPCE identifying the processing run. Subpartitioning field is different from table to table because different tables map different objects and not all the fields are common to all tables.

Using 'solutionid' as partitioning key, you can execute the retraction of data. Sometimes, data you had received are wrong processed at DPCE so you have to remove them from repository and insert the correct ones. DPCE refers to data to remove using 'solutionid'.

Oracle does not have a defragmentation mechanism so you may encounter some space issues when inserting data in parallel in a table. Moreover Oracle offers the possibility to insert data on compressed table to save space. This feature produces drawback to slow down the ingestion. For example the average time spent to ingest a file in compressed tables found in June was about 55s instead the ingestion of a file in not compressed tables is about 20s.

To address these problems another set of procedures have been developed at DPCT. Procedures move partitions/subpartitions from a month's tablespace to another new empty month's tablespace and these table structures are compressed during the moving. These procedures run on data that were stored at least 15 days before. For example data of May 2015 occupied 2795950MB before space recover; but after space recovering, the space occupied was 2213061MB, 500GBs less than before.

Managing these problems, we can have all data online and accessible. Data stored in REPDB are available for scientific analysis (though DAAS subsystem) in every moment. You can scan a table with more than 40 billion of observations and check the magnitude in less than 20 minutes.

4. DATA ACCESS AND ANALYSIS SUBSYSTEM

DAAS provides to scientists mechanisms to perform specific data analysis in support of Gaia mission assigned to DPCT. DAAS provide services for both basic data analysis, based on standard numerical and statistical libraries and processing oriented data analysis. It is composed by a set of tools accessible through a web interface.

One of the most important features of DAAS is allowing creating packets for scientific analysis. These packets contain input data used by scientific software and they are saved in gbin file format. With these packets scientists can execute some particular processing or they can execute some bug fixing or 'ad hoc' analysis without impacting the operational environment and they can use their preferred developing platform. There are different kinds of packets, depending on the scientific software module.

Scientists can execute packets generations they need autonomously without operator's intervention.

Gbin is a java serializable file and it is the most used format in all Gaia project. Data are received in gbin format at DPCT.

If you want to extract some data as it was received by DPCT without inserting it in a packet, this is obviously possible but in this case the scientist has to create a 'data request' because an operator must check the feasibility of the request. In fact too many data can be requested or some particular 'data request' needs creation of a specific index to be completed.

Inside the data request the scientist must specify: interested data, ranges of data and related fields, output format and periodicity of request. Scientists can request different output format: text file, gbin and tabular format (you can extract data and put it on a demand database that can be created from a snapshot of an operating database or recreating a subset of an operating database or extracting and reimporting data from a database to a custom database).

In the last weeks, Oracle 12c multitenancy was tested. That is a new feature introduced with Oracle 12c. Through multitenancy users can move a database across two different sites maintaining data, permissions and logical structures. This is a very interesting feature in an on-demand database.

Nominal data processing produces output data, stored in databases, and a lot of plots accessible via DAAS using a web interface. Scientists can select a date from DAAS interface and they can view data produced by workflows executed in that date and they can download plots if needed. Furthermore scientists can download reports produced at the end of each processing run.

Another important tool is the MDBExplorer developed in Gaia Project. Using this tool you can consult Gaia main remote databases hosted in Madrid, local databases at DPCT, and you can open and inspect gbin files. MDBExplorer has the capability to decode and read all binary fields of an object (it is not important if the object is stored in file or in database). No other sql client can inspect the content of a binary field. Using this tool you can extract data and save them in gbin file.

Today IDL is the most relevant COTS software selected for data analysis at DPCT.

DAAS offers an IDL library making the Gaia tools available inside any IDL scripts, which are particularly useful for scientists to analyse data.

IDL library permits the scientist to communicate with DPCT datastores, allowing IDL programmer to deserialize binary data into collection of IDL primitive types, which can

be analysed, plotted and collected by the powerful IDL libraries.

More generally the library bridges Java to IDL thus allowing the usage of any Java Gaia library into IDL.

At DPCT we have also developed TECSEL2 project, which has the aim to correlate Gaia data to some public solar datasets [2]. To analyse this relation we have exported data from Oracle databases and put it in Cassandra (a noSql database). This is an interesting technique to exploit Gaia data. To put data in Cassandra we used an ETL (Extract, Transform, Load) tool called Talend Open Studio that allows to insert data in the proper Cassandra structures (or it creates Cassandra structures if they are missing) and to even execute some custom transformations during data movements from Gaia database to Cassandra.

5. CONCLUSIONS

In this paper we describe how data received are stored and managed at DPCT. We store all data received in both relational database and file system. Data stored in RDBMS are always available for scientists that can perform analyses. We describe tools we provide to DPCT scientific team and how these tools can interact with users and data.

Moreover we explain how huge data are managed through Oracle 11G database and we highlight the possibility that new Oracle 12c introduces for exploiting and for spreading data across multiples centres.

In closing we touch upon how data can be analysed in noSql database that it is what happens in TECSEL2 project.

6. ACKNOWLEDGMENTS

The authors are members of the Gaia Data Processing and Analysis Consortium (DPAC) and this work has been supported by the ASI contracts 2013/030/I.0 and I/058/10/0 (Prof. M. Lattanzi, INAF-OATo, is the Italian Principal Investigator).

7. REFERENCES

[1] Messineo, R., Morbidelli, R., Martino, M., Pigozzi, E., Mulone, A. F., and Vecchiato, A., “Italian dpc infrastructure and operations to support of the italian participation to the gaia data processing” in [Space Telescopes and Instrumentation 2012], Proc. SPIE 8451-13 (2012).

[2] De March R., Busonero D., Messineo R., Bemporad A., Vaccarino F., Mulone A.F., Fonti A., and Lattanzi M.G., “TEmporal Characterization of the remote SEnsors response to radiation damage in L2”, Proc. Big Data from Space (BiDS16), IEEE, 2016.

[3] Riva, A., Gai, M., Lattanzi, M. G., Russo, F., Buzzi R., “BAM: A metrology device for a high precision astrometric mission”, RMXAC, Vol. 45, 35 (2014).

[4] Busonero, Deborah, “Astrometric instrument modeling in the context of Gaia astrometric verification: tasks and activities”, in [Space Telescopes and Instrumentation 2012], Proc. SPIE, Vol. 8449, 84490F (2012).

[5] Alberto Vecchiato, Ummi Abbas, Marilena Bandieramonte, Ugo Becciani, Luca Bianchi, Beatrice Bucciarelli, Deborah Busonero, Mario G. Lattanzi, and Rosario Messineo, “The global sphere reconstruction for the Gaia mission in the Astrometric Verification Unit”, *Proc. SPIE* 8451, Software and Cyberinfrastructure for Astronomy II, 84513C (2012).

INTEROPERABILITY OF MULTISCALE VISUAL REPRESENTATIONS FOR SATELLITE IMAGE BIG DATA

François Merciol, Antoine Sauray and Sébastien Lefèvre

Université Bretagne Sud – IRISA
Campus de Tohannic, BP 573, 56017 Vannes Cedex, France

ABSTRACT

In this paper, we propose an interoperable solution for dealing with hierarchical representations of satellite images. Computationally intensive construction of the tree representation is performed on the server side, while the need for computational resources on the client side are greatly reduced (tree postprocessing or visualization). The proposed scheme is interoperable in the sense that it does not impose any constraints on the client environment, and API for C++, Java and Python languages are currently available. The communication is performed nodewise in a binary format using array structure for limiting the memory footprint.

Index Terms— Hierarchical image representation, Tree, Client-server framework, Interoperability.

1. INTRODUCTION

Recent years have seen Earth Observation entering in the Big Data era. This brings new challenges related to the V's: volume, variety, velocity, as well as veracity and visibility. Very high spatial resolution (e.g., a Pleiades scene is $40,000 \times 40,000$ pixels) as well as high temporal revisiting frequency (e.g., Sentinel missions will offer worldwide updates of a scene every 5 days) lead to massive datasets to be processed in order to produce land cover map, detect objects of interest or changes, retrieve specific patterns, or perform (manual or assisted) visual interpretation. To do so, very efficient solutions have to be designed, requiring algorithms of appropriate complexity relying on efficient data structures and operating on adequate architectures (e.g., cloud, HPC, etc). Representing images through multiscale representations based on tree structures has been proved to be a relevant framework for efficient processing of large image data [1]. Besides, trees have been successful in addressing various tasks in remote sensing [2, 3, 4].

However, one of the major bottlenecks of such a hierarchical image analysis strategy is the need for prior computation of the tree representation. The computational cost of building a tree from a large dataset is indeed significantly high, even in the case of efficient algorithms [5]. In this paper, we propose a new strategy based on a client-server model, where the tree is built on a server before being used by the client. Such a strategy presents many advantages: extensive (and possibly scalable) computational resources are needed only on the server side, since the client will only process the tree structure (requiring much less resources than dealing with raw data); similarly to computing resources, network resources are required on the server side between the data source and the processing server, not on the client side; if the satellite image comes with access restrictions, which is the case for most commercial satellites that do not allow distributing the raw data, it is possible to provide a tree-based product (non bijective image representation) to the client; any compatible client can be used, even running in some environments not known for their performance; finally, if the client is only interested by some part of the image (and not the full data), only a subset of the tree can be sent by the server.

The authors acknowledge the support of the French Agence Nationale de la Recherche (ANR) under reference ANR-13-JS02-0005-01 (Asterix project).

2. INTEROPERABILITY

The client-server strategy is designed here in such a way that there is no need for a unique environment (software, coding language, operating system) on both sides. Interoperability is ensured through several wrappers allowing to access the API from the client side using various environments (C++, Java, Python, etc.). Any postprocessing or visualization tool can then be used on the client side. Let us note that the proposed solution can also be used on the client side only allowing different tools for tree construction and manipulation. The only requirement is a standard scheme to support the distribution of the (full or partial) tree from the server to the client.

There are many interoperable formats, the most famous being XML. However, exchanging large datasets using a verbose language such as XML is not efficient. So we have rather used byte arrays containing scalar values (either integers or real numbers). It requires some appropriate interface between client and server components (i.e. to use the same coding format: number of bits, byte ordering). To illustrate the interoperability offered by our framework, we provide 3 code snippets for C++, Java and Python in Fig. 1. We can observe that the calls are very similar from one language to the other.

```
// C++ snippet
CppSession<Component, Leaf> *cppSession
  = CppSession<Component, Leaf>::getInstance
    (serverUrl, imageUrl);
cppSession->buildTree (algo, metric);
cppSession->createLeaves ();
cppSession->copyTree ();

// Java snippet
Session s = Session.getInstance (urlServer, urlImage);
s.buildTree (algo, metric);
s.createLeaves ();
root = s.copyTree (connectedComponentCreator);

// Python snippet
s1 = PythonSession (urlServer, path)
s1.buildTree (algo, metric)
s1.createLeaves (leafCreator)
s1.copyTree (connectedComponentCreator)
```

Fig. 1. Client requests for tree construction in C++, Java and Python.

3. IMPLEMENTATION

For the sake of performances, the server component has been coded in C++. It benefits from significant optimizations with no overhead for tree coding (only the minimal number of quartets needed is required). API for tree communication with C++, Java, and Python-compliant clients have been designed. The system is made freely available to the scientific community[1].

The proposed framework is illustrated through an example (Fig. 3). We assume here that nodes contain scalar data, but the solution can be easily adapted to more advanced representations.

First, a given client has to choose the appropriate server based on its needs. Indeed, a pool of servers might be available, each server coming with its specific algorithms (kind of tree, kind of data to be processed, etc). The client establishes a connection to a server through a Kernel object. Once the communication link has been established, messages are exchanged between the client and the server through a dedicated protocol. To do so, it relies on the interface given in Fig. 2. Let us observe that the communication is here of type unconnected socket (e.g. with one TCP connection by HTTP request). This protocol relies on the BOOST library, both for socket management and vector serialization.

```
class UTP {
public:
  static UTP *getInstance (const string &url);

  int openSession (const string &url);
  void closeSession (int sessionId);
  void closeAllSession ();

  void buildTree (int sessionId, int tree, int weight);
  vector<int> getInfos (int sessionId);
  double getMillisec (int sessionId);
  vector<float> getValues (int sessionId);
  vector<int> getNodeChildren (int sessionId, int id);
  vector<int> getNodePixels (int sessionId, int id);
 //
};
```

Fig. 2. Interface available to the client (C++ example).

Once the Kernel object has been created, the input image has to be loaded on the server, resulting in a new Session object. The protocol for iterative communication of array structures is actually implemented within this Session object. As shown in the Snippets from Fig. 1, the creation of both Kernel and Session objects are embedded in a single function call. Depending on the application context, various data sources are to be considered. Indeed, the client might already store the image to process. In this case, the image is to be sent to the server over the network, before receiving back its tree representation. More interestingly, the server might already store the image data (e.g. in some business or research environments). As such, there is no network cost to transfer the original image. A last and more realistic scenario, especially when considering the current initiatives such as the EU Copernicus programme, considers that original data are actually made available through dedicated geospatial hubs (such as the one from ESA or from national space agencies). In this last case, we can reasonably assume that there exist a fast network connection between the server and the data hub.

Once the image is stored on the server, the latter computes its tree representation. To do so we assume that some efficient tree construction algorithms [5, 6] are available on the server. Besides the tree, several data structures are stored: the list of nodes in the image (array with 10 cells in the figure), the list of inner nodes or nodes with children (array with 4 cells, noted C*), the list of leaves or nodes with pixels (array with 8 cells, noted L*). These additional structures increase the overall storage cost on the server side. But let us recall that once the tree has been built, there is no need to store the initial image on the server anymore. Indeed, the tree can then be used at any time to answer a client request (both for the full image or only a part of it). If the tree is used in a single user session, the server will store the tree and the additional data structures until the end of the user session. If the image data shall remain available (through its tree representation) for future users, then all structures but the original data are kept.

After tree construction, the server sends a first array providing the image height and width (3×3), the number of nodes in the tree (10), the number of inner nodes or nodes with children (4), and the number of leaves or nodes with pixels (8). The client is then able to allocate the right amount of memory and initialize the local data structure (root, nodes, pixels). On the client side, an iterator is used to request the server to send information for each node of the tree. Inner nodes are considered first ($C0$ to $C3$). The server then send iteratively the information for each C_i through an array containing the ID of the node followed by the ID of its children (e.g., node 0 has children 1, 5, and 2). Once all inner nodes have been transferred, the client iterates on leaves ($L0$ to $L7$). It thus requests the server to provide information that consists for each node of an array containing the ID of the node followed by the ID of the pixels (e.g., node 2 points to pixel 3). The (full or partial) tree is then available on the client side.

4. EXPERIMENTS

It is rather difficult to assess the impact of network performances on the overall system. Thus we have decided to run the client and the server on the same computer to avoid being affected by unpredictable network behavior. Furthermore, we have also discard the effect of image loading over the network by storing the images directly on the server. We report the results obtained on a standard workstation (8 GB RAM, 4 cores Intel i5-3320M@2.60GHz).

We have compared between several configurations. The server is running in C++, and various clients are used: C++, Java and Python. For reference we have also included a Java standalone application (not client-server). We have considered various satellite image size (from 80,000 to 13,000,000 pixels) and types (panchromatic, multispectral), and use the notation $X \times N$ to denote an image containing X pixels with N spectral bands. Results are given in Fig. 1. We use here the α-tree implementation from Havel et al [5].

We can derive several observations from these experimental results. First, when small images are considered, the proposed client-server framework is not mandatory. Indeed, a standalone application is not too costly in this case, and can achieve a similar level of performance than the client-server mode (e.g. very similar times obtained for the full Java and C++/Java mode). Even on such small images, we can notice the strong impact of the core algorithm (e.g. 0.11 for the C++ construction, while 0.40 for the Java one). When image size increases, computation cost becomes an issue and the proposed framework achieves satisfying performances. Indeed, it can benefit from a significant optimization on the server side (i.e., compare between the full Java and C++ server). Furthermore, the cost for array communication and tree reconstruction on the client side is reasonable w.r.t. the overall cost. This allows the proposed system to outperform a standalone application for both C++ and Java clients. Let us remark that if a more powerful server would have been considered (and not run on the same computer than the client), the improvement would have been much higher. For a given client, the overall cost

[1] See http://www.irisa.fr/obelix for more details.

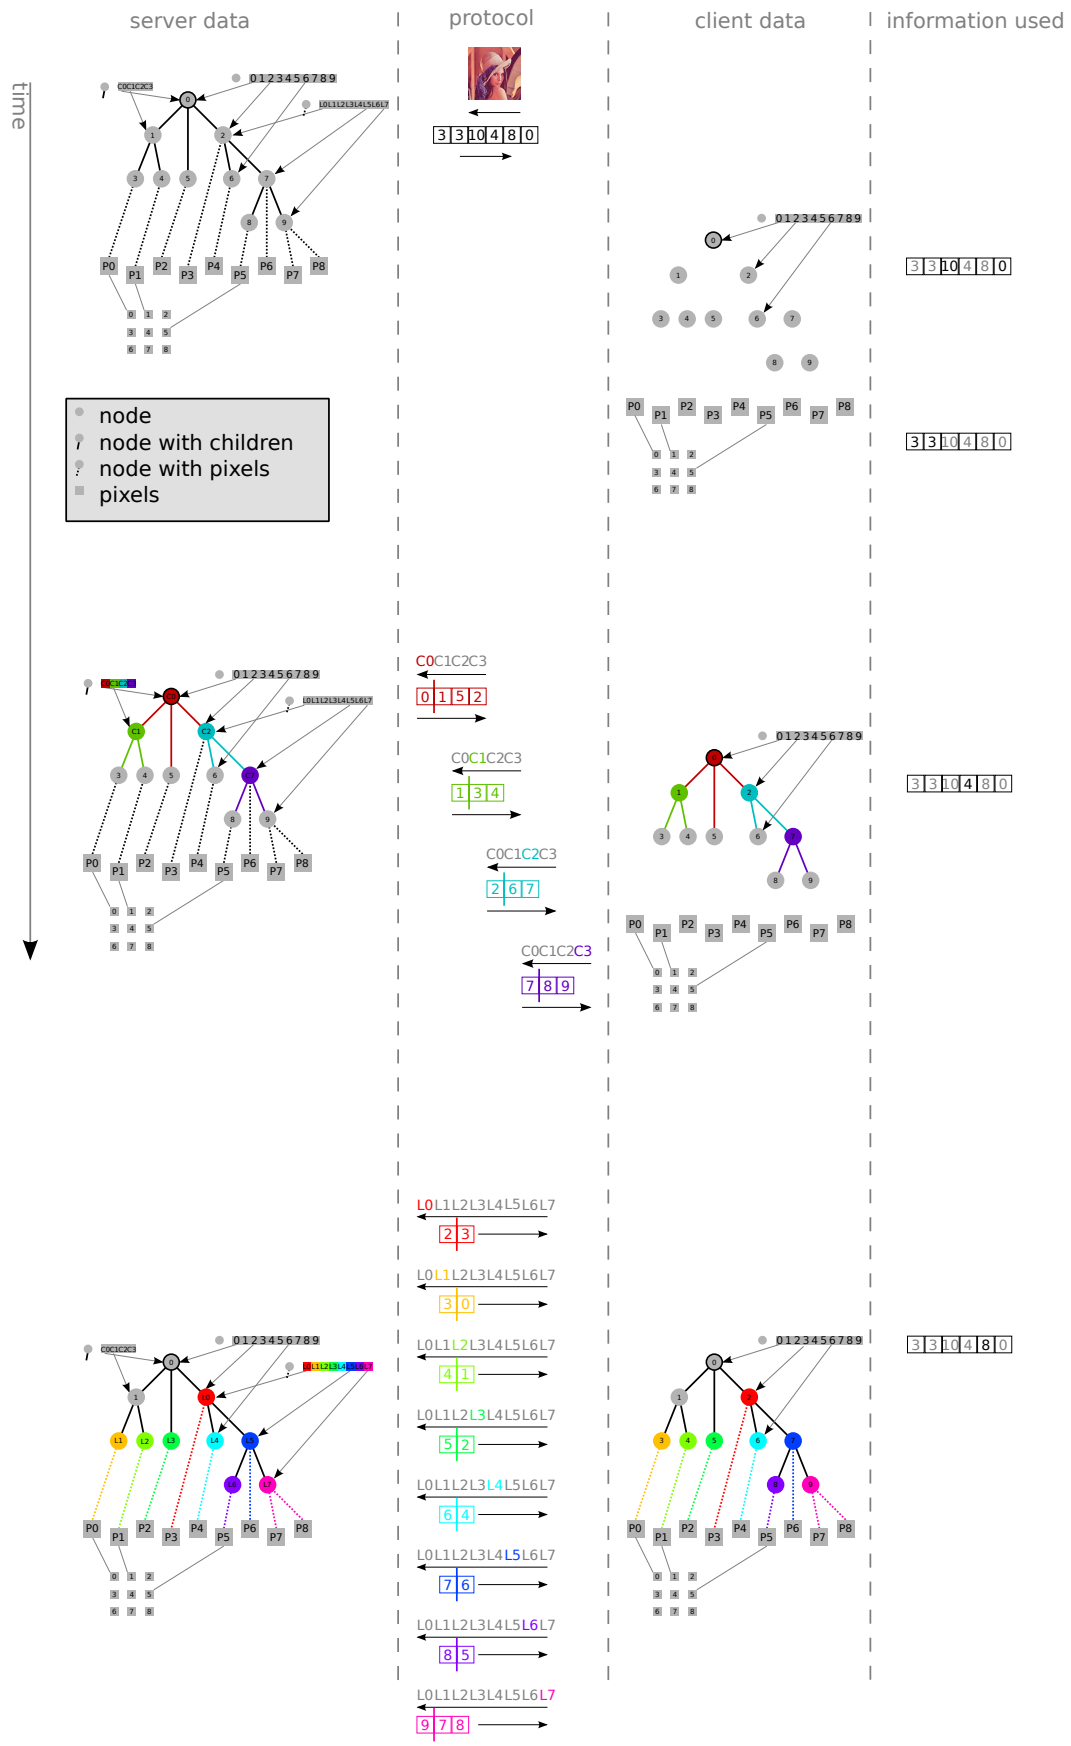

Fig. 3. Illustration of the communication process for tree transfer.

doi: 10.2788/854791

Method	$80k \times 4$	$200k \times 4$	$8M \times 1$	$13M \times 1$
full Java	0.40	0.85	22.11	—
C++ server	0.11	0.31	15.57	26.36
C++ client	0.05	0.13	4.26	7.05
C++/C++	**0.16**	**0.44**	**19.83**	**33.41**
Java client	0.31	0.31	6.84	19.87
C++/Java	**0.42**	**0.62**	**22.41**	**46.23**
Python client	1.03	2.88	99.64	—
C++/Python	1.14	3.19	115.21	—

Table 1. Comparison between standalone computation and the proposed client-server framework with various clients.

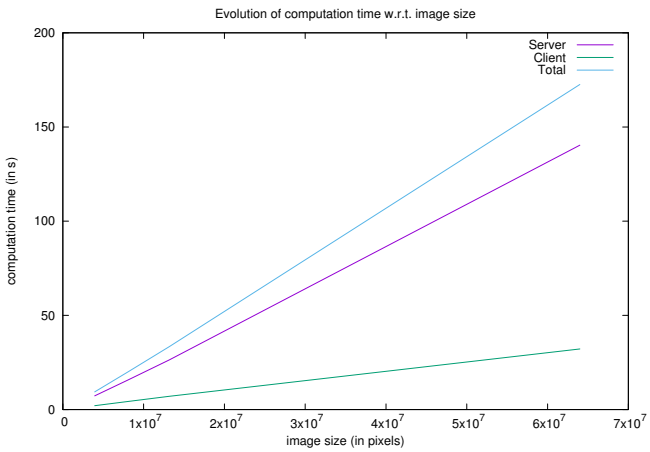

Fig. 4. Quantitative evaluation of the most efficient architecture (both server and client are in C++).

then tends to the client cost only.

Furthermore, we can notice some significant differences between the clients. Indeed, each client comes with its own memory management system, and its performance has a direct influence on the cost needed to create data structures on the client side. As such, we can see that Python might not be able to address very large images, even in this client-server mode. But let us note that for the largest image in consideration, the full Java application also fails due to memory limitations. Only the proposed framework, both with C++ and Java clients, is able to address the large image.

Finally, we have measured the computation time for both C++ server and client, for another set of satellite images with increasing size. We can see in Fig. 4 that the overall process has a linear complexity, and the client cost is always much lower than the server one. This shows the relevance of the proposed framework.

5. DISCUSSION

The proposed solution has appealing properties in terms of memory cost. As already indicated, a naive approach would have been to perform a full transfer of a tree through an XML representation. But the textual representation, as well as the data embedding through XML tags, would have led to a prohibitive cost. Byte coding appears as a relevant alternative. It is indeed possible to consider coding the tree in a proprietary binary format. Unfortunately, there is no native serialization scheme in all languages, and worse no interoperability between the existing serialization formats. Without serialization, byte coding can still be implemented to ensure the communication process between the client and the server. But this would also require an additional cost if interoperability is sought. Indeed, in this case the end-user aims to manipulate the tree in an environment different from the one on the server side. Thus a new tree structure is needed on the client side to map the tree sent from the server in a format compliant with the client environment. The proposed scheme (see Fig. 3) introduces node-based messages. Such messages have a size that is negligible w.r.t. the tree size, and can be removed as soon as they have been used on the client side to update the tree reconstruction. Since a tree representation might come with a very large memory footprint (e.g. in case of large images), the proposed strategy allows allocating most of the client memory to the storage of the tree with no significant overhead due to the server communication.

We have considered a complete transfer of the tree from the server to the client. If the latter is interested in only a part of the image, two different strategies can be followed (top-down and bottom-up). On the one side, the user can select a subtree, but it requires to identify the node of the tree acting as the subtree before transferring the subtree. While this second step is easily achievable with the framework proposed in this paper, performing node selection on a client-server paradigm is not trivial. On the other side, the selection can be done in a bottom-up way. In this case, the user defines the area of interest directly in the image, from which the related subtree (or forest, i.e. set of (sub)trees) is extracted. Among the available solutions for allowing such user selection, we can rely on our previous work [7] where the user input consists in a bounding box. The selected tree is then the largest subtree (or set of subtrees) totally included in the bounding box.

Future work will include demonstration of this framework within realistic remote sensing scenarios, as well as enrichment of the server component with additional tree construction algorithms.

6. REFERENCES

[1] G.K. Ouzounis, V. Syrris, L. Gueguen, and P. Soille, "The switchboard platform for interactive image information mining," in *Image Information Mining Conference*, 2012.

[2] S. Lefèvre, L. Chapel, and F. Merciol, "Hyperspectral image classification from multiscale description with constrained connectivity and metric learning," in *IEEE Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 2014.

[3] F. Merciol and S. Lefèvre, "Fast building extraction by multiscale analysis of digital surface models," in *IEEE International Geoscience and Remote Sensing Symposium*, 2015, pp. 553–556.

[4] E. Aptoula, M. Dalla Mura, and S. Lefevre, "Vector attribute profiles for hyperspectral image classification," *IEEE Transactions on Geoscience and Remote Sensing*, 2016, to appear.

[5] J. Havel, F. Merciol, and S. Lefèvre, "Efficient schemes for computing α-tree representations," in *International Symposium on Mathematical Morphology*, 2013, pp. 111–122.

[6] E. Carlinet and T. Geraud, "A comparative review of component tree computation algorithms," *IEEE Transactions on Image Processing*, vol. 23, no. 9, pp. 3885–3895, 2014.

[7] F. Merciol and S. Lefèvre, "Buffering hierarchical representation of color video streams for interactive object selection," in *International Conference on Advanced Concepts for Intelligent Vision Systems*, 2015, pp. 864–875.

AUTOMATED DTM GENERATION AND SUPER-RESOLUTION RESTORATION FROM NASA MRO CAMERAS AND IN FUTURE FROM EMTGO16 CASSIS

Y. Tao, P. Sidiropoulos and J.-P. Muller

Imaging group, Mullard Space Science Laboratory, Department of Space & Climate Physics, University College London, Holmbury St Mary, Surrey, RH5 6NT, UK
(yu.tao@ucl.ac.uk; p.sidiropoulos@ucl.ac.uk; j.muller@ucl.ac.uk)

ABSTRACT

An automated DTM production pipeline for planetary mapping and a separate one for super-resolution restoration algorithm/pipeline are introduced. The objectives are to provide fully automated processing chains for very large – scale production of many Tbytes of MRO CTX and HiRISE imaging data and explore how to best to exploit existing planetary data volumes, as well as use SRR for future missions.

Index Terms: *HiRISE, CTX, stereo, automated, DTM, Super-resolution, CaSSiS, Gotcha*

1. INTRODUCTION

Understanding the role of different planetary surface formation processes within our Solar System is one of the fundamental goals of planetary science research. There has been a revolution in planetary surface observations over the last 15 years, especially in 3D imaging of surface shape. This has led to the ability to be able to overlay different epochs back to the mid 1970s, to examine time-varying changes, such as the recent discovery of mass movement, tracking inter-year seasonal changes and looking for occurrences of fresh craters.

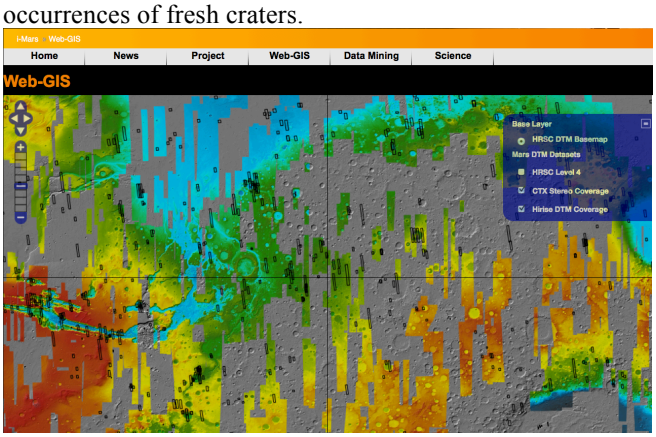

Figure 1 i-Mars.eu/web-GIS showing coverage of HRSC DTMs and corresponding CTX and HiRISE stereo coverage

To track changes on the planet Mars, it is important to be able to process data from different sensors and therefore address issues of processing large datasets with different image resolution, lighting conditions, coverage and locational accuracy. The goal of this work is to be able to maximize the exploitation of the available planetary data volumes to enable improved capability for geological and geomorphological study through the (a) generation of high quality co-registered digital terrain models using data from different instruments; (b) generate 2-5x super-resolution resolved texture images for near rover-scale feature analysis.

We introduce here recent results from a wide range of research and development activities, achieved within the EU FP-7 iMars and PRoViDE projects, in the area of automated 3D reconstruction/DTM generation and super-resolution restoration, both based on the use of the 5th generation of an adaptive least squares correlation and region growing matcher, called Gotcha [1]. Gotcha provides accurate and robust sub-pixel tie-points, which can be used in stereo disparity calculation and multi-angle super-resolution motion estimation. The experiments focus on the NASA MRO CTX and HiRISE instruments but could be extended in future ESA Mars missions such as ExoMars TGO 2016.

2. AUTOMATED DTM PIPELINE

The automated DTM processing chain, called Co-registered Ames Stereo Pipeline [2] with Gotcha refinement and Optimization (CASP-GO), takes ISIS formatted "left" and "right" MRO images (HiRISE or CTX) and reference HRSC orthorectified images (ORIs) as inputs. Using a combination of the UCL-Gotcha and NASA Ames Stereo Pipeline (ASP) ORIs and DTMs are generated which are co-registered to HRSC (and thence to a MOLA reference frame). The complete workflow has 10 steps (Fig. 1): (a) ASP "left" and "right" image pre-processing (image normalisation, LoG filtering, pre-alignment); (b) ASP disparity map initialisation (pyramid cross-correlation and building a rough disparity map); (c) Maximum likelihood sub-pixel refinement and building of a sub-pixel initial disparity map; (d) ASP sub-pixel correlation; (e) Rejection of mis-matched disparity values and erode matching gaps; (f) ALSC sub-

doi: 10.2788/854791

pixel refinement; (g) Gotcha (ALSC with region growing) densification of disparity maps; (h) Co-kriging grid-point interpolation to generate ORI and DTM as well as height uncertainties for each DTM point; (i) ORI co-registration/geocoding with reference to HRSC orthoimage and DTM adjustment; (j) Generation of Object Point Cloud (OPC) for 3D real-time visualisation on GPU using Pro3D®, specifically for HiRISE products.

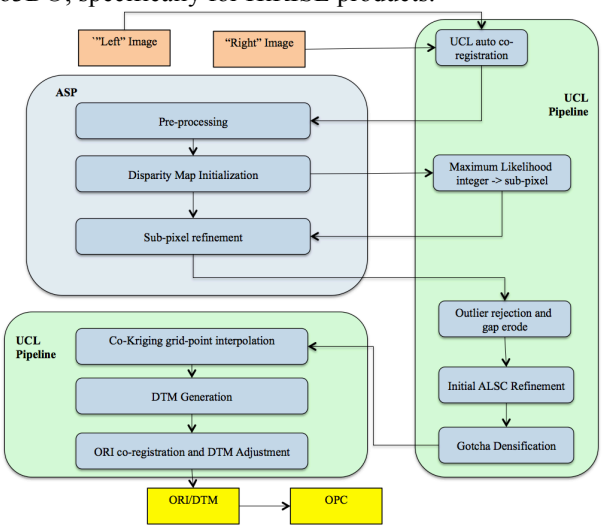

Figure 2 Flow diagram of CASP-GO auto DTM pipeline.

In CASP-GO, we have developed a fully automated processing system with improved performance compared to the original ASP system: (a) Co-registered geo-spatial coordinates w.r.t HRSC (and MOLA) data; (b) Improved DTM completeness for unmatched areas (see Fig. 2); (c) Reduced DTM artefacts; (d) Improved DTM accuracy; (e) Fully documented uncertainty for every height.

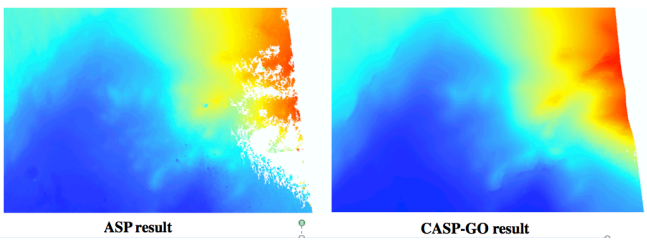

Figure 3 Example of MSL DTMs showing unmatched gaps, which have been matched/reconstructed using CASP-GO

The CASP-GO processing system has been initially applied to MER, MSL, and MC11-E CTX stereo images and is being streamlined to be applied to a large fraction of the HiRISE and CTX stereo pairs (≈3,000 pairs each) in the iMars project. (Fig. 3) shows an example of a colourised hill-shaded DTM and orthorectified image in 3D over Gale crater (MSL). These products will all be publicly available.

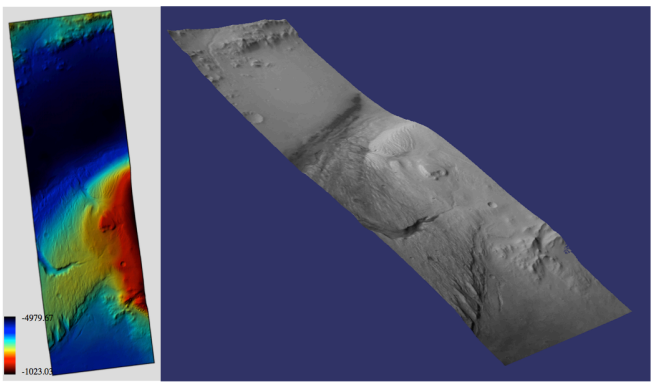

Figure 4 Example of CTX DTM and ORI (in 3D) over MSL site.

3. GPT SUPER-RESOLUTION RESTORATION

Higher spatial resolution imaging data is always desirable to the international community of planetary scientists interested in improving understanding of surface formation processes. For example, studying an area on Mars using 12m panchromatic HRSC allows you to be able to visualise the "big picture" whilst 6m CTX images allows you to see important mineralogical and geomorphological information which you can't see in HRSC and finally for a tiny percentage of the Martian surface (≈1%), 25cm HiRISE allows you to see details of surface features such as fine-scale layering. However, 25cm is not high enough resolution to be able to view individual features with diameters less than 0.75m or see the types of sedimentary features that MSL Curiosity has found in rover-based imagery.

We have also developed a novel super-resolution restoration (SRR) algorithm/pipeline, called Gotcha-PDE-TV (GPT), to be able to restore higher resolution images from the non-redundant sub-pixel information contained in multiple lower resolution raw remotely sensed images (3, 4). The GPT-SRR technique was developed in the PRoViDE project (3) to obtain improved scientific understanding of the Martian surface using a fusion of orbital and rover imagery in order to be able to provide better support in future for several critical engineering rover operations, such as landing site selection, path planning, and optical rover localisation. The technique is unique, since (a) we not only use sub-pixel information from small translational shifts but also restore pixels onto an orthorectified grid from different (comparably large) viewing angles, and are therefore able to achieve a 2-5x enhancement in resolution; (b) we use a novel segmentation-based approach to restore different features separately; (c) apply a state of the art Gotcha matcher and PDE-TV regularization to provide accurate and robust (noise resistant) restoration. GPT-SRR is applicable

whenever there exist sub-pixel differences and there are comparably large view zenith angle differences, which is always the case in orbital images, even between multiple image acquisitions taken at different times with different solar illumination conditions. Each view is subject to different atmospheric blurring and scattering but as long as the atmospheric transparency is sufficiently high, Gotcha-PDE-TV SRR can be applied.

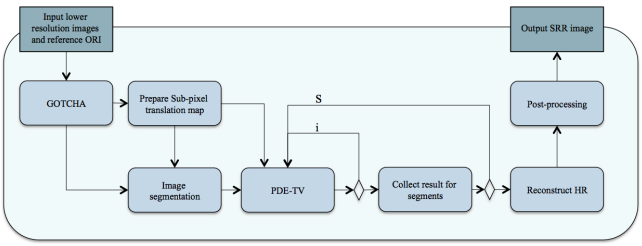

Figure 5 Flow diagram of GPT SRR processing chain.

The GPT SRR technique has been initially applied to a stack of 8 MER-A HiRISE images at the Homeplate area but later on applied to image areas comprising the entire Mars rover traverse for MER-A, MER-B and MSL providing an enhancement factor of 2-5x.

We have processed a rock size distribution study using a similar method to [5] and compared these distributions with what we can see and measure automatically on 25cm HiRISE, 6.25cm SRR image and 0.5cm (down-sampled to 2.5cm) Navcam mosaics (JPL vertical projected RDR).

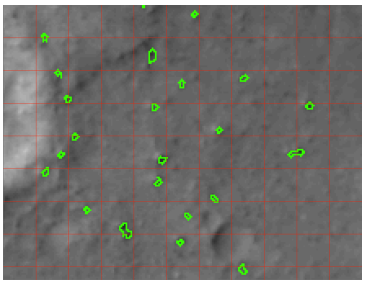

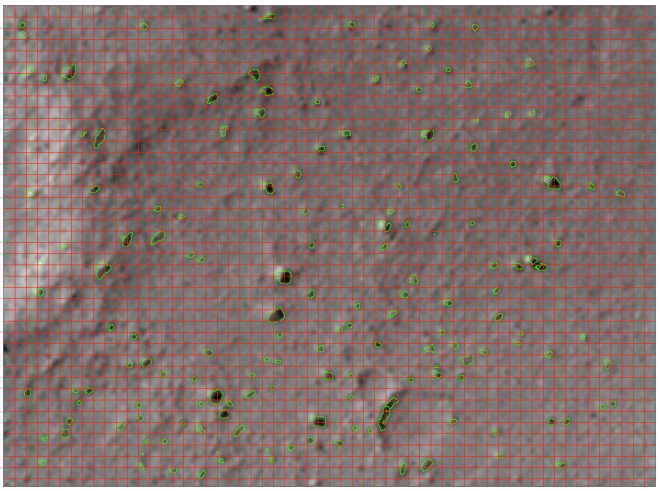

Figure 6 Automatically detected rocks (labelled green) of 25cm HiRISE image (PSP_001513_1655) (top) with 20pixel grid (5m) and 5cm SRR image (bottom) with 20pixel grids (1m) around an impact crater close to MER-A traverse at ~ (175.51045º, -14.58461º).

Table 1 Accumulated number of rocks in HiRISE and SRR image.

Rock diameter (D)	Accumulated number of rocks in HiRISE	Accumulated number of rocks in SRR image
D≥150cm	22	33
D≥50cm	23	144
D≥30cm	23	153

As shown in (Table. 1), for rocks with diameters larger than 150cm, there are 22 rocks detected from the original HiRISE image and only 1 rock detected in the range 50cm<diameter<150cm. On the other hand, in the SRR image, there were 33 rocks with diameter larger than 150cm, 111 rocks with 50cm<diameter<150cm, and 9 rocks with 30cm<diameter<50cm. This experiment has demonstrated that there are a large number of rocks, which are not clear enough for either automated detection/classification or manual measurement in the original HiRISE image. However, with SRR, a much greater number of rocks can be detected and therefore provide stronger evidence to support an application such as the selection of a future landing site.

4. FUTURE WORK

For the automated CASP-GO DTM pipeline, we are currently about to start processing a large fraction of the CTX and HiRISE stereo pairs accumulated to date in a matter of several months on a large linux cluster of 224 cores. We also aim to process all available image datasets in

doi: 10.2788/854791

future where we have repeat multi-view imagery starting with HiRISE first and then apply these SRR techniques to CTX after porting the GPT SRR software onto a GPU. We plan to develop the capability for the ExoMars Trace Gas Orbiter 2016 CaSSiS instrument (from 4m up to ≤1m/pixel) including both 3D and SRR images from multiple overlapping colour stereo views.

5. HANDLING BIG DATA

At MSSL, we have mirrored the archives of HRSC, CTX, HiRISE, as well as MER, MSL EDR and RDR data volumes in a local shared storage system in order to speed up the production process with an option such that if data is unreachable it can be read from the original source again. On the other hand, our developed software is installed in a shared directory, which is accessible from 14 Linux processing blades (16 cores/blades for 10 blades with 48GB and 24 cores/blade for 4 blades with 96GB RAM). Jobs are controlled from a local desktop machine and distributed to the 14 processing blades with multiple sessions of multi-threaded processing. Processed results are stored in several 1TB RAID storage partitions and logged back to the local controlling desktop. Failed jobs can be examined through detailed log filse and in the future will be reprocessed automatically with different processing parameters. Once all jobs are done, the resulting products are fetched by a PostgreSQL database and displayed in an interactive web-GIS system, called PRoGIS 2.0 installed in a map server at MSSL. PRoGIS 2.0 is able to call Pro3D® (the 3D rendering/analyzing software developed by VRVis in PRoViDE project), which is installed in a GPU server at MSSL. The GPU has a K-1 GRID and a K-2 GRID card with 0.5TB RAM and 24TB local disk space and supports up to 20 simultaneous users.

6. BIG DATA CHALLENGES

We have experienced several challenges during the processing chain development and testing stage with handling large data sizes and large data volumes. For example, to address large data size problem in SRR, we use tiling for large HiRISE scenes and mosaic afterwards. To address single processing computational complexity, we use computationally inexpensive algorithms for initial processing and use more complex methods to refine and improve the final results. Also, in the future, we will port most of our existing processing chains to the GPU server, to further accelerate single session processing time. Automated task scheduling system is also important to handle large data volumes. An scheduler/task distribution daemon based system was successfully tested previously for MER and MSL stereo panorama processing. The automated scheduling system is able to retrieve and initialize jobs from a catalogue database, monitor the processing status of jobs on server blades and update the database with processing results. A task distribution daemon competes for processing jobs created by the Scheduler, downloads input data if not already locally available, start/invokes processing software, retrieves results and processes logging information and stores them in local storage, creates production description file including software versions and all processing parameters and notifies the Scheduler that processing is done or has failed. Moreover, an automated image quality assessment mechanism has been developed at MSSL to distinguish different levels of image quality for HiRISE, CTX and HRSC images in order to improve the job success rate and save processing time from low quality input data.

ACKNOWLEDGEMENT

The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under iMars grant agreement n° 607379. This work has also received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement No. 312377 PRoViDE. Partial support is also provided from the STFC "MSSL Consolidated Grant" ST/K000977/1.

REFERENCES

[1] Shin, D. and J.-P. Muller, Progressively weighted affine adaptive correlation matching for quasi-dense 3D reconstruction. Pattern Recognition, 2012. 45(10): p. 3795 - 3809.

[2] Moratto, Z. M., M. J. Broxton, R. A. Beyer, M. Lundy, and K. Husmann, 2010. Ames Stereo Pipeline, NASA's Open Source Automated Stereogrammetry Software. Lunar and Planetary Science Conference 41, abstract #2364.

[3] Tao, Y. and J.-P. Muller, 2015. Towards a better understanding of Martian surface processes: zooming in for a quantitative assessment of key geomorphological features from super-resolution HiRISE images in comparison to overlapping rover Navcam image. European Planetary Science Conference, Nantes, 28.9-2.10.15

[4] Tao, Y. and J.-P. Muller, 2015 "A novel method for surface exploration: super-resolution restoration of Mars repeat-pass orbital imagery," *Planetary and Space Science*, http://www.sciencedirect.com/science/article/pii/S0032063315003591.

[5] Golombek, M., Grant, J., Kipp, D., Vasavada, A., Kirk, R., Fergason, R., et al. (2012). Selection of the Mars Science Laboratory Landing Site. Space Science Reviews, 170(1-4), 641–737. doi:10.1007/s11214-012-9916

BIG DATA FROM THE J-PAS AND J-PLUS SURVEYS

D. Cristóbal-Hornillos, J. Varela, A. Ederoclite, H. Vázquez Ramió, A. J. Cenarro, M. Moles, A. Marín-Franch, C. López-Sanjuan, A. Yanes-Díaz, A. López-Sainz, J. Hernández-Fuertes, T. Civera , D. Muniesa

Centro de Estudios de Física del Cosmos de Aragón, Plaza San Juan 1, Planta 2, E-44001 Teruel, Spain

ABSTRACT

The Observatorio Astrofísico de Javalambre (OAJ) is a Spanish astronomical facility designed to carry out large surveys. It has two main telescopes: JST/T250, a 2.55 m telescope with a FoV of 3 deg diameter, and JAST/T80, an 83 cm with a 2 deg FoV, both of them equipped with panoramic cameras, amounting more than 1.3Gpix together. During the first 6-7 years of operation, JST/T250 will be devoted to complete the Javalambre-PAU Astrophysical Survey (J-PAS), a photometric survey with 56 narrow-band ($\sim$14nm width), contiguous, optical filters (from $\sim$ 350 nm to $\sim$ 900 nm), plus several broad band filters covering an area of 8500$\square^\circ$. The JAST/T80 started operations in 2015, being devoted to perform the J-PLUS survey, covering in $\sim$ 3 years the same area as J-PAS using a system of 12 narrow and broad-band filters. Both surveys will produce $\sim$ 1 PB of raw data; a single data release of the final combined mosaics is $\sim$ 850 TB. The archive and processing is being done in a dedicated data center equipped with a processing system and $>$ 5 PB combining disks and tape technologies in order to process and store the data produced by the surveys.

Index Terms— Data management, Image processing, Catalogs, Observatorio Astrofísico de Javalambre

1. INTRODUCTION

CEFCA is a Center for Astrophysics and Cosmology located in Teruel (Spain) which is in charge of construction, operation and scientific exploitation of the Observatorio Astrofísico de Javalambre (OAJ, [1], [2]).

The OAJ is a Spanish Singular Scientific and Technological Infrastructure (ICTS) for astronomy located at 1957 m in the Sierra de Javalambre, Teruel (Spain), that has been conceived to carry out large sky multi-filter surveys. The main telescope at the OAJ is the JST/T250 (Javalambre Survey Telescope), a large-etendue 2.55 m telescope with a FoV of 3 deg diameter. The telescope is currently in the last stages of commissioning with a temporary verification camera, awaiting for the arrival of JPCam in June 2016. JPCam is a large panoramic camera constituted by a mosaic of 14 large format e2v CCDs of 9.2 k$\times$9.2 k pix, covering the large focal plane of the telescope. With a pixel size of 10$\times$10 microns, JPCam provides a plate scale of 0.23 arcsec/pix at the Cassegrain focus. It includes five filter trays with 14 positions each to host up to 70 different filters simultaneously.

In addition, the OAJ includes and operates the JAST/T80 (Javalambre Auxiliary Telescope), an 83 cm telescope with a 2 deg FoV. T80Cam is the panoramic camera operating in this telescope, it has one large-format CCD by e2v of 9.2 k$\times$9.2 k pix, identical to the ones in JPCam, and providing a plate scale of 0.55 arcsec/pix. T80Cam mount 12 different filters simultaneously in 2 filter wheels. JAST/T80 and T80Cam are in operation since the beginning of 2015. During the second half of 2015 sciencific quality images have been acquired for several projects and the first images of the J-PLUS survey have been collected.

2. J-PLUS AND J-PAS SURVEYS

During the first years of operation, both telescopes will be mainly dedicated to complete two large sky multi-filter surveys. JST will be devoted to conduct the Javalambre-PAU Astrophysical Survey (J-PAS [3]), an extragalactic survey of 8500$\square^\circ$ in an optical filter system of 56 contiguous, narrow-band filters of 14nm width (from $\sim$ 350 nm to $\sim$ 900 nm), plus several broad band filters. In the end, J-PAS will provide a low resolution ($R \sim 50$) spectrum for each observed pixel in the sky, hence performing like a low-resolution IFU. This data will allow to measure accurate photometric redshifts for several hundred millions of galaxies. The central goal of the project is to determine the equation of state of the Dark Energy through the analysis of radial Baryonic Acoustic Oscillations[4]. Nevertheless J-PAS will provide valuable data for different studies in Galaxy Evolution, SNe, and Milky Way structure among others[3].

During the next 3 years JAST/T80 and T80Cam will carry out the Javalambre Photometric Local Universe Survey (J-PLUS, Cenarro et al. in prep), aimed to observe the same J-PAS area with 12 broad, medium, and narrow-band filters, the later located at key features for stellar classification and stellar population studies. The main goal of J-PLUS is to procure the photometric calibration for J-PAS. At the same time small modifications in J-PLUS data collection strategy

will produce valuable data for scientific projects in different fields of astronomy. For example, in two bands gSDSS and rSDSS the area will be visited in 3 different epochs allowing for variability studies. J-PLUS has started to obtain the first scientific data during late 2015.

3. DATA ACQUISITION AND TRANSFERENCE

The image acquisition at OAJ is controlled using Robotic Observing Blocks (ROB). ROB is an operation mode of the Control Execution Module included in the Observatory Control System, which is essential to optimize the observations by the different OAJ telescopes. When ROB mode is activated the whole observatory prepare all the systems to automatically execute queued observations from a database. ROB points the telescopes to the fields detailed in the observing blocks, control the hexapods in order to focus, set the required filters and parameters in the cameras and other subsystems to perform expositions. The system is prepared to be used with a Sequencer and Scheduler modules which are in charge of feeding the observing blocks into the database taking into account a predefined survey strategy, ephemerids and atmospheric conditions.

During the operations of J-PAS and J-PLUS up to 1.5 TB of data can be collected during a night. The OAJ have dedicated servers to automatically manage the transmission of the images from camera servers to the OAJ storage of $\sim$ 100TB, and produce an online processing using lightweight versions of the pipelines. The data are automatically transferred to the Unidad de Procesado y Archivo de Datos (UPAD) at CEFCA using a dedicated radio-link. During daytime automatic backup copies of the night data will be done in magnetic tapes.

4. DATA HANDLING AND PROCESSING SOFTWARE

The data Management Software is designed to handle the enormous data flow produced by the panoramic cameras making optimal use of the hardware deployed at the datacenter. On one hand, daily the Data Management Software shall automatically process the data collected during the night to check its quality, update the survey's databases and feed the Scheduler to compute the telescope targets of the following nights. On the other hand, the pipelines reprocess all the images when the proper master Calibration Frames are available to guarantee the scientific quality of the data.

When both telescopes are operational they will produce $\sim$ 8000 images per night, to sustain this data rate the pipelines have been designed to work mostly automatically. Once the raw data reach the main archive of UPAD they are automatically uploaded to a management database which stores all the metadata of the collected images. This database is used to execute the different stages of the pipelines. During the pipeline jobs, the steps done over each image and all the meaningful parameters of the products are recorded in the management database. This allows to control when the inputs for the next stage are ready and automatically prepare the jobs to be executed. The information recorded in the management database also permits to recover the processing history of any individual or coadded image.

The image processing pipelines have three main stages. The first one is related with the generation of master Calibration Frames (CF) which afterwards are used to correct the instrumental imprint on the individual images. The Calibration Frames are generated by higher order scritps that need as inputs the target Frame (e.g. master, bias, dome or sky flat, fringing pattern, zero point map), the time interval and the instrumental configuration (telescope, camera and detector parameters, read-out configuration, etc). Using this information the pipeline queries the database for the needed images and process the master CF. Once the CF are generated they are available in a Web Tool connected to the database and filesystem. This tool allows for a quick verification of the products. Figure 1 shows some visual information available from the Web during a flat validation. Finally, the master CFs have to be validated by some user and only validated CFs are used in further processing of science images.

Fig. 1. Capture of the interactive view in the Web Tool devoted to validate a master dome flatfield.

The second stage is related with the processing of the different individual images. The pipeline corrects the instrumental signature using the validated Calibration Frames, masks contaminants (cosmic rays, satellite traces) and calibrates the astrometry and the photometry of the images. The desired corrections (bias, flat, fringing, etc) to apply for each image type are coded in the configuration files. The pipeline automatically queries the management database to find the appropriate Calibration Frames for each image depending on the needed corrections, observation epoch and instrumental configuration. For the detection and masking of cosmic ray hits we are using an implementation of L.A.Cosmic [5]. Once the

individual images are processed a first catalog using SExtractor [6] is created over the individual images.

In order to produce an online processing after observation it is possible to enable a pre-processing option that allows to use validated Calibration Frames from a previous block. Those preprocessed images are used mainly for a quick validation but also by science projects that need a quick response after image acquisition.

Finally, the Tiling pipeline combines the processed images from the same sky area to provide a deeper photometric calibrated image for each filter. To perform the astrometric calibration and image coadding we are using the Astromatic packages Scamp [7] and Swarp [8]. The combined images for a Tile in the whole set of filters compound a Datacube. SExtractor source catalogs are extracted from the final Tiles and Datacubes and are stored in an internal database system. Two kind of catalogs are produced. In the first one, the source detection and segmentation is done independently for the combined images of each filter. In the second one, the detection and aperture definition is done in the Sloan r' band used as reference image. The survey strategy will guarantee a better atmospheric seeing during the acquisition of the reference filter. In order to compute the catalogs a PSF homogenization is done to match each filter with the reference one. This technique was already used to produce the photometric catalogs for the ALHAMBRA survey [9].

To verify the status of the raw or processed frames, the Web Tool allows to navigate through the Datacubes, Tiles, individual images, and CFs. This web interface uses the administrative database to provide an agile access to the status, processing history, metadata, measurements and statistics of the images and plots that allows the validation of the image quality and the data treatment. Figure 2 shows the rendered frame produced by the pipeline for an image of M33.

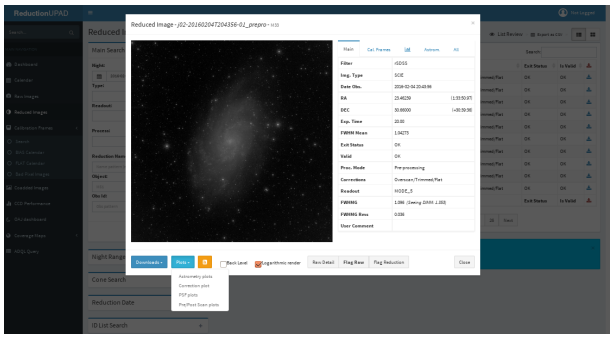

Fig. 2. Capture of the rendered view of an individual processed image displayed in the Web Tool.

5. UPAD STORAGE AND PROCESSING FACILITIES

The J-PAS survey will collect of the order of 10 million raw images, but it is not only the large number of images produced by the OAJ telescopes what imposes a bigdata management challenge. The large format detectors used in T80cam and JPCam produce individual images of $\sim$ 90Mpix, JPCam producing 14 of them simultaneously. Those individual images in raw format are $\sim$ 170MB, but once processed amount to $\sim$ 340MB. To handle a great amount of images of several hundreds of MB is a IO bound problem where the transference of the images to the processing nodes has a big impact in the processing times.

A new datacenter (UPAD) has been deployed to store, archive and automatically process the OAJ data. UPAD storage and processing hardware has been designed in order to handle the large images produced by the telescopes. At the UPAD datacenter the main storage is divided in two tiers. The data that is accessed frequently by the pipelines is kept in a disk storage cluster. The actual implementation of the UPAD cluster provides more than 1.0 PB of net storage capacity. In order to feed data to the processing nodes, the storage cluster runs a distributed filesystem providing access to the data through several servers. The storage cluster and the core network provide an aggregated bandwidth larger than 50 Gbps. The permanent archive and the backup of the data will be done in a robotic tape library with $\sim$ 4 PB. The storage cluster and robotic library are interconnected with a Hierarchical Storage Managenement (HSM) system solution. The HSM allows us to define politics to backup or archive in the tape library the data with low access requirements.

Fig. 3. Picture of the main storage and processing resources at the UPAD datacenter.

Regarding data processing, several nodes with more that 450 cores in total will be used to run the pipelines using SGE as a batch-queue manager. The processing servers have a large RAM per core, as the pipeline has to manage large files and most of the steps are done avoiding writing to disk. Also

the computing nodes have several TB of internal disk storage to retain local copies of files systematically used during the data processing. As an example, the master CFs will be cached in each processing node to speed up the processing, reducing the overall network load.

6. THE ACCESS TO THE DATA

Our development regarding data distribution follows the successful experience of the Sloan Digital Sky Survey, and it is based on standards defined in the Virtual Observatory (VO). This allows us to follow established standards and protocols and use software already developed inside the VO framework.

Once J-PAS is completed the Scientific database will have of the order of 6×10^8 records. For each astronomical object in the catalog, ~ 1400 properties will be stored. For the J-PLUS and J-PAS surveys we are using the relational database PosgreSQL.

To offer data access to the astronomical community, a web portal has been developed which offers both the possibility to "navigate" across the images and to query the database. The web-portal navigator (largely inspired in the SDSS Sky Server) allows the user to see the available footprint of the survey and access the data by an interactive sky-map. The user can browse through RGB composite images and overlay marks of a search based on magnitudes, colours and redshift. All the services which are available through the web-portal are also accessible through VO-protocols. Finally, the user has the possibility to do asynchronous queries to the database.

For the publication of the J-PLUS and J-PAS data we are deploying a dedicated hardware system. This system consists mainly of redundant web-servers with large internal storage capabilities, and a pool of database servers. The web-servers run the web-portal and manage the communication mechanism to retrieve the data from UPAD as demanded by the users. The database servers are optimized to respond the user queries and accumulate more than 30 TB of storage to keep online at least the last two releases of the source catalogs.

7. SCIENTIFIC ANALYSIS CHALLENGES

The final scientific analysis of the huge OAJ datasets is also a big data challenge. Thanks to the J-PLUS and J-PAS volumes, statistical uncertainties will be low compared with the systematics in our analysis techniques that will dominate the error budget in our final measurements. With the usual photometric techniques prone to biases ([10], [11]), and a too low resolution (R ~ 50) to successfully apply usual spectroscopic techniques, new and suited methodologies are mandatory to extract unbiased galaxy distributions and cosmological constraints from the J-PAS data.

In order to extract statistical information from the data, posterior Probability Distribution Functions (PDFs) are recognized as the right way to deal with photometric redshifts (e.g., [12], [13], [14]). Nowadays Bayesian inference is widely used to estimate galaxy properties (e.g. [15], [16]). Nevertheless, current distribution estimators assume galaxies with a fixed redshift, luminosity, stellar mass, etc. Given the probabilistic nature of the photometric redshifts, any galaxy property becomes probabilistic and thus the posterior PDF of galaxy properties must be tracked along all the analysis process. We are developing novel and optimised techniques that use posterior PDFs to analyse multi-filter survey data, with the aim of dissecting the fundamental galaxy properties and of producing unbiased and accurate posterior distributions for galaxy evolution and cosmology studies.

Acknowledges: The OAJ is funded by the Fondo de Inversiones de Teruel, supported by both the Government of Spain (50%) and the regional Government of Aragón (50%), FEDER funds, and the Brazilian funding agencies FINEP, FAPESP and FAPERJ. The UPAD storage and processing infrastructure is funded by the Spanish MINECO (FCDD10-4E-867), cofunded by the European Fund for Regional Development (FEDER) and Fondo de Inversiones de Teruel (FITE). This work has been funded by through the PNAYA, under grant AYA2012-30789.

8. REFERENCES

[1] A. J. Cenarro and et al 2014, 9149, Proc. SPIE, 91491I.

[2] A. J. Cenarro and et al. 2010, 7738, Proc. SPIE, 77380V.

[3] N. Benitez and et al. *ArXiv e-prints: 1403.5237*, 2014.

[4] N. Benítez and et al. ApJ, 691, 241–260, 2009.

[5] P. G. van Dokkum PASP, 113, 1420–1427, 2001.

[6] E. Bertin and S. Arnouts A&AS, 117, 393–404, 1996.

[7] E. Bertin 2006, 351, *ASPCS*, 112.

[8] E. Bertin and et al. 2002, 281, *ASPCS*, 228.

[9] A. Molino and et al. MNRAS, 441, 2891–2922, 2014.

[10] R. K. Sheth and G. Rossi MNRAS, 403, 2137–2142, 2010.

[11] G. Rossi and et al. MNRAS, 401, 666–676, 2010.

[12] Fernández-Soto and et al. MNRAS, 330, 889–894, 2002.

[13] C. E. Cunha and et al. MNRAS, 396, 2379–2398, 2009.

[14] S. J. Schmidt and P. Thorman MNRAS, 431, 2766–2777, 2013.

[15] G. Kauffmann, T. M. Heckman, and et al. MNRAS, 341, 33–53, 2003.

[16] G. A. Blanc and et al. ApJ, 798, 99, 2015.

AN UNSUPERVISED IMPLEMENTATION OF THE P-SBAS DInSAR ALGORITHM FOR PROCESSING LARGE DATA VOLUMES THROUGH DISTRIBUTED COMPUTING INFRASTRUCTURES WITHIN OPERATIONAL ENVIRONMENTS

Zinno I., Casu F., De Luca C., Lanari R., Manunta M.
CNR IREA, Via Diocleziano 328, 80124 Napoli, Italy

ABSTRACT

In the last decades Earth Observation (EO) from space has very fast evolved through the development of remote sensing data-acquisition systems, contributing to the creation of a Big EO Data scenario. In this paper we discuss the unsupervised implementation of the P-SBAS DInSAR algorithm, a data and processing intensive EO technique that permits us to generate ground displacement time series from a set of SAR acquisitions, which is able to run on distributed computing infrastructures by effectively exploiting a large number of computing resources. In particular, we introduce the P-SBAS deployment within the ESA Geohazards Exploitation Platform. Moreover, we present the P-SBAS cloud computing implementation within the Amazon Web Services cloud. The goal is to demonstrate that it is possible to automatically process, in short time frames, very large SAR datasets, such as those expected by the Sentinel-1 constellation, to provide the scientific community with new EO tools and services and finally lower the barrier for accessing to data and processing intensive EO techniques.

1. INTRODUCTION

In the last decades Earth Observation (EO) from space has very fast evolved through the development of remote sensing data-acquisition systems and techniques capable to extract value added information from the collected data.

The current EO scenario is characterized by a huge availability of Synthetic Aperture Radar data that have been acquired during the last 25 years by past and present sensors. These data include the long-term ESA archives collected by the ended missions ERS-1, ERS-2 and Envisat. Moreover we have at disposal the data provided by the currently operational X-band generation SAR sensors, such as COSMO-SkyMed (CSK) and TerraSAR-X (TSX). Furthermore, a massive and ever increasing data flow is supplied by the recently launched (April 2014) Sentinel-1A (S1A) SAR satellite operating within the framework of the Copernicus program of the European Union. This sensor will be paired during 2016 with the Sentinel-1B twin system that will allow halving the satellites revisit time from 12 to 6 days. Therefore, due also to its very large spatial coverage, Sentinel-1 huge data archives relevant to extended areas on Earth are already available and will be further enlarged thanks to a totally "free and open access" acquisition data policy.

Moreover, among the SAR data exploitation methodologies, we focus on Differential SAR Interferometry (DInSAR), which is a well-established microwave remote sensing technique that allows us to estimate the ground deformations with centimeter to millimeter accuracy [1]. Over time, DInSAR methodology has moved from the analysis of single deformation episodes towards the study of the temporal evolution of the detected deformations, also thanks to the availability of the above-mentioned large SAR data archives. A very well known DInSAR algorithm is the one referred to as Small BAseline Subset (SBAS) [2], which is able to generate mean deformation velocity maps and displacement time series from multi-temporal SAR datasets; it is, besides, capable to perform analyses at different spatial scales and with multi-sensor data.

Recently, an advanced algorithmic parallel computing solution, referred to as P-SBAS, that encompasses diverse parallelization strategies, both multi-nodes and multi-cores, and has shown to have good scalable performances, has been developed [3], [4]. P-SBAS implements the whole SBAS processing chain, starting from raw data focusing and ending with displacement time-series generation, in a totally automatic way. More specifically, it is capable to run without needing external expert interventions, once some initial parameters are properly set. Such an unsupervised implementation, besides, has been designed to effectively exploit High Performance Computing (HPC) and distributed computing infrastructure (cluster, grid, or cloud computing platforms). Therefore, the advanced P-SBAS algorithm is very suitable to run in operational environments aimed at providing the scientific community with new EO services and products.

In this work we present the deployment of the P-SBAS processing chain within two different environments. The first one is the ESA Geohazards Exploitation Platform (GEP), which aims at exploiting satellite EO data and techniques for mapping the hazard-prone land surfaces and monitoring the terrain deformation. The P-SBAS algorithm within GEP is an on-demand web tool made available to scientists that are not expert on interferometric SAR data processing. It allows them to generate, quickly and in an

unsupervised way, displacement time series and surface deformation mean velocity maps of both ERS and ENVISAT datasets. Moreover, very recently, we implemented a new service for the S1 Interferometric Wide Swath (IWS) data processing within the GEP environment. To deal with the very big data volume expected from S1, we developed a P-SBAS implementation that is ad-hoc designed to manage the S1 data, which are composed of series of bursts that can be considered as separate acquisitions. The S1 implementation of the P-SBAS algorithm is designed to exploit as much as possible the burst characteristics of S1 data in order to achieve a significant granularity level for the multi-node parallelization and therefore to optimize performances.

The second environment in which we deploy the P-SBAS processing chain is the Amazon Web Services (AWS) public cloud. In this case the goal is to demonstrate the capability of P-SBAS to properly exploit a very large number of computing resources to carry out the processing of a big SAR dataset. We prove in this paper the feasibility of such experiment, showing the results relevant to a very large interferometric analysis involving the processing of 334 ENVISAT images, which has been carried out by exploiting 64 computing instances of the AWS cloud.

The proposed P-SBAS implementation, integrated in operational environments as those above-mentioned, represents the starting point for advanced Sentinel-1 services able to provide users with continuously and systematically updated deformation time-series of very large areas on the Earth's surface.

2. THE P-SBAS WEB TOOL WITHIN GEP

Within the current remote sensing scenario characterized by the disposal of both massive SAR data archives and algorithms able to properly exploit them, there is an ever increasing need of environments capable to put them together and, at the same time, to provide high performance computing HPC infrastructures to process the data and make the obtained results available to the EO community. Exactly towards this direction ESA is building an ecosystem of Thematic Exploitation Platforms (TEP) focused on the capitalization of Ground Segment capabilities and Information and Communication Technologies (ICT) to maximize the exploitation of EO data from past and future missions. In particular, a TEP refers to a computing platform that follows a given set of scenarios for users, data and ICT provisioning, aggregated around an Earth Science thematic area. Among the TEPs, it is already started the implementation of the GEP, which aims at exploiting EO data for risk assessment. In particular, the GEP aims at the exploitation of EO in the context of the Geohazard Supersites & Natural Laboratories and CEOS Seismic Hazards and Volcanoes Pilots.

The P-SBAS automatic algorithm has been already deployed within the ESA's Grid Processing on Demand (G-POD) and, in this version, is fully accessible from the GEP. G-POD benefits from the access to the large ESA computing facilities as well as to EO data archives (ESA SAR data Virtual Archive 4), and provides a friendly web user interface that permits to set up an efficient and on-demand P-SBAS processing web tool addressed to scientists that are not expert on interferometric SAR data processing.

The P-SBAS web tool integrated within GEP generates surface deformation velocity maps and time series starting from ERS or ENVISAT datasets; moreover, recently, a new service relevant to the S1A data exploitation has been opened. At the time of writing the S1A data processing is able to automatically produce interferograms, by also jointly processing and combining more than one S1A slice, but very soon also the time series generation will be made available.

In order to show its potential, we present some results that have been achieved by exploiting the P-SBAS web tool through the GEP platform. In particular, the panel (a) of Figure 1 represents the usage of the P-SBAS service during the period January-May 2015. The green circles illustrate the sites that have been processed a single time by the users, whereas red circles represent sites on which multiple runs of the P-SBAS algorithm have been launched. In Figures 1 (b), (c) and (d) the mean deformation velocity maps relevant respectively to the Napoli bay area (Italy), Los Angeles (California) and the Mauna Loa volcano generated through the P-SBAS web tool are represented.

Furthermore, for validation purposes, we compared the P-SBAS results with those obtained by continuous GPS stations. In particular, in Figure 2 the surface deformation time series relevant to the GPS measurements (red stars) and P-SBAS (black triangles) of three GPS stations located in the maximum deformation area of the Campi Flegrei caldera, which is an active volcanic area located in the Napoli bay, are shown. The standard deviations (σ) of the difference between the GPS and the P-SBAS measurements are also reported, demonstrating the very good match between the two measurements.

As for the results provided with the S1A service, in Figure 3 we show the interferogram generated through the GEP platform relevant to the Chile earthquake that occurred on September 16, 2015. It has been obtained by jointly processing three S1A slices, covering an area of about 750 km x 250 km, allowing us to analyze the entire region affected by the earthquake.

doi: 10.2788/854791

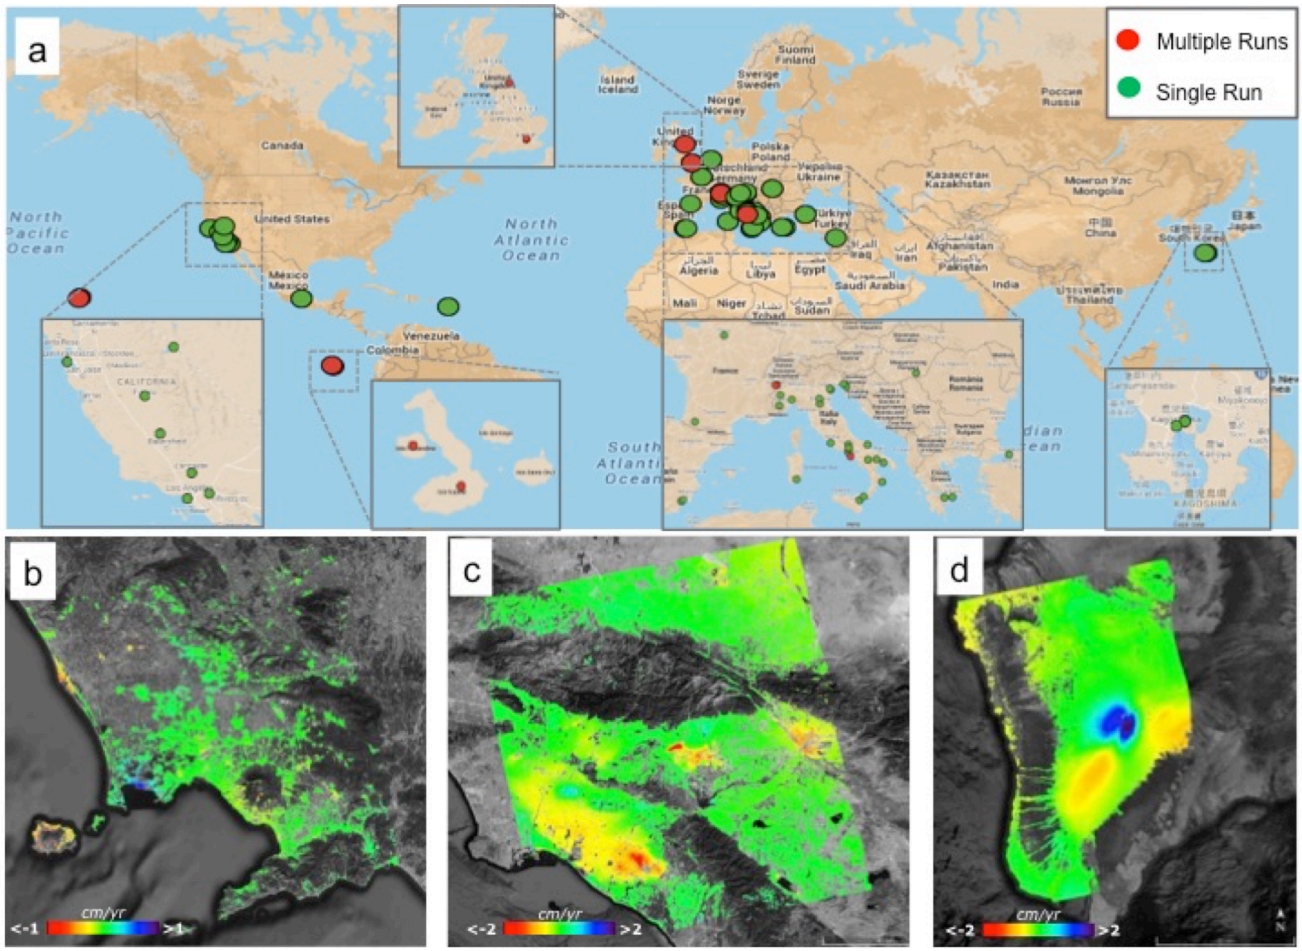

Figure 1. P-SBAS service usage (a) Pictorial representation of the processing results achieved by the users through the P-SBAS web tool within G-POD, during the period Januar-May 2015; (b), (c), (d) Retrieved mean deformation velocity maps relevant to the Napoli Bay (Italy), Los Angeles (California) and Mauna Loa Volcano (Hawaii) areas, respectively.

3. MIGRATION OF P-SBAS TO THE AWS CLOUD

Cloud Computing technologies are well established within the ICT field and can give great benefit also to the scientific application context. Indeed, the exploitation of customized computing infrastructures built up within cloud environments can be crucial in many respects; first of all because of the theoretically unlimited computing facilities they make available, afterwards for the flexibility and the easiness of use they provide, and, last but not least, for the achievement of processing performances improvement. In such a context, we have carried out the migration of the P-SBAS algorithm to the Amazon Web Services (AWS) public cloud environment [4], [5]. In this case our main objective has been to implement the parallel P-SBAS algorithm within the cloud environment in such a way that it could benefit as much as possible from the huge availability of computing nodes. Accordingly, we carefully evaluated the P-SBAS scalable performances as well as the main bottlenecks to such scalability within cloud environments [6]. In order to show the processing capability of the proposed P-SBAS cloud solution, we represent in Figure 4 the mean deformation velocity map of a large area located in the Southern California, including the Los Angeles and San Francisco metropolitan areas, which has been generated by exploiting the AWS resources. In particular, for such an analysis, we processed in parallel 8 ENVISAT standard frames, for a total of 334 ENVISAT acquisitions, covering an area of about 400 km x 250 km. We implemented a computing architecture made up of 64 computing nodes, for a total of 512 CPUs (we used the m4.2xlarge instances of the AWS Elastic Compute Cloud EC2 [7]) and, concerning the storage, 64 Provisioned IOPS SSD volumes, guaranteeing high I/O performances. The overall processing lasted less than 8 hours and cost about 440 USD.

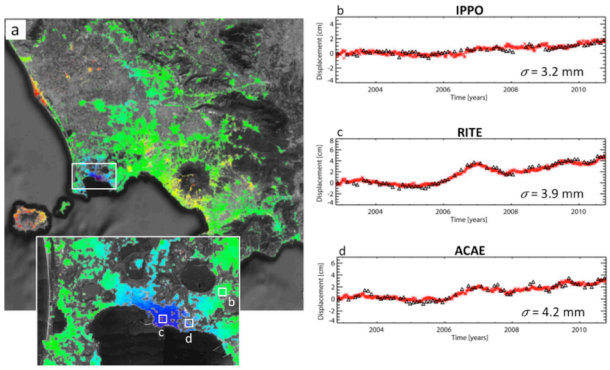

Figure 2. Unsupervised P-SBAS processing chain validation experiments on the Campi Flegrei caldera (Napoli) test site. We show: (a) the mean deformation velocity map retrieved from the P-SBAS analysis with the locations of three selected GPS stations; (b), (c) and (d) the comparisons between the GPS (red stars) and the P-SBAS (black triangles) surface deformation time series relevant to the stations identified in the corresponding map; the standard deviations (σ) of the difference between the GPS and the DInSAR measurements.

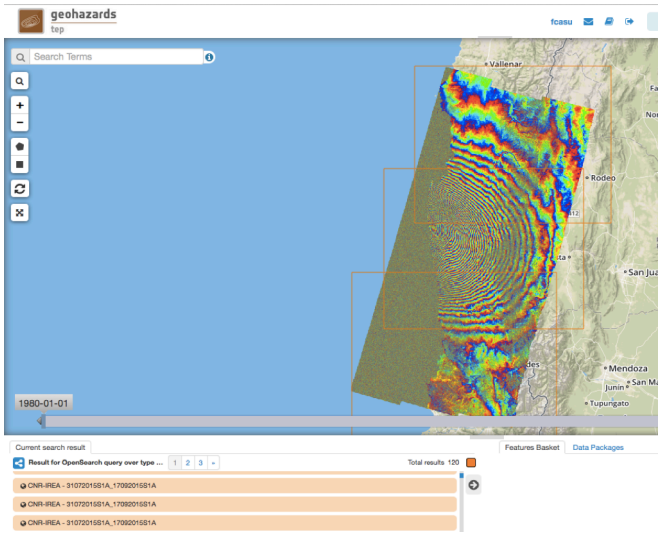

Figure 3. Interferogram produced through the S1A service within GEP relevant to the Chile earthquake that occurred on September 16, 2015. It has been generated by jointly processing three S1A slices that are represented by the rectangles in the figure.

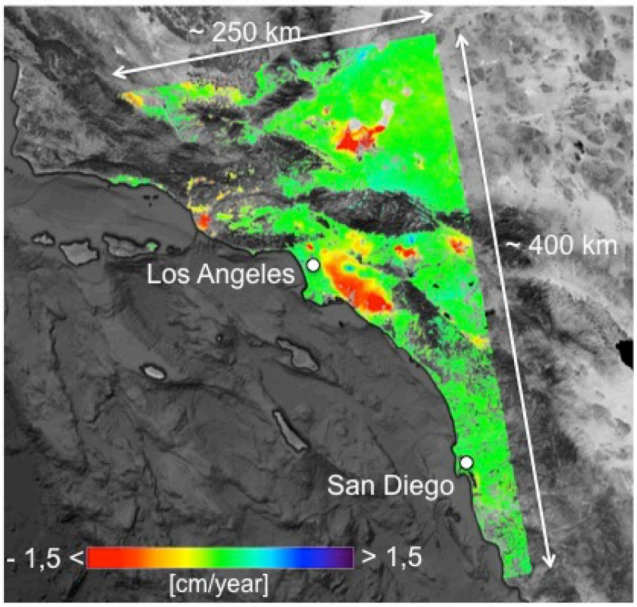

Figure 4. Mean deformation velocity map relevant to an area of the Southern California generated by processing in parallel 334 ENVISAT images through 64 computing nodes within the AWS cloud.

4. REFERENCES

[1] D. Massonnet and K. L. Feigl, "Radar Interferometry and its application to changes in the Earth's surface," Rev. of Geophys., vol. 36, pp. 441–500, 1998.

[2] P. Berardino et al., "A new algorithm for surface deformation monitoring based on small baseline differential SAR interferograms," IEEE Trans. Geosci. Remote Sens., vol. 40, no. 11, pp. 2375–2383, Nov. 2002.

[3] Casu, F. et al., "SBAS-DInSAR Parallel Processing for Deformation Time-Series Computation," in Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of, vol.7, no.8, pp.3285-3296, Aug. 2014.

[4] Zinno, I. et al., "A First Assessment of the P-SBAS DInSAR Algorithm Performances Within a Cloud Computing Environment," in Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of, doi: 10.1109/JSTARS.2015.2426054.

[5] Zinno, I. et al., "Cloud Computing for Earth Surface Deformation Analysis via Spaceborne Radar Imaging: a Case Study," in Cloud Computing, IEEE Transactions on, doi: 10.1109/TCC.2015.2440267.

[6] Zinno, I. et al., "New advances in intensive DInSAR processing through cloud computing environments," in Geoscience and Remote Sensing Symposium (IGARSS), 2015 IEEE International , vol., no., pp.5264-5267, 26-31 July 2015.

[7] Amazon EC2 - aws.amazon.com/ec2

DATA-INTENSIVE COMPUTING IN RADIATIVE TRANSFER MODELLING

Efremenko D.S., Loyola D., Doicu A., Trautmann T.

Remote Sensing Technology Institute, German Aerospace Center (DLR),
Oberpfaffenhofen, Wessling, Germany

ABSTRACT

The operational processing of remote sensing data requires high-performance radiative transfer model (RTM) simulations. To date, considerable success has been achieved in dimensionality reduction techniques as well as in heterogeneous multi-CPU/GPU computing for highly intensive parallel computations. We have developed several techniques for accelerating the radiative transfer solver. They include (1) analytical methods which allow to compute set of atmospheric scenarios in one RTM call; (2) dimensionality reduction of the datasets, and (3) GPU-computing using CUDA framework. These techniques provide almost 300x cumulative speed-up for the RTM with respect to the original single-threaded CPU code. In this paper, we analyze the applicability of the proposed methods to a practical problem of total ozone column retrieval from UV-backscatter measurements.

***Index Terms*—** Radiative transfer models, discrete ordinate method, CUDA, heterogeneous computing, dimensionality reduction

1. INTRODUCTION

Massive amounts of spectral information are expected from the new generation of European atmospheric sensors (Sentinel 5 Precursor, Sentinel 4 and Sentinel 5). They impose new challenges to data driven algorithms. In this regard, a fast processing of the data in the UVNS spectral domain is required.

The radiative transfer modelling (RTM) is a critical part in the processing chain from the raw instrumental data (level 0) to the geophysical products (level 2) and is the major performance bottle-neck for the retrieval algorithms. Furthermore, the processing of satellite-measured atmospheric composition data involves many computational loops. These are shown in the following serial-CPU pseudo-code:

```
for each pixel:
  for each wavelength:
    for each cloud_fraction:
      for each geometry:
        call RTE_solver();
```

We have developed several techniques for RTM performance enhancement with particular application to trace gas retrievals. Some of them are used to accelerate the radiative transfer solver itself [1] while others are designed to optimize the loops containing the radiative transfer solver [2, 3]. They are described in the following sections. In this study we also investigate the cumulative the performance enhancement obtained by using all these methods together.

2. ACCELERATION TECHNIQUES FOR THE DISCRETE ORDINATE METHOD

The radiative transfer equation (RTE) in a one-dimensional medium is well-known [4]. The discrete ordinate method (DOM) for solving the RTE is numerically stable for arbitrary optical thicknesses in a multi-layer stratified medium. An important parameter controlling the computational time and the accuracy of computations is the number of streams in the polar hemisphere N_{do}. RTMs are called multi-stream if $N_{\text{do}} \geq 2$, and two-stream if $N_{\text{do}} = 1$. Two-stream RTMs are based on closed-form solutions for the radiance, and are therefore considerably faster than multi-stream models. However, two-stream accuracy is not high enough for practical applications in remote sensing. In contrast, multi-stream models are computationally expensive.

To speed-up the RTM, we implemented several acceleration techniques for the discrete ordinate method including:

1. The computation of the inverse of the eigenvector matrix by first scaling the original matrix to yield a symmetric matrix, and then by calculating the inverse of the symmetric matrix by means of the left eigenvector matrix.

2. The use of the telescoping technique, which consists in the reduction of the linear algebra system to the active layers of the clouds and for azimuthal modes $m \geq 3$.

3. The use of an additional discrete ordinate with zero weight in the direction of the line of sight in order to avoid the post-processing step of the conventional discrete ordinate method (source integration along the line of sight).

The most time consuming part of the radiative transfer solver is the eigenvalue problem which is related to the scattering properties of the atmosphere. The code is designed in such a way, that it aggregates the scenarios with common scattering properties and then computes spectral radiances for a set of solar zenith angles, viewing zenith angles and relative azimuth angles at no additional computational costs.

3. LOOP OPTIMIZATION

3.1. Dimensionality reduction of the input data

To optimize performance over spectral loops, we extended the RTM with principal component analysis (PCA) of optical parameters [5]. This approach has the following features: (a) a two-stream model is used to compute the approximate spectrum; (b) differences (or "correction factors") between the approximate and exact solutions are expressed through a second-degree polynomial in the optical parameters; (c) PCA is used to map the initial data set of optical properties to a lower-dimensional subspace, in which the computation of the correction factors is performed.

Assume, the optical parameters, representing the input parameters of the radiative transfer code, are encapsulated in the vector $\mathbf{x}$. High-dimensional real data often lies on or near a lower-dimensional manifold. The fundamental issues in dimensionality reduction are the modeling of the geometry structure of the manifold, and the design of an appropriate embedding for data projection. For the N-dimensional data set $\{\mathbf{x}_w\}_{w=1}^{W}$, where $\mathbf{x}_w \in R^N$ and W is the number of wavelengths, let $\bar{\mathbf{x}} = (1/W)\sum_{w=1}^{W}\mathbf{x}_w$ be the sample mean of the data. The goal of a linear embedding method is to find an M-dimensional subspace ($M < N < W$) spanned by a set of linear independent vectors $\{\mathbf{a}_k\}_{k=1}^{M}$, such that the centered (mean-removed) data $\mathbf{x}_w - \bar{\mathbf{x}}$ lie mainly on this subspace (manifold),

$$\mathbf{x}_w \approx \bar{\mathbf{x}} + \sum_{k=1}^{M} y_{wk}\mathbf{a}_k = \bar{\mathbf{x}} + \mathbf{A}\mathbf{y}_w, \quad w = 1, \ldots, W, \qquad (1)$$

here $\mathbf{A} = [\mathbf{a}_k]_{k=1}^{M}$ is an $N \times M$ matrix comprising the column vectors $\mathbf{a}_k$, and y_{wk} is the kth component of the vector of parameters $\mathbf{y}_w \in R^M$.

In the operational processor, the radiance I is computed from

$$\ln \frac{I(\mathbf{x}_w)}{I^{\mathrm{TS}}(\mathbf{x}_w)} = f_{\mathrm{I}}(\mathbf{x}_w), \qquad (2)$$

where I^{TS} is the radiance computed by the two-stream model, and f_{I} is a correction factor. Setting

$$\Delta\mathbf{x}_w = \sum_{k=1}^{M} y_{wk}\mathbf{a}_k, \qquad (3)$$

we approximate $f(\mathbf{x}_w)$ by a second-order Taylor expansion around $\bar{\mathbf{x}}$, that is,

$$f(\mathbf{x}_w) \approx f(\bar{\mathbf{x}} + \Delta\mathbf{x}_w) \approx f(\bar{\mathbf{x}}) + \Delta\mathbf{x}_w^T \nabla f(\bar{\mathbf{x}}) + \frac{1}{2}\Delta\mathbf{x}_w^T \nabla^2 f(\bar{\mathbf{x}})\,\Delta\mathbf{x}_w, \qquad (4)$$

where ∇f and $\nabla^2 f$ are the gradient and the Hessian of f, respectively. Using central differences to approximate the first and the second-order directional derivatives gives

$$f(\mathbf{x}_w) \approx f(\bar{\mathbf{x}}) + \frac{1}{2}\sum_{k=1}^{M}\left[f(\bar{\mathbf{x}} + \mathbf{a}_k) - f(\bar{\mathbf{x}} - \mathbf{a}_k)\right] y_{wk} + \frac{1}{2}\sum_{k=1}^{M}\left[f(\bar{\mathbf{x}} + \mathbf{a}_k) - 2f(\bar{\mathbf{x}}) + f(\bar{\mathbf{x}} - \mathbf{a}_k)\right] y_{wk}^2. \qquad (5)$$

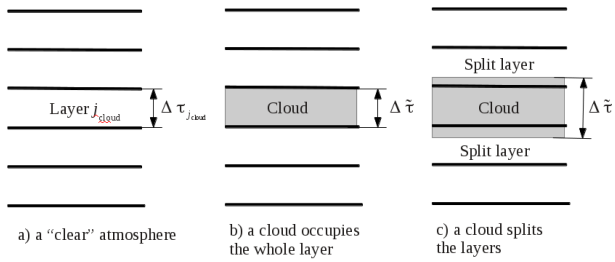

Fig. 1. Cloud position with respect to atmosphere layers.

From Eq. (5) it is apparent that the computation of the correction factor requires $2M + 1$ calls of the full- and two-stream models. As a result and taking into account that $M \ll W$, we are led to a substantial reduction of the computational time.

A similar approach is used to compute derivatives of the radiance (Jacobians) with respect to atmospheric parameters. Forward-model RTM simulations for total ozone retrieval in the wavelength domain 325–335 nm (Huggins bands) containing 88 spectral points were obtained by calling the multi-stream model with 8 streams per hemisphere only 5 times and the faster two-stream model 93 times. The speed improvement was about 8, with the maximum radiance error smaller than 0.2%.

3.2. Computations under cloudy conditions

In the independent-pixel approximation for cloud-contaminated scenes, radiances are computed as a linear superposition of two solutions for the clear-sky and fully-cloudy scenarios, requiring two RTM calls. We developed two methods based on the re-use of results from clear-sky RTM calculations to speed up corresponding calculations for the cloud-filled scenario [6]. The first approach is numerically exact, in that results from the clear sky computation can be saved in memory and reused for all non-cloudy layers in the second computation involving clouds. The insertion of a cloud layer in a clear sky atmosphere will affect the atmospheric layering scheme. This depends on the cloud-top height and the cloud geometrical thickness, and the possible options are illustrated in Figure 1.

The simplest case involves a cloud with optical thickness $\Delta\tau$ introduced into the layer j_0 which has clear-sky optical depth $\Delta\tau_0$ (Figure 1b). In this case, when solving the clear-sky problem we store the temporary matrices for all layers

$j \neq j_0$, as well as, the source vectors for all layers $j \neq j_0$. When solving the cloudy-sky problem, we use the clear-sky layer equations for all layers $j < j_0$; for the layers $j > j_0$ we take account the change in attenuation of the direct solar beam. If the boundary of the cloud splits a layer, as shown in Figure 1c, then these corrections are applied to clear-sky layers situated above and below the split layers. For the split layers themselves, we must store the solutions to the homogeneous RTE obtained for the original clear-sky calculation.

The second approach is (for the cloudy scenario) to generate a spectral correction applied to the radiation field from a fast two-stream RTM. We propose the following computational formula for the multi-stream solution in a cloudy atmosphere

$$I_{\text{cloud}}(\lambda) \approx I_{\text{clear}}(\lambda) \frac{I^{\text{TS}}_{\text{cloud}}(\lambda)}{I^{\text{TS}}_{\text{clear}}(\lambda)} K(\lambda). \tag{6}$$

Here, $I^{\text{TS}}_{\text{cloud}}(\lambda)$ is the two-stream solution for the cloudy scenario, $I^{\text{TS}}_{\text{clear}}(\lambda)$ is the two-stream solution for clear-sky scenario and $K(\lambda)$ is the correction factor to be determined. Second, applying the dimensionality reduction techniques for computing the multi-stream solution for a clear sky, i.e., $I_{\text{clear}}(\lambda) \approx I^{\text{TS}}_{\text{clear}}(\lambda) f(\lambda)$, Eq. (6) becomes

$$I_{\text{cloud}}(\lambda) \approx I^{\text{TS}}_{\text{clear}} f(\lambda) \frac{I^{\text{TS}}_{\text{cloud}}(\lambda)}{I^{\text{TS}}_{\text{clear}}(\lambda)} K(\lambda). \tag{7}$$

With this approach, the computation of the multi-stream solution for a cloudy sky I_{cloud} requires only one additional call to of the two-stream solution $I^{\text{TS}}_{\text{cloud}}$ for each spectral point.

Correction factors $K(\lambda)$ are pre-computed for various values of the cloud parameters using the following inverted form of Eq. (6)

$$K(\lambda) = \frac{I_{\text{cloud}}(\lambda)}{I_{\text{clear}}(\lambda)} \frac{I^{\text{TS}}_{\text{clear}}(\lambda)}{I^{\text{TS}}_{\text{cloudy}}(\lambda)}. \tag{8}$$

Next , $K(\lambda)$ is interpolated in the spectral domain as

$$K(\lambda) \approx k\left(\lambda - \bar{\lambda}\right) + b + v\sigma_{\mathtt{O}_3}(\lambda) \tag{9}$$

where k, b and v are constants, the values of which are stored in look-up tables, $\bar{\lambda}$ is the mean wavelength for the spectral window (in this case $\bar{\lambda} = 330$ nm), $\sigma_{\mathtt{O}_3}$ is the $\mathtt{O}_3$ absorption cross section at temperature 270 K convolved with the GOME slit function to the 88-point wavelength grid in our 325-335 nm window. Although this method involves some approximations, it still provides radiance accuracy better than 0.2%, with a speed-up factor of approximately 2 compared with time taken for two separate RTM calls.

3.3. GPU-accelerated radiative transfer model

To optimize the loop over ground pixels, we designed a RTM code using the GPU architecture of modern graphical cards. To implement GPUs, the original CPU code has been redesigned using the C-oriented Compute Unified Device Architecture (CUDA) developed by NVIDIA. To reduce the CPU/GPU communication overhead, we exploited the asynchronous data transfer between host and device. To obtain optimal performance, we also used overlapping of CPU and GPU computations by distributing the workload between them.

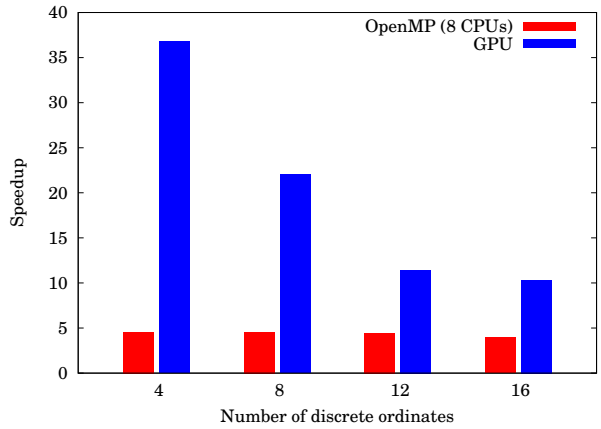

Fig. 2. Speed-up of the GPU-implemented DOM with the matrix operator technique compared to sequential CPU code execution.

Typically, values of $N_{\mathtt{do}} = 4 \div 8$ are chosen for simulations of scattered sunlight in the UV spectral range. The dimensions of matrices involved in the computations are mostly $N_{\mathtt{do}} \times N_{\mathtt{do}}$. Our numerical simulations regarding basic matrix operations evidence that for matrix sizes 8×8 the highest performance is achieved when all arrays required for the RTM solver are placed into registers. Speed-up factors are plotted as functions of the number of discrete ordinates in Figure 2. The speed-up for $N_{\mathtt{do}} = 16$ is less than that for $N_{\mathtt{do}} = 4$. With a low number of discrete ordinates, the kernel consumes a small number of registers, and so a large number of kernels can run simultaneously and the occupancy of GPU is high.

For the algorithm consisting of n parts with corresponding workloads W_i and speedups S_i, the total speedup for the whole algorithm reads as

$$S_{\mathtt{total}} = \left(\sum_{j=1}^{n} \frac{W_j}{S_j}\right)^{-1}. \tag{10}$$

Let's also introduce the "reduced workload" $\bar{W}_i$ as

$$\bar{W}_i = \frac{W_i/S_i}{\sum\limits_{j=1}^{n} \frac{W_j}{S_j}}. \tag{11}$$

Values of workloads, corresponding speedups as well as reduce workloads are given in Table 1 for $N_{\mathtt{do}} = 8$. For presented numbers, the theoretical speedup is $S_{\mathtt{total}} \approx 15$. In

Table 1. Workload of the PCA-based RTM and the corresponding speedup for $N_{\mathtt{do}} = 8$.

	Workload	Speedup	Reduced workload
Multi-stream RTM	50%	22	34.8%
Two-stream RTM	25%	53	7.2%
PCA	20%	6	51%
Rest	5%	∼10	6.12%

our computations, we obtain the speedup factor of $S_{\mathtt{total}} \approx 12$. Note that, PCA has the largest reduced workload. The main part of the PCA is the eigenvalue problem. According to Amdahl's law, the eigenvalue solver is a main limitation factor of the performance. PCA could be implemented on GPU. However, the standard eigenvalue solver from CULA library [7] shows poor performance for small matrices (see benchmarks at www.culatools.com). Moreover, it cannot support the computations in the batched mode.

With GPUs (Tesla K20 with 2496 cores), we achieved a 20x-40x speed-up for the multi-stream RTM, and 50x speed-up for the two-stream RTM, these figures with respect to performance with the original single-threaded CPU codes run on Intel Xeon CPU E5-1620 3.60GHz. The speed-up of the PCA-based RTM is of about 12 times.

4. CUMULATIVE PERFORMANCE ENHANCEMENT

Above mentioned acceleration techniques have been implemented in a common framework. The resulting code has been validated against the codes DISORT [4] and LIDORT [8]. The error imposed by the acceleration techniques is less than 0.1% in the spectral radiances. The code has been applied to the problem of ozone retrieval. The performance enhancement for considered techniques is given in Table 2. The obtained cumulative performance enhancement is of about 300 times which is 85% of a theoretical maximum estimated as a product of speed-up rates of all methods. The performance enhancement excluding GPU computations is of about 25 times. Our analysis shows that the considerable speed-up can be achieved by tuning and optimizing the RTM code to a specific remote sensing problem. It is required to make a lot of substantial changes to the underlying codebase to make it efficient. These changes affect the memory organization, data flows, the generation of appropriate look-up tables, the data compression algorithms and sparse matrix computations.

5. ACKNOWLEDGMENT

Part of this work was supported by the Bavarian Ministry of Economic Affairs and Media, Energy and Technology grant 07 03/893 73/ 5 /2013.

Table 2. The performance enhancement for the acceleration techniques

Acceleration technique	**Speed-up rate**
Dimensionality reduction	8
Parallel computing on GPU	12
Optimization of cloudy scenarios	1.9
Telescoping	1.5
Left eigenvectors	1.3
Cumulative speed-up	320

6. REFERENCES

[1] D. Efremenko, A. Doicu, D. Loyola, and T. Trautmann, "Acceleration techniques for the discrete ordinate method," *J Quant Spectrosc Radiat Transfer*, vol. 114, pp. 73–81, January 2013.

[2] D.S. Efremenko, A. Doicu, D. Loyola, and T. Trautmann, "Optical property dimensionality reduction techniques for accelerated radiative transfer performance: Application to remote sensing total ozone retrievals," *J Quant Spectrosc Radiat Transfer*, vol. 133, pp. 128–135, 2014.

[3] D.S. Efremenko, D.G. Loyola, A. Doicu, and R.J.D. Spurr, "Multi-core-CPU and GPU-accelerated radiative transfer models based on the discrete ordinate method," *Computer Physics Communications*, vol. 185, no. 12, pp. 3079 – 3089, 2014.

[4] K. Stamnes, S.C. Tsay, W. Wiscombe, and K. Jayaweera, "Numerically stable algorithm for discrete-ordinate-method radiative transfer in multiple scattering and emitting layered media," *Appl Opt*, vol. 12, pp. 2502–2509, 1988.

[5] V. Natraj, R.L. Shia, and Y.L. Yung, "On the use of principal component analysis to speed up radiative transfer calculations," *J Quant Spectrosc Radiat Transfer*, vol. 111, no. 5, pp. 810–816, 2010.

[6] D.S. Efremenko, D. Loyola, R.J.D. Spurr, and A. Doicu, "Acceleration of radiative transfer model calculations for the retrieval of trace gases under cloudy conditions," *J Quant Spectrosc Radiat Transfer*, vol. 135, pp. 58–65, 2014.

[7] J.R. Humphrey, D.K. Price, K.E. Spagnoli, A.L. Paolini, and E.J. Kelmelis, "CULA: hybrid GPU accelerated linear algebra routines," *Proc. SPIE*, vol. 7705, pp. 770502–770502–7, 2010.

[8] R.J.D. Spurr, "LIDORT and VLIDORT. Linearized pseudo-spherical scalar and vector discrete ordinate radiative transfer models for use in remote sensing retrieval problems," in *Light scattering reviews*, A.A. Kokhanovsky, Ed., vol. 3, pp. 229–275. Springer, 2008.

IMAGEQUERYING – EARTH OBSERVATION IMAGE CONTENT EXTRACTION & QUERYING ACROSS TIME AND SPACE

Dirk Tiede[1], Andrea Baraldi[1,2], Martin Sudmanns[1], Mariana Belgiu[1], and Stefan Lang[1]

[1]Department of Geoinformatics – Z_GIS, University of Salzburg, Austria
[2]Department of Agricultural and Food Sciences, University of Naples Federico II, Portici (NA), Italy

ABSTRACT

ImageQuerying (IQ) is an innovative incremental system for automatic near real-time image content extraction and content-based image retrieval (CBIR) in big Earth observation (EO) image databases. The IQ is built upon a hybrid (combined deductive and inductive) inference engine for low- and high-level EO image understanding, independent of the multi-spectral (MS) imaging sensor. Each MS image or image time-series stored in the EO database is automatically provided with at least one (or more) thematic layer(s), consisting of semantic image-objects in compliance with the increasingly popular object-based image analysis (OBIA) paradigm. Through the IQ graphic user interface, a user can: (1) perform semantic CBIR, where the comprehensive archive of EO image-derived thematic layers is queried rapidly; (2) develop original decision rules, where spatial and temporal operators are input with available image-derived qualitative (categorical) or quantitative (numeric) information layers. New thematic maps generated by user-defined decision rules steadily expand the knowledge base of the incremental IQ system.

***Index Terms*—**Earth observation (EO) image, big data, semantic querying, automatic real-time image understanding, content-based image retrieval.

1. INTRODUCTION

To date a very low percentage (estimated as ≤ 10%) of the Earth Observation (EO) images available in EO big data archives, made accessible free-of-cost by space agencies and public institutions, is ever downloaded for use. It means that the remote sensing (RS) community is currently overwhelmed by sensory data it is unable to transform into operational, timely and comprehensive knowledge/ information products. In conventional EO image databases, image retrieval is limited to querying metadata information (e.g., acquisition time, lat-long coordinates of the depicted Earth surface), quality flags and summary statistics, if any, e.g., overall cloud coverage, and/or by allowing a potential user to look at an image-specific "thumbnail" preview. Existing EO big data archives are provided with no EO image understanding intelligence, to overlay images with their own content layers. Consequently, there is no opportunity for a user to semantic querying an EO big data archive based on available image-content layers. In practice, vast amounts of information-as-data-interpretation are "hidden" in big EO image databases of ever-increasing quality and quantity, whose information richness and complexity increases monotonically with new spaceborne missions, such as Sentinel-2.

Figure 1. Top: Sentinel-2A (S2A) image of Austria, acquired on 2015-08-18. Bottom: Automatic mapping of the S2A image onto a legend of 96 color names (spectral categories), depicted as pseudo colors (green as vegetation, blue as water or shadow, etc.), by a prior knowledge-based SIAM vector quantizer.

To tackle the challenge and opportunities of a steady and massive acquisition of multi-source spaceborne/airborne EO images in EO big data archives, we are presenting ImageQuerying (IQ), an innovative automatic near real-time EO image understanding and semantic querying system. Built upon a hybrid (combined deductive and inductive) inference system for low- and high-level computer vision, IQ provides each EO image stored in the database with at least one (or more) content map(s) generated automatically.

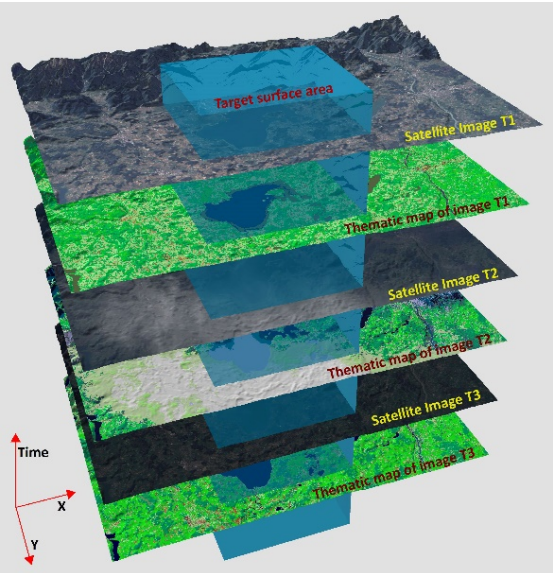

Figure 2. Overview of the IQ's EO image and image-derived map querying system for a user-defined AOI through time. For each EO image stored in the database, there is at least one (or more) image-derived thematic layer(s) automatically generated by IQ. In practice, concise content layers overlap with massive raw images to enhance high-level information queries.

In agreement with the human visual system which leads toward symbolic reasoning, IQ can track through time pixels and image-objects (segments) as combined trajectories of semantic labels, geometric (shape and size) properties and inter-object spatial relationships. The IQ graphic user interface (GUI) allows a user to interactively augment the IQ knowledge base, consisting of either categorical (nominal, qualitative) or quantitative image-derived variables, by querying the existing database of thematic layers, quantitative variables (e.g., leaf area index, spectral indexes, etc.) and/or raw images as operands. These operands are: (i) combined through algebraic, relational, and spatial (e.g., proximity) and/or time-domain (e.g., overlay) operators, (ii) subject to integrity constraints, either static or temporal, generated from a semantic network, i.e., a conceptual model of geospatial entities and relationships through time [1].

2. METHODS

To fill the information gap from sub-symbolic pixels to symbolic image-objects, the IQ's hybrid inference engine comprises the following library of automatic algorithms (parameter-free, requiring no user supervision to run) for low-level vision applications. (i) A battery of sensor-specific EO image radiometric calibrators, e.g. for Sentinel-2A, Sentinel-3, Landsat-5 to Landsat-8, MeteoSat, RapidEye, etc., required to transform digital numbers into a radiometric unit of measure, specifically, top-of-atmosphere reflectance (TOARF) (ii) A novel algorithm for image color constancy (color normalization), inspired by human vision. It is typically applied to non-calibrated multi-spectral (MS) images, such as those acquired by consumer-level RGB cameras mounted onboard Unmanned Aerial Vehicles (UAVs). (iii) An innovative fully automatic panchromatic/MS image segmentation algorithm, consistent with human vision and the OBIA paradigm. (iv) An expert system (prior knowledge-based decision tree) capable of partitioning a radiometrically calibrated MS data space into a pre-defined dictionary of color names. Known in recent literature as the Satellite Image Automatic Mapper (SIAM) [2], it transforms a radiometrically calibrated MS image into a preliminary classification map, whose legend is a dictionary of up to 96 color names (spectral categories, semi-symbolic labels), see Fig. 1. Noteworthy, spectral categories cannot always be inverted to unique land cover (LC) class names. For example, e.g., spectral category "*MS green*" is related to the land surface type *vegetation*, which includes many LC classes such as *deciduous forest*, *evergreen forest*, *grassland*, etc. Spectral category "*MS blue*" is related to either land surface types either *water* or *shadow*, etc. (v) An expert system, called RGB Image Automatic Mapper (RGBIAM), capable of partitioning an RGB data cube into a pre-defined dictionary of color names. (vi) A bi-temporal SIAM-based post-classification change/no-change detector.

Starting from the information primitives generated as output by these standard information processing blocks, IQ automatically extracts both quantitative and categorical (nominal) variables (incl. geometry information of the pre-classification segments) from each single-date MS image stored in the EO database. In greater detail, the legend of a standard categorical map automatically generated by IQ from a single-date EO image complies with the initial Dichotomous Phase of the Land Cover Classification System (LCCS) taxonomy, adopted by the Food and Agriculture Organization of the United Nations (FAO) [5]. Eight major LC types are distinguished as a combination of three-dichotomous mapping criteria: (i) vegetation/non-vegetation, (ii) terrestrial/aquatic, (iii) managed/natural or semi-natural. The eight dichotomous LC types are listed below (see also Fig. 3):

doi: 10.2788/854791

(1) Cultivated and Managed Terrestrial (non-aquatic) Non-vegetated Areas
(2) Natural and Semi-Natural Terrestrial Vegetation
(3) Cultivated Aquatic or Regularly Flooded Areas
(4) Natural and Semi-Natural Aquatic or Regularly Flooded Vegetation
(5) Artificial Surfaces and Associated Areas
(6) Bare Areas
(7) Artificial Waterbodies, Snow and Ice, and
(8) Natural Waterbodies, Snow and Ice.

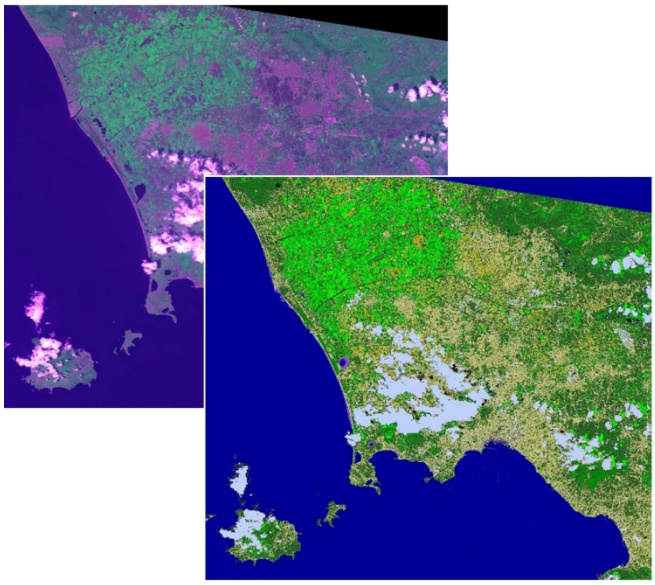

Figure 3. Top: 4-band (B, G, R, NIR) ALOS AVNIR-2 image of Campania, Italy, radiometrically calibrated into TOARF values and depicted in false colors (R = MIR, G = NIR, B = Visible Blue), 10 m resolution. Bottom: Automatic mapping of the image onto 10-class classification map based on a convergence-of-evidence approach, in compliance with the FAO-LCCS 3-level Dichotomous Phase.

The IQ's GUI allows the user to access, select and combine the EO image-derived quantitative and qualitative information layers available in the database to accomplish CBIR or geospatial data analysis applications, constrained by a user-defined information pair of an area of interest (AOI) and a target time (TT) interval, see Fig. 2.

To accomplish efficient geospatial data analysis through time within a user-defined AOI and TT, the current IQ implementation adopts an array-DBMS (specifically, Rasdaman, see below) within a client-server solution together with a web-based query interface.

3. IMPLEMENTATION

IQ is designed as a scalable, web-based client-server architecture and deployed on a Red Head Enterprise Linux machine with Rasdaman [3][4] as main component. It stores a set of multi-sensor EO images radiometrically calibrated into TOARF values - in the current implementation Landsat 5/7/8 images, but not limited to - plus their image-derived thematic layers.

Unlike standard relational databases, the images are stored within the array-DBMS using an array data model with semantically annotated dimensions. The definitions of the dimensions are based on OGC standards and are accessible using an URI. Thus, the data can also be accessed using the OGC WCS standard. The IQ-importer converts the individual images (flat files) into the IQ data model consisting of tiled multi-dimensional discrete data (MDD) cubes with the X and Y dimensions as spatial coordinates and the Z dimension as temporal coordinates, see Fig. 2. According to the three different sources of input information, specifically, spaceborne/airborne images, categorical variables and continuous variables, there are three different MDD cubes. Thus, the cubes will grow in the Z (temporal) dimension with each new image of the same area.

With this setup, IQ encompasses a tailored storage strategy, which is optimized for querying and analyzing the data through space and time. The user performs ad-hoc queries by graphically combining a set of predefined query elements, which consists of one or more array operations such as slicing or aggregations along the temporal dimensions as well as predefined and more complex user-defined functions (UDF).

The main graphical user interface (GUI) of IQ is a web application, which is accessible with a standard browser. However, a unified API of the IQ backend allows simultaneous access by multiple clients (e.g., web application, ArcGIS Toolbox) which are suitable for different type of workflows. Fast response times of IQ allow an interactive use of the system and, allow creating the rules easily and refining them interactively. Even complex queries and rules can be conducted easily within the comprehensive and scalable query panel, which connects the queries and their results visually.

Since the rules are encoded as XML in the IQ frontend, they can be parsed by the IQ backend for security and syntax checks and are subsequently converted into safe database queries. Optionally, the rules including its metadata can be stored as XML in the knowledge base.

In general, human expertise is required to bridge the semantic gap from low- to high-level image-derived information as a combination through time and space of low-level semantic image-objects. The IQ's GUI supports the user to develop a feedforward (top-down) knowledge transfer paradigm (from humans to the IQ system), followed by a feedback (bottom-up) information-from-data retrieval phase. To this end, the IQ's GUI comprises a graphical query mode (see Fig. 4) supporting (i) the graphic selection of existing geospatial semantic queries and (ii) the graphic

generation of new geospatial semantic queries. The application domain of geospatial semantic queries is twofold: (1) CBIR. For example, retrieve EO images that are cloud-free across the AOI, or those where a specific vegetation type is found in the AOI, etc. (2) Generation of new information layers. For example, detect flooded areas as a post-classification combination through time of available single-date thematic maps; detect-through-time evergreen vegetation; detect clear cuts as a vegetation cover change in time, etc.

4. RESULTS AND DISCUSSION

The current IQ implementation allows expert users to build and save complex semantic queries of the EO big database investigated through time and space. Non-expert users are enabled to select any of the predefined queries according to their target sensor, spatial and temporal intervals. Examples of semantic queries implemented during our IQ system test phase are listed below.

- Retrieve all images with cloud cover < x % in the selected area-of-interest (AOI) and target-time (TT) window.
- In the selected AOI and TT, detect all areas labelled as evergreen forest, etc.
- Detect changes in an AOI trough time, such as deforestation phenomena, flooded areas, burned areas, etc.
- Extract lakes as water areas with a certain combination of fuzzy memberships in geometric properties, such as size, compactness etc. (see Figure 4).

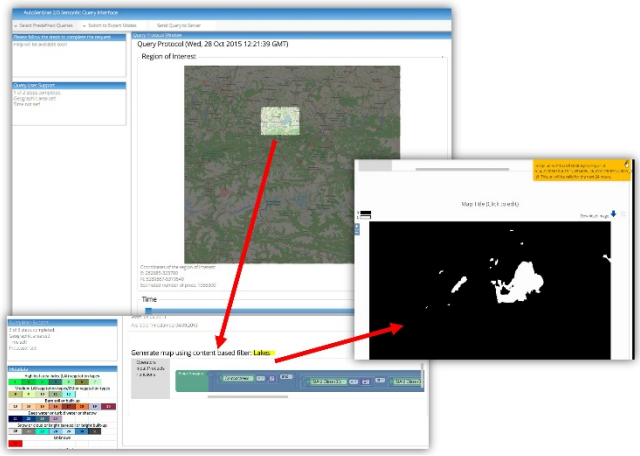

Figure 4. Web-based GUI implemented in a test IQ realization where a Landsat-8 time-series is analyzed by a semantic query to extract lakes as water areas of a certain size and compact form from the database of EO images overlaid with image-derived content maps.

The IQ's GUI allows user-defined semantic queries to be saved and repeated on different AOIs and TTs. Results of a query employed for CBIR purposes is a list of one or more image, quantitative variable and/or qualitative variable stored beforehand in the EO database. Results of a query employed for image analysis are new information layers, either quantitative or qualitative, which can be saved as georeferenced layers to be employed as input variables by further queries of the incremental IQ system.

The tested IQ implementation was found to be computationally efficient, in both the CBIR and the image-content extraction operating mode. Building blocks of the IQ system were selected and developed to be scalable and extensible (e.g., automatic and near real-time low- and high-level vision capabilities, array database system, capability to address multiple clients). Testing the IQ system upon large-scale EO image databases will be the subject of our future research and IQ system development.

5. ACKNOWLEDGEMENTS

The research has received funding from the Austrian Research Promotion Agency (FFG) under the Austrian Space Application Programme (ASAP) within the project AutoSentinel-2/3 (contract no: 848009). The test implementation of the IQ system was awarded with the *Copernicus Masters 2015 - T-Systems Big Data Challenge.*

6. REFERENCES

[1] Cabot, J., Olivé, A., & Teniente, E. "Representing temporal information in UML". In: *«UML» 2003-The Unified Modeling Language. Modeling Languages and Applications*. Springer Berlin Heidelberg, pp. 44-59, 2003.

[2] A. Baraldi, L. Durieux, D. Simonetti, G. Conchedda, F. Holecz, and P. Blonda, "Automatic spectral rule-based preliminary classification of radiometrically calibrated SPOT-4/-5/IRS, AVHRR/MSG, AATSR, IKONOS/ QuickBird/ OrbView/ GeoEye and DMC/SPOT-1/-2 imagery. Part I: System design and implementation," *IEEE Trans. Geosci. Remote Sens*., vol. 48, no. 3, pp. 1299-1325, 2010.

[3] Baumann, P., Dehmel, A., Furtado, P., Ritsch, R., & Widmann, N. "The multidimensional database system RasDaMan". In: *ACM SIGMOD Record*, vol. 27, no. 2, pp. 575-577, 1998.

[4] Baumann, P. "Accelerating computationally intensive queries on massive earth science data: (system demonstration)". *Proceedings of the EDBT/ICDT 2011 Workshop on Array Databases*. ACM, pp. 31-35, 2011

[5] Di Gregorio A. and Jansen L., "Land Cover Classification System (LCCS): Classification Concepts and User Manual", FAO: Rome, Italy, FAO Corporate Document Repository, 2000. [Online]. Available: http://www.fao.org/DOCREP/003/X0596E/X0596e00.htm

A PILOT FOR BIG DATA EXPLOITATION IN THE SPACE AND SECURITY DOMAIN

Sergio Albani (1), Michele Lazzarini (1), Manolis Koubarakis (2), Efi Karra Taniskidou (2), George Papadakis (2), Vangelis Karkaletsis (3), George Giannakopoulos (3)

(1) European Union Satellite Centre, Apdo de Correos 511, 28850 Torrejón de Ardoz, Spain
(2) University of Athens, University Campus Ilisia, 15784 Athens, Greece
(3) National Centre for Scientific Research 'Demokritos', 15310 Aghia Paraskevi Attikis, Greece

ABSTRACT

In the framework of the Horizon 2020 project BigDataEurope (Integrating Big Data, Software & Communities for Addressing Europe's Societal Challenges), a platform comprising key open-source technologies has been set up in order to meet the Big Data requirements of seven communities representing the Horizon 2020 Societal Challenges (Health, Food and Agriculture, Energy, Transport, Climate, Social Sciences and Secure Societies).

The BigDataEurope platform is currently validated by implementing relevant pilots; for the Secure Societies challenge particular importance has been given to the integration and fusion of data coming from remote and social sensing in order to add value to the current data exploitation practices.

Index Terms *- Earth Observation, Big Data, heterogeneous data sources, multi-temporal and multi-sensor analysis, space and security*

1. INTRODUCTION

The rapidly increasing amount and variety of data coming from satellites and other sources is raising new issues such as the management and exploitation of extremely large and complex datasets (Big Data); the main challenge in the Space and Security domain is to improve the capacity to extract in a timely manner operational (i.e. useful and clear) information from a huge amount of heterogeneous data.

A number of public and private initiatives are taking place at European level with the ambition of taking advantage of the opportunities offered by Big Data technologies. The Horizon 2020 BigDataEurope[1] project (Integrating Big Data, Software & Communities for Addressing Europe's Societal Challenges) aims at providing support mechanisms for all the major aspects of the data value chain in terms of employed data and technology assets, the participating roles and the established or evolving processes.

BigDataEurope is focusing on two coordination and support measures:

- Engaging with a diverse range of stakeholder groups representing the Horizon 2020 Societal Challenges[2] Health, Food & Agriculture, Energy, Transport, Climate, Social Sciences and Secure Societies.
- Collecting requirements for the ICT infrastructure needed to design, realize and evaluate a Big Data Aggregator platform infrastructure that comprises open source technologies in the framework of the lambda architecture [1].

The Secure Societies Societal Challenge has been defined for the protection of freedom and security of Europe and its citizens. Key aims of this challenge are to enhance the resilience of our society against natural and man-made disasters, to develop novel solutions for the protection of critical infrastructures, to improve border security and to support the Union's external security policies; a major activity in supporting these aims is the provision of geospatial products and services, mainly resulting from satellite data.

The pilot for Secure Societies implemented to validate the BigDataEurope platform is focusing on the integration and fusion of data coming from remote and social sensing in order to add value to the current data exploitation practices; this is key in the Space and Security domain, where useful information can be derived not only from satellite data but also from data coming from social media and other sources.

[1] BigDataEurope has received funding from the Horizon 2020 programme under grant agreement n° 644564

[2] https://ec.europa.eu/programmes/horizon2020/en/h2020-section/societal-challenges

2. Pilot Description

According to the user requirements collected from a number of stakeholders in the Secure Societies domain during the first phase of the BigDataEurope project, there is a need for exploiting the increasing amount of data coming from space and other sources, with a major contribution of open data and tools. Automatic tools for data management and processing are one of the key aspects, where the adopted solutions have to be integrated in the whole data chain in order to reduce the human effort, to reduce the amount of needed economic resources and to efficiently perform data analysis.

The BigDataEurope pilot for Secure Societies has been developed in the Space and Security domain following the above requirements; it covers all the issues related to Big Data, namely Volume (large satellite images), Variety (heterogeneous data such as satellite images and textual content), Velocity (fast-paced social data and news stream), Veracity (cross-verification of the sources) and Value (adding useful information).

The pilot considers two different workflows of data:

- The first workflow, called the **change detection workflow,** ingests satellite images to detect areas with changes on land cover or land use by using change detection techniques; the identified Areas of Interest are then associated with social media and news items and presented to the end-user for cross-validation.
- The reverse procedure is applied to the second workflow, called the **event detection workflow**. Event detection is triggered by news and social media information, where trending topics (i.e. document clusters) with geospatial connotation constitute a time- and space- localized event; provided such event, the corresponding satellite images are acquired and processed in order to check for changes in land cover or land use.

The pilot has been developed on the basis of existing tools targeted to expert users and suitable for small-scale (i.e. serial) processing; in the context of the pilot, the functionality of these tools becomes fully automated and parallelized according to parallelization and distributed storage principles for higher throughput. For remote sensing, specific tools for image (pre-) processing have been adapted from the Sentinel Application Platform (SNAP)[3]. For social sensing, a set of crawling and machine learning tools lying at the core of the NewSum summarization [2] application were employed; these tools involve specialized techniques for gathering news items and social media data from the Web and for effectively clustering them into events using text mining methods.

3. Pilot Architecture

The architecture of the pilot was designed to accommodate both workflows and consists of the three components depicted in Figure 1: the user interface, which is a modified version of the web-application Sextant [3]; the storage component, which involves Apache Cassandra[4] and Strabon [4]; the core component, which consists of two modules, one for change detection in land cover / land use from satellite images and one for event detection in news items and social media.

The first component runs on the client-side (i.e. on the user local computer) and constitutes a web application that can be deployed on a variety of platforms. The last two components run on the server-side (i.e. on the BigDataEurope infrastructure) in order to offer scalability and high efficiency.

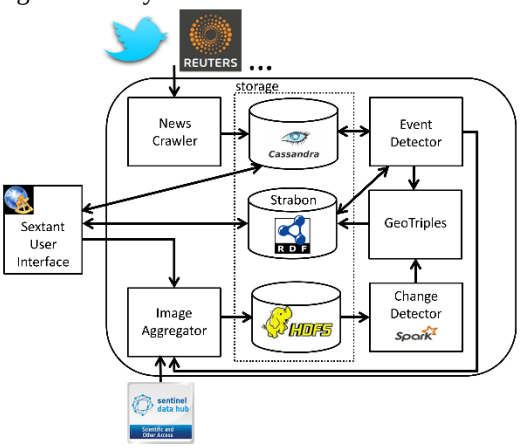

Figure 1: Pilot architecture

For the **change detection workflow**, Sextant offers an interface similar to Google Earth[5], allowing users to select an Area of Interest by forming a rectangle on an Earth map or inserting the geographic coordinates of such area. In both cases, the coordinates of the specified area are extracted and forwarded to the core component of the pilot. The user has the possibility to determine the dates of interest, i.e. one date before and one after the change should have taken place.

The Image Aggregator receives this information as input and issues the appropriate query to the Sentinel Scientific Data Hub[6] repository of satellite images. The parameters of the query are automatically configured, ensuring that only images (within the specified dates) with specific characteristics (e.g. resolution, polarization, orbit and view angle) are downloaded and stored as one file in the local file system.

The selected images are then compared by the Change Detector (implementing the functionalities offered by the ESA SNAP toolbox) in order to identify areas with changes in land cover or land use. At the moment, the

[3] http://step.esa.int/main/toolboxes/snap

[4] http://cassandra.apache.org

[5] https://www.google.com/earth

[6] https://scihub.copernicus.eu

Change Detector is only able to perform a binary classification, distinguishing areas into changed and unchanged ones; in the future it will be enriched with the ability to perform a classification in order to identify the type of change (e.g. an area with forest which becomes urban or an area with water where a manmade structure is built).

In this process, GeoTriples [5] is used in order to convert all the information that have been produced by the Change Detector (e.g. coordinates and type of change) into RDF statements (triples). Subsequently, GeoTriples stores the resulting triples into Strabon, an effective spatiotemporal RDF store that supports the main extensions of SPARQL, offering a large variety of queries over evolving geospatial data.

Finally, Sextant retrieves the stored geocoded areas from Strabon and presents them to the user. It also gets from Strabon the textual events identifiers related to the coordinates and geo-locations of the change. Then it can retrieve from the storage of the event detection workflow (Apache Cassandra) the event content, essentially links to the original news items and content. This is done through a SPARQL query in Strabon, where the crawlers and clusterers of the textual information gather items on a near real-time fashion, updating Strabon with any events detected. Consequently Strabon has all the information to answer queries related to historic data and to current events (e.g. within the last day), given a geo-location specification.

For the **event detection workflow,** the input is automatically given by the news stream, i.e. the RSS feeds from Reuters[7] and the status updates from Twitter which are continuously gathered from the Web by the News Crawler. The Twitter crawler (or *listener*) follows either the Twitter Stream API, or can use using specific keywords (Search API) or user accounts (e.g. news agencies).

The gathered content of the news stream is regularly processed at specific short time intervals by the Event Detector in order to cluster news items into events, based on topic clustering (i.e. news items talking about the same topic form an event cluster). Named Entity Recognition (NER) is used to determine the locations mentioned in the news items as well as in tweets and events; these locations and any explicit location annotation by the news publisher are mapped to geocoded locations through Strabon. Users can specify date and location for retrieved events and, in the future, also topics by appropriate keywords. The geocoded Areas of Interest that correspond to the detected events are forwarded to the change detection workflow in order to verify the events in the satellite images. The rest of the workflow operates as described above, with the changes (if verified) displayed to the user via Sextant.

[7] http://www.reuters.com

4. DATA

Remote and social sensing data have been selected for the pilot to address the fusion of information from heterogeneous sources.

For remote sensing, Sentinel-1 (a satellite launched in 2014 carrying a C-band Synthetic Aperture Radar) Level-1 GRD (Ground Range Detected) images were chosen for this first implementation phase. Sentinel-1 will benefit, among others, services related to: monitoring land-surface for motion risks; mapping for forest, water and soil management; mapping to support humanitarian aid and crisis situations [6]. The data acquisition can be configured with different modes: the default mode over land has a swath width of 250 km and a ground resolution of 5 x 20 m, while the Stripmap mode provides a continuity of ERS and Envisat data offering a 5 x 5 m resolution over a narrow swath width of 80 km. The access and use of Copernicus Sentinel Data and Service Information is regulated under EU law; the free, full and open data policy adopted for the Copernicus programme foresees access available to all users for the Sentinel data products, via a simple pre-registration on the Sentinel Scientific Data Hub.

With regard to social sensing, the focus is on two complementary sources of information: social media, represented by Twitter, and news agencies, represented by Reuters. The former involves plain text (Twitter messages) along with metadata in JSON format, while the latter includes news articles that are made available through RSS feeds in XML format. Recently, Twitter has emerged as a major platform for on-time detection of events, both in research and in industry (e.g. [7]). This pilot uses the free Twitter Public Streams API[8], which provides a random sample of its content. Twitter provides many options for parameterization, thus allowing the pilot to make the most of the retrieved content even if this is a subset of the tweets generated. For example, broad Areas of Interest can be specified, or specific users can be monitored without limitation. No private statuses are retrieved by the Public API. The content provided by news agencies through public RSS feeds is also free of charge. Within the pilot the content is used for the extraction of metadata, while the user is pointed to the original source for more information, to abide by intellectual property rights requirements. Given that all major news agencies release their content as RSS feeds, the pilot approach can be extended straightforwardly to more public content.

5. PRELIMINARY RESULTS

The main effort in this stage of the pilot deployment has been dedicated to the change detection workflow and, in particular, to the Change Detector module. In fact this

[8] https://dev.twitter.com/streaming/public

module has to process in each user request at least two very large satellite images, which occupy several GBs even after compression. The Change Detector module implements in the pilot the functionalities offered by the ESA SNAP toolbox both for the pre-processing step (i.e. co-registration) and the change detection algorithm, achieving high effectiveness in terms of precision and recall. Nevertheless this implementation still requires a significant amount of time for processing a pair of images as a specific parallelization approach to increase efficiency has still to be developed.

To scale the requirement of serving numerous user requests at the same time in the context of Big Data, the functionality of the Change Detector has been parallelized following the straightforward approach that is implemented in Calvalus[9] by the developers of SNAP: a separate node using the MapReduce framework has been assigned to every set of images. This means that if the pilot runs on a cluster with N nodes, this approach is able to process N-1 different user requests in parallel (as one of the nodes operates as the master). The only difference is that the current state-of-the-art in MapReduce, namely Apache Spark[10], has been used, whereas Calvalus employs Apache Hadoop[11] for this purpose. For the event detection workflow, the News Crawler uses the Cassandra storage to persist data and metadata, while the NewSum clustering functions over Apache Spark.

6. Conclusions And Future Works

In the framework of the Horizon 2020 BigDataEurope project, a pilot for Big Data exploitation in the Space and Security domain (deployed on a Big Data platform) has been presented. The pilot considers the fusion and analysis of information from remote sensing (satellite data) and social sensing (news from Reuters and Twitter), where the analysis of satellite images to detect areas with changes on land cover or land use (e.g. construction of critical infrastructures or exploitation of natural resources) can be associated with information provided by social media and news items.

The next steps of the deployment will be focused on the optimization of the whole data management chain. A more elaborate parallelization approach adapting each individual stage in the functionality of the Change Detector to the MapReduce framework will be developed. In this way, the processing time for a pair of images will be reduced to a significant extent that is almost linear to the number of available resources/nodes. This approach involves higher network overhead, due to the distribution of tiles among the available nodes, but its throughput is expected to be higher. Qualitative control on the final output will be performed, e.g. the verification of the change detection accuracy. In the future, once the workflows will be in use, the possibility to ingest other image processing algorithms (e.g. on the Change Detector module) and to process other types of images (e.g. Sentinel 1 Single Look Complex images and Sentinel 2 data) will be considered. Concerning the textual aspect of analysis, online text clustering and summarization methods using a distributed paradigm (e.g. using Apache Spark) will be evaluated in order to increase the throughput and scalability of the system (in alignment to the BigDataEurope aims). Finally, Sextant will be extended with a keyword-search functionality enabling users to filter the detected events according to topics of interest, in addition to the current filtering possibilities which rely on space and time constraints.

7. References

[1] N. Marz, and J. Warren, Big Data, *Principles and best practices of scalable realtime data systems*, Manning Publications, USA, 2014.

[2] G. Giannakopoulos, G. Kiomourtzis, and V. Karkaletsis, "NewSum: “N-Gram Graph”-Based." *Innovative Document Summarization Techniques: Revolutionizing Knowledge Understanding*, pp. 205-230, 2014.

[3] C. Nikolaou, K. Dogani, K. Bereta, G. Garbis, M. Karpathiotakis, K. Kyzirakos, M. Koubarakis, "Sextant: Visualizing time-evolving linked geospatial data", *Journal of Web Semantics*, pp. 35-52, 2015.

[4] K. Bereta, P. Smeros and M. Koubarakis, "Representing and Querying the Valid Time of Triples for Linked Geospatial Data", *Proceedings of the 10th Extended Semantic Web Conference (ESWC 2013)*, Montpellier, France. May 26-30, 2013.

[5] K. Kyzirakos, I. Vlachopoulos, D. Savva, S. Manegold, M. Koubarakis, "GeoTriples: a Tool for Publishing Geospatial Data as RDF Graphs Using R2RML Mappings", *International Semantic Web Conference (Posters & Demos)*, 2014.

[6] R. Torres, P. Snoeij, D. Geudtner, D. Bibby, M. Davidson, E. Attema, P. Potin, B. Rommen, N. Floury, M. Brown, I. Navas Traver, P. Deghaye, B. Duesmann, B. Rosich, N. Miranda, C. Bruno, M. L'Abbate, R. Croci, A. Pietropaolo, M. Huchler, F. Rostan, “GMES Sentinel 1 Mission”, *Remote Sensing of Environment*, 120, pp. 9-24, 2012.

[7] T. Sakaki, M. Okazaki, Y. Matsuo, "Earthquake shakes Twitter users: real-time event detection by social sensors", *Proceedings of the 19th international conference on World Wide Web*, New York, USA, 2010.

[9] http://www.brockmann-consult.de/calvalus
[10] http://spark.apache.org
[11] https://hadoop.apache.org

SCALABLE GEOSPATIAL SERVICES FOR THE PRODUCTION OF TIME SERIES AND VALUE-ADDED MAPS IN AGRICULTURE AND WATER QUALITY MONITORING

Athanasios Karmas and Konstantinos Karantzalos

Remote Sensing Laboratory
National Technical University of Athens
Zographou campus, 15780, Athens, Greece
thanasis.karmas@gmail.com karank@central.ntua.gr

ABSTRACT

In order to fully exploit the current open, high resolution, multispectral massive streams of satellite imaging data, the rapid online delivery of time series information and value-added maps is of significant importance for several applications. To this end, in this paper, a scalable and robust platform has been designed, developed and evaluated for the online and real-time harvesting of valuable information from earth observation big data. The developed system builds upon recently developed geospatial tools and can perform online, on the server-side data processing and time series information extraction. The core of our system is based on the *Rasdaman* Array Database Management System for data storage and the Open Geospatial Consortium Web Coverage Processing Service for data querying. Advanced remote sensing algorithms and open source libraries have been integrated, while the delivered value-added maps and products contribute to applications like precision agriculture and water quality monitoring.

***Index Terms*—** remote sensing, earth observation, satellite, big data, Landsat, Sentinel, WCPS

1. INTRODUCTION

Earth observation (EO) is shifting towards a new paradigm not only because of the fast growing volume and increasing variety of EO datasets but also due to the increasing complexity of algorithms required for their processing [1]. These algorithms must be validated against big datasets and competing algorithms, a task that, also, requires vast computing and human resources. Furthermore, to fully exploit these massive streams of EO data, high performance, scalable infrastructures as well as the implementation of novel processing chains and algorithms are required towards delivering validated value-added products for several geospatial applications.

In order to meet the aforementioned demands and ensure that users, software and products are brought closer to the data, both the public and private sector are seeking cutting-edge solutions towards efficient EO big data infrastructures. Such solutions can allow the direct and rapid data manipulation and analytics. Google Earth Engine[1] and Amazon S3 (AWS) for instance, offer access and processing capabilities for some of the most high resolution open datasets (*e.g.,* Landsat & Sentinel archive), thereby eliminating obsolete and time consuming network data transfer. Moreover, the Helix Nebula Science Cloud[2] brings together leading technology providers and research centres in order to provide computing capacity and services that are able to address the increasing demand for computing power. In order to exploit the available computing power and resources, various frameworks and tools like *Rasdaman*[3], MonetDB[4] and MrGeo[5] (NGA & DigitalGlobe) have been utilized for various earth science fields and GIS applications [2–5].

To this end, a scalable and robust platform has been designed, developed and evaluated for the online and real-time harvesting of valuable information from EO big data for precision agriculture and water quality monitoring applications. In particular, based on an automated pipeline regarding the data acquisition, storage and pre-processing, the developed platform can handle several open and commercial satellite datasets (*e.g.,* Landsat 8, Sentinel-2, RapidEye), while employing in (near) real-time, on the server-side several algorithms towards time-series analysis and value-added geospatial maps.

2. THE DEVELOPED GEOSPATIAL PLATFORM

The core of the developed system is based on *Rasdaman* Array Database Management System (DBMS) which has been already successfully validated in several studies *e.g.,* PlanetServer [2], EarthServer [6], RemoteAgri [7] and other [5], [8] with quite promising results. Expanding similar solutions [5, 7, 8] the overall system architecture comprises of a

[1] *https://earthengine.google.com/*
[2] *http://www.helix-nebula.eu/*
[3] *http://rasdaman.org/*
[4] *https://www.monetdb.org/Home*
[5] *https://github.com/ngageoint/mrgeo*

'Data Integration' stage as well as of the *Rasdaman* database that hosts the EO data. The system employs the OGC WCPS interface standard either for processing or for quering any dataset, while several developed 'Server-Side Scripts' manage the production of the value-added products. A Web Client forms the front-end of the framework and is responsible for handling all interactions with the user.

2.1. Data Integration

The system is able to seamlessly integrate various EO free & open access as well as commercial datasets of different spatial, spectral and temporal resolutions due to the abstraction layers offered by the *Rasdaman*'s data model. Currently, data from the US Landsat Data Continuity Mission (LDCM)[6] and ESA's Sentinels Scientific Data Hub[7] are systematically stored to our system. Data are downloaded, stored and pre-processed automatically through our system with the aid of several developed *python* scripts which control, facilitate and automate the entire operation. Specific multidimensional datacube structures have been also implemented for commercial satellite imaging datasets and in particular for Pleiades, Worldview-2, RapidEye and Deimos. After the pre-processing stages each scene is inserted in the *Rasdaman* database in a dedicated 3D data cube (one data cube per Worldwide Reference System (WRS) path/row defined scene) as a time slice which features a complex pixel type that serves the respective satellite's (*e.g.,* L8's) spectral bands. This architectural design pattern provides the potential for performing not only remote sensing analysis on just one image but also for time series analysis on multiple images simultaneously based on users's requests.

2.2. Geospatial Services

The developed framework provides in its current version geospatial services for the online and near real-time processing of EO big data (*i.e.* multispectral satellite imagery) for certain applications like precision agriculture and water quality monitoring. It is worth mentioning that a similar framework has been employed for (near) real-time, on the server-side land cover classification [8]. The system is capable of accurately detecting vegetated areas and estimating crop canopy greenness as far as precision agriculture is concerned. It is also highly effective in the detection of inland water bodies as well as in the estimation of inland chlorophyll concentration levels in regard to water quality monitoring. Moreover, the system exhibits potential for time series analysis as it is capable of providing multitemporal information for any remote sensing index across the stored datasets. The aforementioned services which are detailed in the following paragraphs are implemented as WCPS queries.

[6] *http://LANDSAT.usgs.gov/*
[7] *https://scihub.esa.int/dhus/*

```
for c in (path_row_data_cube) return encode(
    (float) coverage histogram over $pt t(low:high)
    values  aggregate_function( remote_sensing_index
        [Lat(low:high), Long(low:high),ansi($pt)]),
"csv")
```

Fig. 1. The developed WCPS query for the processing of multispectral satellite imagery from L8 satellite's OLI sensor towards the extraction of multitemporal values of remote sensing indices through the application of an aggregate function on various data cube's bands.

1) Color Composites, Band Ratios and Indices: For the purpose of visualizing the stored EO datasets the creation of color composites upon request is possible. Moreover, the online calculation of certain band ratios and standard broadband indices is also available.
2) Canopy Greenness Estimation: Based on the vegetation detection and the NDVI computation this service is delivering Canopy Greenness maps. Optical observations on LAI can be well correlated with vegetation indices like NDVI for single plant species which are grown under uniform conditions. However, for mixed, dense and multilayered canopies these indices have nonlinear relationships and can only be employed as proxies for crop-dependent vegetation parameters such as fractional vegetation cover, LAI, albedo, emissivity, etc. The service proceeds with a further classification for those areas that have been detected to contain vegetation towards estimating the different canopy greenness levels which can be associated with the vegetative canopy vigour, biomass, leaf chlorophyll content, canopy cover and structure.
3) Inland Chlorophyll Estimation: The concentration of the photosynthetic green pigment chlorophyll-*a* in inland water bodies is a proven indicator of the abundance and biomass of microscopic plants (phytoplankton) such as unicellular algae and cyanobacteria. Chlorophyll data are useful over a range of spatial scales for monitoring the water quality and environmental status of water bodies. We employ both the Blue and the Coastal spectral bands, as well as the Green one for lower absorption rates. For visualisation purposes the output is determined by zoning the different estimated chlorophyll levels.
4) Time series: This particular service leverages the platform's potential towards the accomplishment of time series analysis as it is capable of providing multitemporal information for any remote sensing index across the stored datasets. More specifically, the system is capable of calculating online and displaying in a time series chart multitemporal values of the NDVI index as well as of the chlorophyll-*a* concentration for any AOI that the user has specified. For the selected AOI and for each time slice of the 3 dimensional data cube that the AOI belongs to, the average value of the NDVI index or of the chlorophyll-*a* concentration is calculated separately for each time slice and is then returned to the Web Client of the

doi: 10.2788/854791

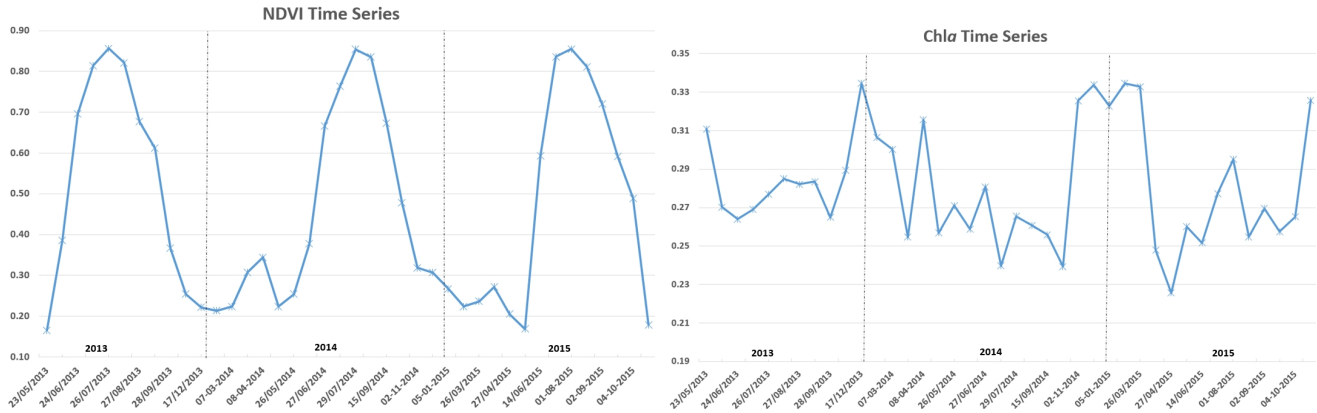

Fig. 2. Time series charts of the NDVI index (left) and chlorophyll-*a* concentration (Chla, right) after the application of the Time Series service on the L8 surface reflectance up to 20/10/2015 for respective AOIs from Figures 3 and 4.

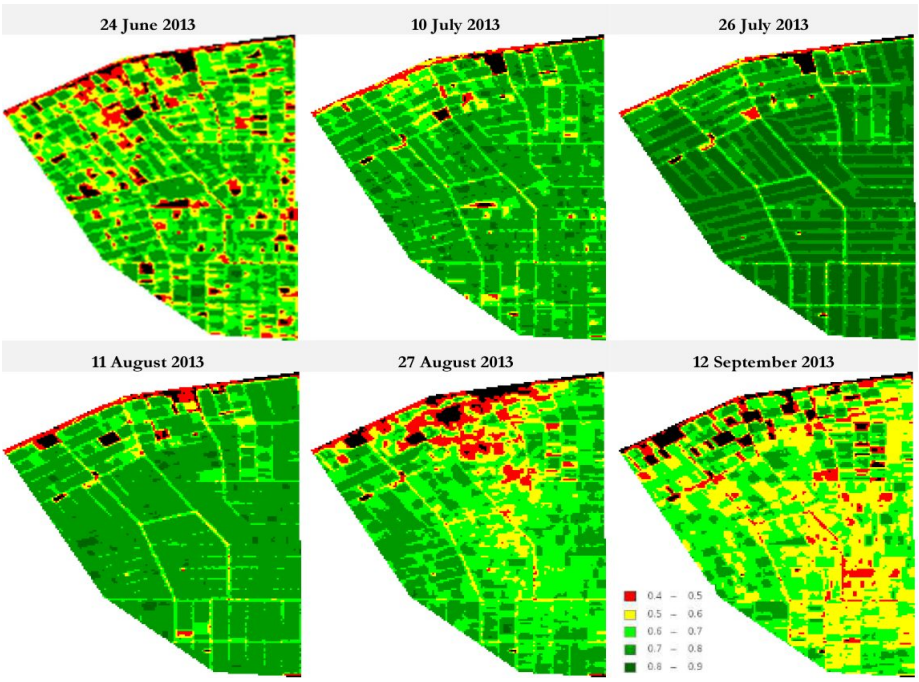

Fig. 3. Multitemporal crop canopy and vigorousness levels after the application of the Canopy Greenness Estimation service.

system which is responsible for displaying the resulting multitemporal values in a user-friendly and intuitive way through a time series chart. This chart can be used in conjunction with the Canopy Greenness maps and the Chlorophyll Estimation maps in order to estimate the potential impact of anomalies either on crop production or water quality respectively.

3. EXPERIMENTAL RESULTS AND EVALUATION

The experimental validation was performed by applying the developed services on numerous regions targeting both applications. In Figures 2, 3, 4 indicative results of the developed system are presented.

In Figure 2 results after the application of the Time Series service on the L8 surface reflectance dataset are presented. From the produced time series chart of the NDVI index it is relatively easy to monitor the life cycle of certain crops (here mainly rice) as well as to extract life cycle patterns for different types of crops and identify possible occurring anomalies. As the entire process of retrieving and making available for processing newly acquired remote sensing data from a satellite sensor, is almost fully automated, the timely addition of new points in the chart is feasible. This leads to the instantaneous identification of anomalies on crop production. Similarly, the produced time series chart of the chlorophyll-*a* concentration in an inland water body can be used to monitor fluctuations in water quality over time.

doi: 10.2788/854791

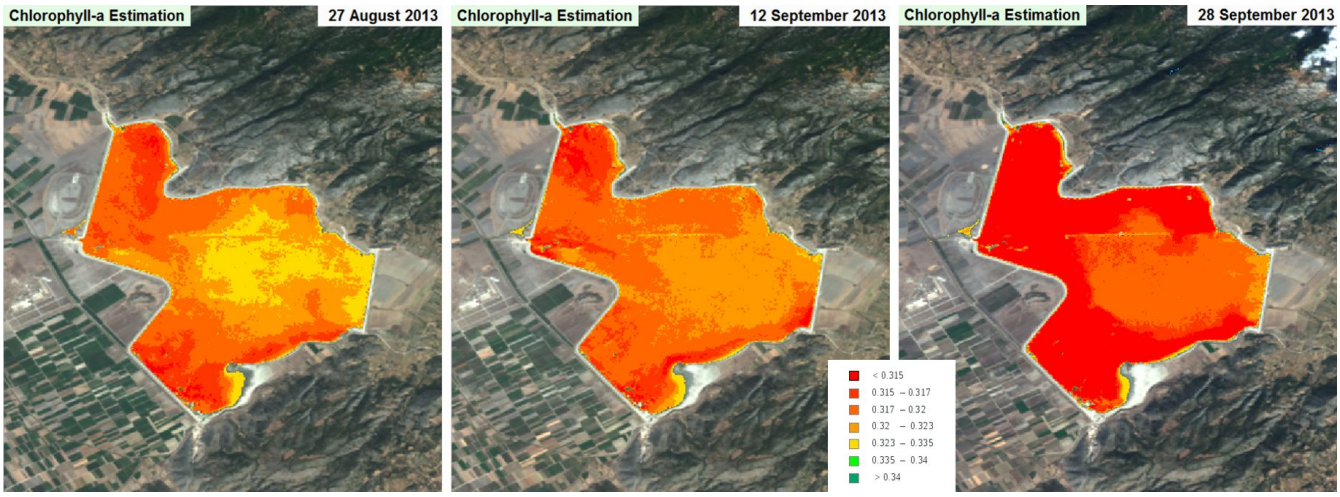

Fig. 4. Multitemporal chlorophyll-*a* concentration levels after the application of the Inland Chlorophyll Estimation service.

In Figure 3 results from 6 different acquisition dates are demonstrated for the Canopy Greenness Estimation service for an agricultural region located near the Axios Delta in the region of Central Macedonia, Greece. The summer rice crops are dominating the area (around 70%) while cotton and corn crops follow. As one can observe the most vigorous crops have been detected during late July which is in fact the period of peak vigorous state for most of the crops. This observation is in accordance with the time series chart of the NDVI index presented in Figure 2 for this area.

Figure 4 pictures results after the application of the Inland Chlorophyll Estimation geospatial service on a sensitive and shallow inland water body i.e., Lake Karla in Magnesia Prefecture, Greece. These maps are in accordance with the in-situ measurements which are periodically performed according to the EU water directives.

The following demonstration page offers a more comprehensive list of results and presents more extensively the overall functionality and performance of the developed system: http://users.ntua.gr/karank/Demos/GeoSp_Demo.html

4. CONCLUSION

In this paper, a framework capable of handling and analysing online diverse and multi-dimensional EO data was presented and evaluated. The framework establishes an automated workflow as well as a processing chain for the production of time series and value-added maps in agriculture and water quality monitoring.

Future work spans several different areas. An immediate goal consists of the reimplementation of the back-end of the platform with different frameworks like the ones from the Apache Ecosystem (*e.g., Hadoop, HBase, etc.*) and approaches (*e.g., NGA's MrGeo*). This action will enable us to perform extensive benchmarking against the current system configuration that features *Rasdaman* as its core platform in order to fully determine in action the advantages and disadvantages of each technology and computational scheme when manipulating and processing EO big data. Moreover, towards improving significantly the response time of the developed geospatial services, exploitation of GPUs computing capabilities is a top priority.

5. REFERENCES

[1] Yan Ma, Haiping Wu, Lizhe Wang, Bormin Huang, Rajiv Ranjan, Albert Zomaya, and Wei Jie, "Remote sensing big data computing: Challenges and opportunities," *Future Generation Computer Systems*, vol. 51, no. 0, pp. 47 – 60, 2015, Special Section: A Note on New Trends in Data-Aware Scheduling and Resource Provisioning in Modern HPC Systems.

[2] J.H.P. Oosthoek, J. Flahaut, A.P. Rossi, P. Baumann, D. Misev, P. Campalani, and V. Unnithan, "Planetserver: Innovative approaches for the online analysis of hyperspectral satellite data from Mars," *Advances in Space Research*, pp. 219–244, 2013.

[3] M. Koubarakis, C. Kontoes, and S. Manegold, "Real-time wildfire monitoring using scientific database and linked data technologies," in *16th International Conference on Extending Database Technology*, March 2013.

[4] Chris Mangold, "Mobile device visualization of cloud generated terrain viewsheds," in *ASPRS 2014 Annual Conference: Geospatial Power in Our Pockets*, 2014.

[5] Athanasios Karmas, Angelos Tzotsos, and Konstantinos Karantzalos, "Scalable Geospatial Web Services through Efficient, Online and Near Real-Time Processing of Earth Observation Data," in *IEEE Int. Conf. on Big Data Computing Service and Applications*, March 2015, pp. 194–201.

[6] Angelo Pio Rossi, Peter Baumann, Jelmer Oosthoek, and Federico Cantini, "PlanetServer/EarthServer: planetary data exploration and exploitation through array databases," in *ESAC GIS Workshop, 5-7 May 2015, ESAC, Madrid (Spain)*, 2015.

[7] A. Karmas, K. Karantzalos, and S. Athanasiou, "Online analysis of remote sensing data for agricultural applications," in *OSGeo's European Conference on Free and Open Source Software for Geospatial*, July 2014.

[8] K. Karantzalos, D. Bliziotis, and A. Karmas, "A Scalable Web Geospatial Service for Near Real-Time, High-Resolution Land Cover Mapping," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. Special Issue on 'Big Data in Remote Sensing', 2015.

RHETICUS: FUSION OF EARTH OBSERVATION AND INSPIRE DATA FOR THE ENVIRONMENTAL MONITORING

Diomede Illuzzi, Paolo Manunta, Vito De Pasquale, Cristoforo Abbattista, Leonardo Amoruso, Luigi Agrimano, Sergio Samarelli, Mauro Casaburi

Planetek Italia s.r.l.

ABSTRACT

This paper talks about Rheticus, the Planetek cloud-based data and services hub. It is described how this system applies the Big Data concepts through the fusion of EO and ancillary data – with a particular attention to INSPIRE data - operated through a Big Data infrastructure that supports the batch processing of a continuously increasing volume of data for the generation of environmental indicators and alerts.

Index Terms— data fusion, environmental indicators, soil sealing, coastal monitoring, INSPIRE, GPU.

1. INTRODUCTION

Rheticus is the Planetek cloud-based data and services hub designed to deliver products through complex automatic processes and, if appropriate, a minimum interaction with human beings. One of the main purposes of the system is to enable Planetek Italia to shift from the provision of data to the provision of services, intended as continuous access to information from the users. This target is achieved by means of programmable components working as different software layers within an enterprise system relying on SOA model.

Every functionality is well defined and encapsulated in a standalone component making Rheticus highly scalable and distributable and allowing different configurations. This approach makes the system very flexible with respect to the services implementation, ensuring the ability to rethink and redesign the whole process with little effort.

Beyond giving hints about some of the solutions employed by Rheticus to tackle the technical challenge of Big Data management, such as the exploitation of GPGPU capabilities, this article aims to highlight the relevance of ancillary data – with a particular attention to INSPIRE [1] data - for the achievement of the level of accuracy that users claim for the services provisioned.

The description of a specific service case – the detection of new artificial areas - gives the opportunity to illustrate the Rheticus capabilities to generate environmental indicators and alerts through the fusion of EO and ancillary data, and to indicate how the integration of INSPIRE data can add value to the final product, in the context of an operational workflow driven by the Rheticus Big Data infrastructure that supports the batch processing of a continuously increasing volume of data.

2. RHETICUS AND INSPIRE

Among the peculiarities of Rheticus, there is the capability to integrate, in the processing algorithms, ancillary data, gathered in different ways from different data sources and channels, either for validation purposes or for increasing the information resolution enabled by the EO data used. In the case of coastal monitoring applications, for instance, Rheticus enables an alerting workflow triggered by crowd-sourced data, whose informative content is processed in real-time and checked against the EO-derived information to eliminate false alarms.

Highly beneficial, in terms of availability of usable ancillary data, is INSPIRE, the Infrastructure for Spatial Information in Europe, a Directive of the European Commission ([1]) that aims to establish an infrastructure for spatial information in Europe, in order to make spatial or geographical information more accessible and interoperable. INSPIRE is based on the infrastructures for spatial information established and operated by the 28 Member States of the European Union. The Directive addresses 34 spatial data themes ([2]) needed for environmental applications, with key components specified through technical implementing rules.

The interoperability of data and services, available through the Spatial Data Infrastructures (SDI) created by the European Member States is important to support the European Community environmental policies and all the activities that may have an impact on the environment.

It is worth to further remark the potential of fusing INSPIRE environmental data. Planetek Italia, in fact, intends to exploit the expertise matured in the development of the European INSPIRE Geoportal (see [3]), customizing the land and coastal applications for scenarios of cross-border environmental monitoring, where the access to the official data made available by the Member States within the INSPIRE ecosystem is a key element for the provision of a valuable service.

3. APPLICATION: FINDING NEW ARTIFICIAL AREAS

doi: 10.2788/854791

Soil is not a renewable resource and it is an indispensable basis for life. Soils that are sealed for human activities lose most of their functions due to disrupted water, nutrient, and biological cycles. This loss is close to be irreversible [4] and the measurement of soil consumption is an important task to address policies, planning activities and regulations. The monitoring of the new artificial areas has many practical applications: from finding illegal building or constructions, till to the monitoring of expansions of urban areas for reporting purposes at local, national or European level. This application has been developed to take advantage from the new generation of Earth Observation data in combination of data coming from not-EO sources, like INSPIRE data.

3.1. Input data

The data used for the proposed application can be divided in two categories: the EO data and the not-EO data.

3.1.1 EO data

The new generation of satellite sensors has brought a new way of thinking the applications. Their open and free access policies, in combination with the huge amount of data available, opens new perspectives in terms of algorithms but, at the same time, brings new difficulties in term of management of a big data volume.

The most important among the on-going space initiatives intended to produce EO-data, is probably the European Sentinel, a family of new space missions developed by European Space Agency (ESA) and the European Commission in the framework of the Copernicus Programme. The Sentinel family consists of five missions each one composed by a constellation of two satellites. The application proposed takes advantages from the quality of data of the first two missions: Sentinel-1 (S1) and Senitnel-2 (S2).

The S1 is a Synthetic Aperture Radar (SAR) mission operating in the C-band of the microwave part of the electromagnetic spectrum. It is an active sensor which does not depend on the sun to operate, thus it can acquire data during both day and night time. Moreover, the microwave wavelengths are marginally affected by the atmospheric condition, and the sensor can provide useful data in all weather conditions. Two S1 satellites are foreseen, one already in orbit and fully operative, the second planned for 2016. When both satellites will be operative, the constellation will acquire the same area every 6 days (12 days with one satellite).

The payload of S2 satellite is the Multispectral Instrument (MSI) which is able to sample the Earth surface in 13 spectral bands, spanning from visible to short wave infrared (SWIR) wavelengths, with a field of view 290 km wide. Four of the 13 bands are acquired by the MSI at 10 m spatial resolution, while 6 bands are acquired at 20 m spatial resolution and three bands at 60 m spatial resolution. The combination of the wide swath and the high resolution is the main characteristic of the sensor. All the continental land surface, the Mediterranean Sea, the bigger islands and all the costal water are systematically acquired by the sensor. This means that every 10 days an acquisition is available on these areas. This interval will be reduced at five days, when the second satellite, S2-B, planned for the mid-2016, will be operative.

The European Commission has decided to release the data with free and open access license. It will allow long-term continued access to the data opening new opportunities in setting up new service based on these data. On the other side, the amount of data to manage will increase dramatically. Each S2 acquisition, downloadable from the ESA catalog, has a size ranging from 2 to 7 GBs. The Level 1C data, used for this application, is partitioned in tiles of 100x100 km2 with a size of 500 MB each. For example, the Italian national territory is covered by 62 tiles, this means a maximum of 93 GB per month of data available on the country (this amount will be doubled when the constellation will be fully operative). For S1 the data used is the Level-1 Single Look Complex (SLC) product, which consists of focused SAR data geo-referenced using orbit and attitude data from the satellite. For each month, potentially more than 100 GB of data have to be processed in order to have full coverage for Italy.

From the technical point of view, the application is based on an object based supervised classification, thus a training dataset is required. For this purpose, the European Environmental Agency High Resolution Imperviousness layer [5] is used. It has been produced in the framework of Copernicus Land Monitoring Services and it maps all the artificial areas like buildings, transport infrastructure etc.

3.1.2 No-EO data

Ancillary data are necessary in order to improve the quality of the results of the application, and the INSPIRE data perfectly fit with this purpose and with the Rheticus operational concept. INSPIRE data, in fact, are published through interoperable network services, with which the Rheticus discovery and download clients can interact to automatically retrieve data of interest, filtering discovered results on the basis of the spatial theme of interest and other attributes specified by the standard metadata format. Once downloaded, the standard format of INSPIRE data enables the automatic processing capabilities of the Rheticus data fusion processor. The following section provides indication of which kind of INSPIRE data may be involved in the different steps of the process.

3.2 Methodology

For a single user, during the service setup phase, three parameters need to be defined: the area of interest, the starting

doi: 10.2788/854791

time and the periodization. The last one is the interval of time after which a new change-map is released, or, in other words, the period in which the area of interest is considered unmodified.

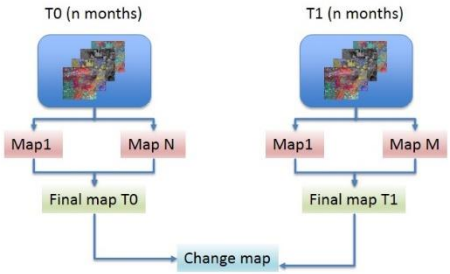

Figure 1: the basic structure of the methodology.

In figure 1 is schematically illustrated the procedure driven by the Rheticus system that downloads all Sentinel-1 (S1) and Sentinel-2 (S2) data available in the period T0 for the given area of interest. All the interferometric acquisitions of S1 are combined one to each other in order to produce several coherence maps. To reduce the noise, naturally present in SAR data, all the coherence maps are averaged producing only one coherence map of the area in the reference period. The optical data are pre-processed calculating the Top of Atmosphere reflectance. Each S2 acquisition is partitioned in three datasets according to the spatial resolution of its bands (three images). All data used, both SAR and optical, are then automatically coregistered. The subsequent step is to individually classify these datasets by means of an object based supervised classification method, which uses the European Environmental Agency Imperviousness layer 2012 to train the machine-learning algorithms. Finally, the resulting maps are then merged, with a pixel-based method, using Bayesian rules in order to derive the final map for the period T0. The ancillary data are used in this phase in order to improve the quality of the map. The Rheticus system is able to connect with the INSPIRE network services to automatically retrieve the data needed. In particular the following INSPIRE themes are used:

Transport networks: the automatic procedure, in some cases, fails to map the secondary roads or other transport infrastructures. The use of these data allows to cover this lack.

Buildings: buildings in low residential areas are not easy to map. If available in the area di interest, the location of the buildings improve the quality of the map.

Production and industrial facilities: the industrial areas are usually located at the edge of cities, in some cases not-vegetated bright soil can be confused with an industrial areas. This data layer can be useful to distinguish the two different land-use classes.

Hydrology: The bottom of shallow rivers is, in general, very bright; the automatic procedure can erroneously classify them, therefore, as built-up area. This layer can help to reduce this type of misclassifications.

The same procedure used for the period T0 is used for the period T1. During the time span of a period, Rheticus runs all the processing steps it can do with the data available. It periodically checks the catalogs for new images. If new data are available it downloads them and runs the appropriate processing. The Rheticus computation engine enables high level of parallelization; it can run, in fact, independent processes for single images or for group of images, triggered on the basis of what is discovered in the configured catalogues.

When T1 is over, and the T1 map is generated, theT0 and T1 maps are compared and the indications of the changed areas are derived. This last step is done by comparing the two maps and applying topological rules and probability rules to remove various type of errors. Moreover, also at this stage data from INSPIRE catalog are integrated in the above rule in order to improve the errors removal process. In particular, the themes that can be used are: Agricultural and aquaculture facilities, Cadastral parcels, Land Use and Land Cover.

A typical results is shown in figure 2, in which the red color refers to the changed areas with very high probability to be occurred, while the green color refers to the low probable changes (probable false positive).

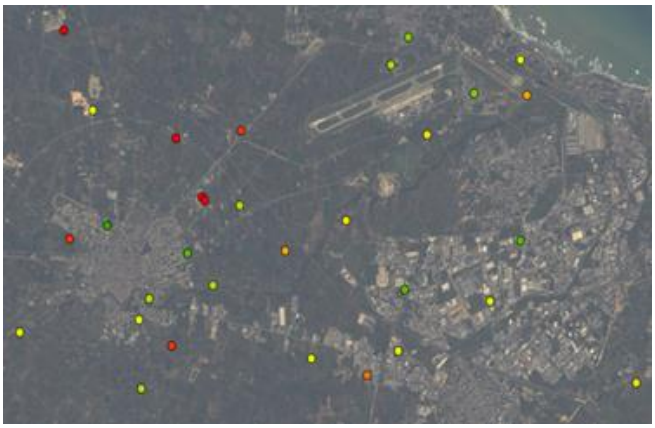

Figure 2: A typical example of the results of the application: the dots are the location of the potential changes detected (green dots indicate low probability, red dots high probability).

4. DATA PROCESSING USING GPU

Most of the Rheticus data processing tasks are highly parallelizable, thus they can take advantage, in terms of processing speed, by the use of GPGPUs (General Purpose Graphic Processing Unit), due to GPGPU architectures that

natively support parallelization ([6]). GPU architectures, in fact, are solutions that employ many cores (many thousands) with lower clock frequencies and lower size with respect to CPUs, providing, thus, massive parallelism on data access and processing, and handling natively the same simple operations or instructions executed at the same time on each core (SIMD instructions). Code parts requiring parallel computation applied to data can be written for GPU, while sequential code can continue to be executed by CPU.

Parallel code executed on GPU can improve from tens to thousands of times the performance of a single CPU. Today, different architectures are available provided by vendors like NVIDIA (Fermi and Kepler architectures), AMD/ATI (Cypress, Polaris) and Intel.

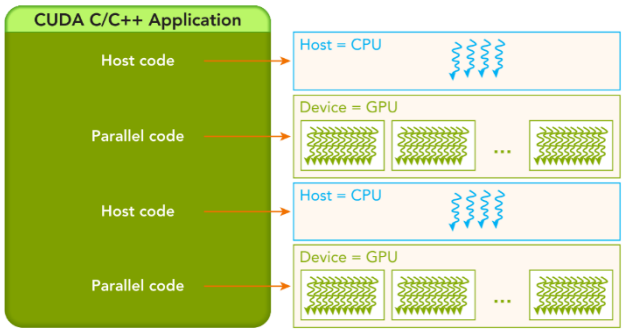

Figure 3: The general scheme of the interaction between CPU and GPU.

From the programmer point of view, different technologies are available for GPU programming. CUDA technology, only for NVIDIA GPUs, exposes API for GPU programming and C extension (Figure 3), while OpenCL (Open Computing Language) is a cross-platform technology for parallel programming (CPU, GPU, Accelerators) enabling to write programs for different GPUs (NVIDIA included).

5. RHETICUS AND CTEP

A further characteristic of Rheticus is the ability to act as an interoperable service node offering processing capabilities within a wider Big Data infrastructure. This concept will be demonstrated in the on-going ESA project CTEP (Coastal Thematic Exploitation Platform). CTEP project is addressing the coastal theme in the context of the Thematic Exploitation Platform (TEP, [7]) initiative, whose purpose is to exploit EO big data for the provision of a set of data and tools for thematic communities of users.

Coastal environment is particularly sensitive to anthropogenic pressure, pollution and climate change. Long-term monitoring can be supported by the generation of homogeneous time series (multi-mission data alignment), but also by automatic trend detection tools.

The CTEP aims, therefore, to provide new services and tools to enable a step change in the monitoring of coastal data, supporting a multi-disciplinary approach, and the provision of long-term data series and innovative services.

This objective is achieved by means of a widely distributed architecture which envisages the cooperation among the mother CTEP platform, and the children CTEPs. In the overall architecture, Rheticus plays the role of child CTEP serving specific needs of the Mediterranean users' community.

The integration of Rheticus in the CTEP federation allows extending both the computational capabilities of the overall platform and the portfolio of products and services offered, adding the processors developed by Planetek, when it is not possible for technical reasons linked to portability or for issues related to license of third-party software used, to deploy them on the mother CTEP. And this extension is done with a minimum integration effort, thanks to employment of interoperable processing interfaces following the Open Geospatial Consortium (OGC) Web Processing Service (WPS) standard.

6. REFERENCES

[1]. Directive 2007/2/EC of the European Parliament and of the Council of 14 March 2007 establishing an Infrastructure for Spatial Information in the European Community (INSPIRE) 14.03.2007

[2]. Commission Regulation (EU) No 1312/2014 of 10 December 2014 amending Regulation (EU) No 1089/2010 implementing Directive 2007/2/EC of the European Parliament and of the Council as regards interoperability of spatial data services

[3]. Planetek Italia, "INSPIRE geoportal", *http://www.planetek.it/eng/projects/inspire_geoportal*, 2015.

[4]. European Commission, Guidelines on best practice to limit, mitigate or compensate soil sealing http://ec.europa.eu/environment/soil/pdf/guidelines/EN%20-%20Sealing%20Guidelines.pdf, 2012.

[5]. European Environment Agency, "Copernicus Land Monitoring Services, Imperviousness," *http://land.copernicus.eu/pan-european/high-resolution-layers/imperviousness*, 2015.

[6]. David B. Kirk, Wen-mei Hwu, "Programming Massively Parallel Processors. A Hands-on Approach", Elsevier Inc, 2010.

[7]. S. Loekken, J. Farres, "ESA Earth Observation Big Data R&D, Past, Present, & Future Activities", CEOS WGISS Meeting #37, Cocoa Beach, USA, 2014.

DAMATS: PROTOTYPE SYSTEM FOR CLASSIFICATION OF SATELLITE IMAGE TIME SERIES

Catalin Cucu-Dumitrescu[1], Florin Serban[1], Adrian Stoica[1], Corina Vaduva[2], Martin Pačes[3], Michele Iapaolo[4]

[1]TERRASIGNALtd., Bucharest
[2]University Politehnica of Bucharest
[3]EOX GmBH, Vienna
[4]ESA ESRIN, Frascati

ABSTRACT

A functional prototype software for analyzing satellite image time series (SITS) is presented. The system is designed to extract classes of pixels with similar evolution in time and asses the dynamic of the Earth's surface. A set of criteria applied over a given repository of EO imagery will determine the generation of the SITS. There are three actions available to the user: **1**. Perform a general classification (for all the pixels in the area of interest), **2**. Visualize and analyze the change detection masks between pairs of scenes in the SITS, **3**. Extract one class starting from an example. All the results presented in the paper are obtained in the context of the ESA GSTP project: Data Mining for Analysis and Exploitation of Next Generation of Time Series (DAMATS) [1].

***Index Terms*—** Satellite Image Time Series, Classes of Evolution, Latent Dirichlet Allocation, Dynamic Time Warping, Change Detection.

1. INTRODUCTION

There is more than a decade of satellite image acquisition. Earth surface has been pictured for many times, creating series of images as tools for long term monitoring. When referring to Satellite Image Time Series (SITS), we picture a pixel in the area of interest and we see it changing from one acquisition moment to another and so on. This process is expressed in terms of a linear sequence made from the spectral data of the individual scenes, in chronological order. The concept of 'classes of evolution' is a natural generalization of the concept of class at scene level (such as Corine Land Cover classes or Urban Atlas classes). No simple labels can be used in labeling the classes of evolution, as they donôt have the structure of a static class, but rather that of a story (e.g. one hypothetic class could be ñSITS-pixels which were in a forest area during the first 3 scenes, then changed into deforest-area for the next 4 scenes and belong to agriculture areas in the last 12 scenesò).

The SITS explain the need to preserve all the data and provide access for further exploitation. The challenge is to extract "categories of evolution", to discover the information able to explain the land cover evolution process. Several methods have been developed in order to fulfill this goal and to offer a valid solution in scenario based applications.

Considering the series as multidimensional signals, the physical changes are statistically modeled through an interactive learning process for user-based semantic annotation in [6]. Clustering trajectories is addressed in [7] by grouping spatio-temporal interdependencies as an issue of random field modeling. In [8] the authors deal with the problem of fast access and retrieval of relevant information from databases by building an index of compressed object database. Compression based techniques are also used in several application scenarios in [9], showing the generality of the concept. Similar to the idea of defining dictionaries, a frequent sequential pattern extraction technique has been adapted for spatio-temporal behavior [10]. A method that exploits the dynamic time warping technique is presented in [3]. The method can deal with sequences that are not equal in length, i.e., temporal samples might miss due to meteorological conditions, technical errors. The generality of the approach is argued also by the tests developed on both SAR and optical images. The methodology described in [2] exploits a wide range of dynamic evolution: short time changes, as well as long time transformations.

Evolution can also be defined as the result of a transformation occurred between two moments. This is modeled through change detection algorithms. There are several approaches that perform unsupervised binary change detection. A method to find the optimal threshold to be applied on the pixel values of the image resulted from the differentiation of the two images is presented in [4]. A different approach is suggested in [5], which aims at finding the directions of change in the images through a Principal Component Analysis.

All the above mentioned procedures deal with single aspects, extracting categories of evolution or detecting changes. The variety and complexity of EO data is one of the main drawbacks. The temporal information increases the

doi: 10.2788/854791

open issues. In search of a solution, the authors present a platform integrating well known techniques in order to provide a tool for complementary processing involving temporal analysis for the exploration and exploitation of SITS. Moreover, a dedicated module for SITS generation enabling data selection by means of specific criteria and a data analytics interface come to complete a compound architecture for a SITS information mining system.

2. FRAMEWORK FOR MULTI-TEMPORAL ANALYSIS

SITS exploitation is a great opportunity for many monitoring applications. The process involved enables the extraction of significant information regarding the complex transformation and evolution processes of Earth surface. Considering the repetitive, yet irregular in time, data acquisition process of the same scene, SITS generation and analysis becomes a real challenge. The available techniques provide solutions for the automatic discovery of regularities, relationships and especially temporal interdependencies. An understanding of the underlying phenomena is created. However, there is more to learn in the process..

The paper proposes an effective solution to the issues arising in this frame. The diagram in Fig. 1 illustrates the main functionalities considered, as well as the critical points where human intervention is required in order to get the reliable and relevant information with high precision.

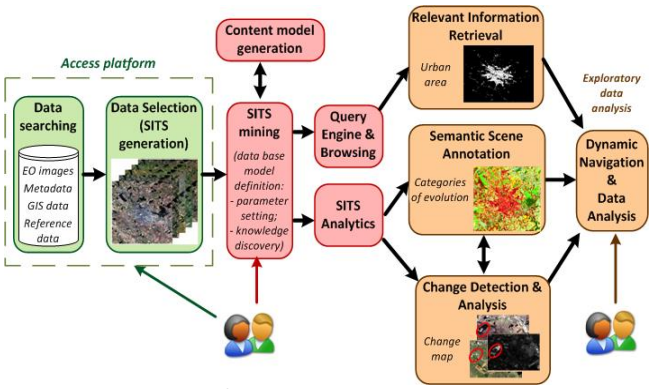

Fig. 1. Framework for multi-temporal analysis. In green, the modules enabling the access to the system. In red, the data processing. In brown, the interface guiding data understanding, visualization and analytics.

The concept was designed such that it provides access to the user via a web platform that will link him to the data repository and the core processing. A graphical interface enable the human machine interaction, including all the required settings and the visual exploration. After connecting to the repository, the general workflow starts with the data searching and selecting the proper criteria towards SITS generation. Further, in order to proceed with the application scenario definition, the SITS created must be analyzed using mining tools, whose parameters must be set by the user based on the preferred functionality. The content will be modeled accordingly, emphasizing the spatio-temporal evolution such that the obtained results to be considered as further input data for searching and analysis modules. The user may also extract relevant information from the scene, given a query example. Groups of pixels with similar temporal evolution will be highlighted, separating a distinctive class. The categories of evolution may be verified by comparison with short time analysis that means change detection between two acquisitions from the SITS. The final step refers to the exploratory data analysis in order to understand land cover transformations over the selected time period.

There is no significant contribution as related to the stat of the art SITS processing methodology. Nevertheless, the concept is an important step forward because it provides an architecture integrating several techniques in order to extract complementary information and discover hidden knowledge that is difficult to perceive at a visual inspection. Moreover, the proposed framework encourages: SITS generation, extraction of categories of evolution, semantic searches and complementary analysis.

2.1. SITS generation

SITS generation is one of the key steps in analyzing the evolution of SITS evolution. In order to extract the maximum amount of information from the SITS, a simple stacking of EO images is not sufficient. The simple visualization of the multi-temporal data does not deliver an indicator of the processes that occur in a certain area, especially when dealing with large amounts of data as it is the case of SITS. Here are the most important steps to be followed in order to obtain the desired information:

- selection of the area of interest - all the images in the SITS must completely cover this area.
- selecting the sensor and acquisition time - the image characteristics and the time period must be in accordance with the desired application.
- selecting cloud coverage - optical images containing a high percentage of cloud coverage should not be taken into consideration, as clouds and their shadows distort the data.

2.2. SITS analysis

Since there is no visual confirmation in the case of temporal evolution increased the interest towards unsupervised techniques for SITS analysis. Processing a large amount of EO data comes at the cost of finding algorithms that work autonomously. Given the lack of reference measurements, additional information offer an advantage to the analysis. Therefore, the study of variations between consecutive acquisitions in the SITS, with focus on

doi: 10.2788/854791

detecting the changes, sustain a better understanding of the data content.

Generally speaking, the multi-temporal analysis of SITS can be seen from different perspectives depending on several criteria, such the number of temporal samples, the type of application scenario considered and the amount of a-priori information held at the moment of analysis. These criteria are interconnected and they must be strongly related to the needs of each user. While the last two criteria a completely defined by the user, the first criteria is considered to be an open issue. In order to cope with the shortcomings of irregular time sampling, the proposed framework integrates two types of analysis:

1) *bi-temporal analysis* – two temporal images are compared in order to determine the existent differences (change detection). Different alternate methods for extracting change detection maps were used [4], [5].
2) *multi-temporal analysis* – more than two images are analyzed in order to extract information about temporal evolution, history of changes in an area and so on. There are two methods implemented for general classifications: one is based on Latent Dirichlet Alocation (LDA) [2], and the other one is based on Dynamic Time Warping (DTW) [3]. The LDA approach has the advantage that it can be applied directly on the SITS, but also on a time series obtained from change maps computed from consecutive image acquisitions from the SITS.

2.3. Specific information retrieval from SITS

A query by example approach, this functionality focuses on the idea to express temporal meaning from the SITS using temporal indexing instead of syntax. The corresponding signatures (corresponding to the temporal evolution of one pixel) will be allocated into appropriate metric space in order to define the similarity between temporal signatures comparable evolutions. The analysis of the collection by a machine is only able to provide similarity by data processing, resulting in a central issue because the meaning of SITS is not self-evident. A temporal signature is compared with the entire collection of temporal signatures. As similarity measures appropriate for specific information retrieval from SITS, we introduce the Euclidean distance and the Levenshtein distance.

3. PRELIMINARY RESULTS

Considering the proposed framework for multi-temporal analysis, we present some preliminary results using a dataset of 13 Landsat 7 image time series covering a 54x87 sq km in Dobrogea, Romania (Fig. 2). The purpose is to provide information able to support a short time monitoring of a region containing: artificial surfaces (C1), agricultural areas (C2), forest (C3), wetlands (C4), water bodies (C5).

6.05.2000	29.10.2000
14.05.2000	9.05.2001
22.05.2000	25.05.2001
7.06.2000	12.07.2001
1.07.2000	29.08.2001
9.07.2000	14.09.2001
18.08.2000	

Fig. 2. Area of interest: Landsat 7 images, Dobrogea, Romania (left), Acquisition dates for SITS (right).

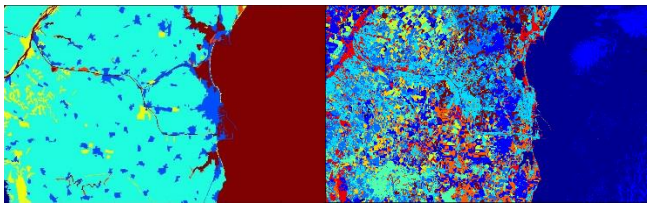

Fig. 3. Reference map ï CLC2000, 5 static classes (left), SITS Classification, 10 classes of evolution (right)

Table 1. Confusion matrix for SITS classification.

	lda1	lda2	lda3	lda4	lda5	lda6	lda7	lda8	lda9	lda10
C1	1.70	4.55	11.42	**59.21**	2.38	0.43	0.67	3.52	10.98	5.14
C2	1.38	18.33	**19.12**	17.70	9.58	6.03	4.99	10.45	3.73	8.69
C3	0.74	15.89	9.50	26.90	1.01	0.53	1.80	1.78	**40.65**	1.20
C4	3.19	8.21	7.37	35.92	4.11	1.06	0.57	0.92	**37.07**	1.58
C5	**85.83**	10.49	0.07	0.81	0.05	0	0	0.01	2.69	0.05

A level 1 Corine Land Cover map (Fig. 3, left side) from the same period of time was selected as a reference of the land cover (static classes). We proposed 3 types of independent analysis in order to exploit the data.

1. *Multi-temporal analysis* - after a LDA processing on the SITS, we have obtained a scene classification based on temporal evolution (Fig. 3, right side). Considering the diversity and dynamic of the region, we opted for 10 classes. The classification result is compared to the CLC map aiming at presenting the way that the static classes can be further divided according to their dynamics over the time period covered by the analyzed SITS (Table 1). For instance, the agriculture areas are characterized by 4 major types of temporal evolution.

2. *Bi-temporal analysis* - a change detection was made in order to establish how much the scene have changed in one year, given a certain period of time: spring (Fig. 4 - a,b,e) and autumn (Fig. 4 - c,d,f). This is a measurement of scene transformation supplementing the first analysis that shows the way the scene has been transformed.

3. *Specific information retrieval* - opposed to the first 2 approaches providing a complete scene processing, this analysis focuses on finding precise structures, with similar dynamics. We have selected the points of interest based on the CLC annotation. Considering the DTW as the similarity measure used to retrieve spatio-temporal evolutions, we preserve from the scene all the pixels whose distance is less than 10% of the maximum distance value computed between the query and all the pixels in the image (Fig. 5).

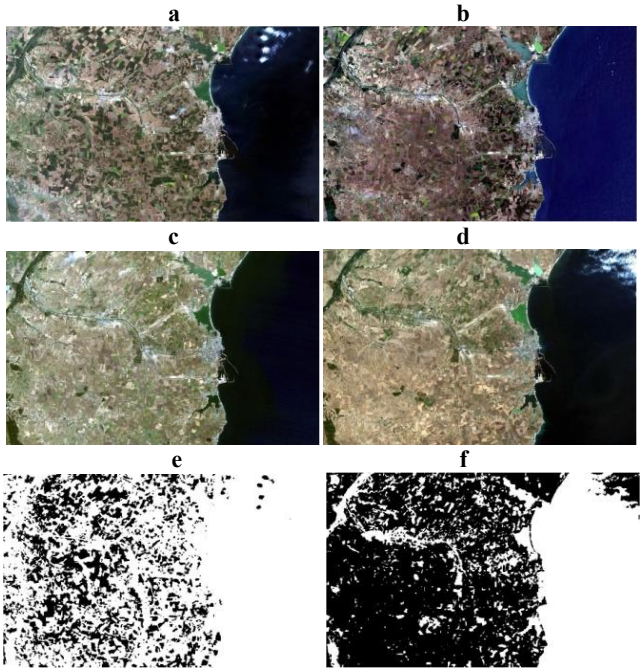

Fig. 4. a) Before-May2000; b) After-May2001; c) Before-September2000; d) After-September2001; e) Change map between a)&b); f) Change map between c)&d). Black means change.

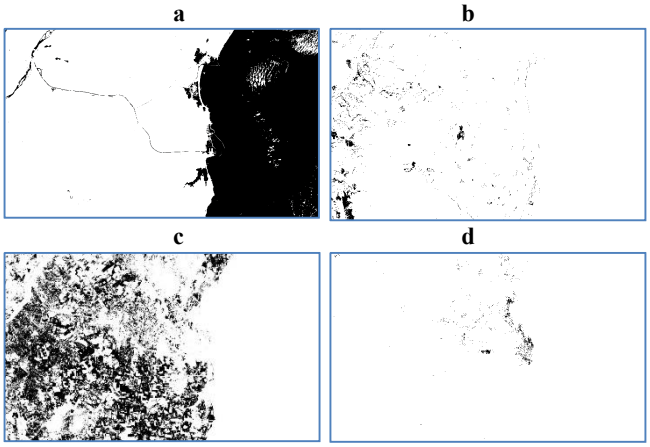

Fig. 5. Temporal evolutions corresponding to: a) water; b) forest and semi-natural; c) agriculture; d)artificial surfaces.

Table 2. Confusion matrix for specific information retrievals.

	C1	C2	C3	C4	C5
C1	**74.28**	4.8	6.36	5.24	0.35
C2	22.08	**91.69**	33.40	25.24	0.14
C3	1.02	3.62	**58.75**	1.43	0
C4	0.82	0.36	0.83	**67.14**	0
C5	1.80	0.28	0.66	0.95	**99.41**

If the results are compared to the CLC map, a certain correspondence between the static classes and the categories of evolution (Table 2. Confusion matrix for specific information retrievals.) is observed. However, the dynamics of the Earth surface is modeling the classification.

4. CONCLUSIONS

The multi-temporal analysis framework presented in this paper was the baseline towards the implementation of a prototype system that comprises general analytical methods for the exploitation of the information contained in SITS. The proposed system relies on data mining for analysis and exploitation of next generation of time series, focusing on the information extraction in the form of ñcategories of evolutionò and in technologies to classify the evolution processes of observed scenes. A short system demonstration for a monitoring application was presented on a Landsat 7 image time series.

5. REFERENCES

[1] DAMATS Website: http://wiki.services.eoportal.org/tiki-index.php?page=DAMATS&highlight=osiride

[2] C. Vaduva, T. Costachioiu, C. Patrascu, I. Gavat, V. Lazarescu, andM. Datcu, ñA Latent Analysis of Earth Surface Dynamic Evolution Using Change Map Time Seriesò, *IEEE Trans. Geosci. and Remote Sens.*, vol. 51(4-1), pp. 2105-2118, 2013.

[3] F. Petitjean, J. Inglada, and P. Gancarski, ñSatellite Image Time Series Analysis Under Time Warpingò, *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 8, pp. 3081-3095, August 2012.

[4] L. Bruzzone and D.F. Prieto, ñAutomatic analysis of the difference image for unsupervised change detectionò, *IEEE Trans. Geosci. and Remote Sens.*, vol. 38(3), pp. 1171-1182, May 2000.

[5] T. Celik, ñUnsupervised change detection in satellite images using principal component analysis and k-means clusteringò, *IEEE Geosc and Rem Sen Letters*, vol. 6(4), pp. 772-776, Oct. 2009.

[6] P. Heas and M. Datcu, "Modeling trajectory of dynamic clusters in image time-series for spatio-temporal reasoning", *IEEE Transactions on Geoscience and Remote Sensing*, vol.43, no.7, pp. 1635-1647, July 2005.

[7] L. Gueguen, M. Datcu, "Image Time-Series Data Mining Based on the Information-Bottleneck Principle", *IEEE Trans on Geosc and Remote Sensing* 45(4): 827-838, 2007.

[8] L. Gueguen and M.Datcu, ĂA similarity metric for retrieval of compressed objects: Application for mining satellite image time seriesò, *IEEE Trans Knowl Data Eng*, vol.20,no.4, pp.562-575,Apr.2008.

[9] Cerra, D., Mallet, A., Gueguen, L., Datcu, M., ĂAlgorithmic Information Theory-Based Analysis of Earth Observation Images: An Assesmentò, *IEEE Geoscience and Remote Sensing Letters*, vol.7, no.1, pp.8, 12, Jan. 2010.

[10] A. Julea, et all, "Unsupervised Spatiotemporal Mining of Satellite Image Time Series Using Grouped Frequent Sequential Patterns", Trans on Geos and Rem Sen, vol. 49, no. 4, April 2011.

THE EUCLID ARCHIVE SYSTEM:

A DATA-CENTRIC APPROACH TO BIG DATA

A.N.Belikov [(1)], O.R.Williams[(2)], B.Altieri[(3)], D.Boxhoorn[(1)], G.Buenadicha[(3)], B. Droege[(2)], J.McFarland[(1)], S.Nieto[(3)], J. Salgado[(3)], P. de Teodoro[(3)], A.Tsyganov[(2)], E.A. Valentijn[(1)]

(1) Kapteyn Astonomical Institute, University of Groningen, Groningen, The Netherlands

(2) Donald Smits Centre for Information Technology, University of Groningen, Groningen, The Netherlands

(3) European Space Astronomy Center, European Space Agency, Spain

ABSTRACT

The Euclid Archive System (EAS) sits in the core of the Euclid Science Ground Segment (SGS). It supports the processing and storage of Euclid data from the input of raw frames to the creation of science-ready images and catalogues. We review the architectural design of the system, implementation progress and the main challenges in the building of the EAS.

***Index Terms*— Euclid, Data Storage, Information System**

1. INTRODUCTION

Euclid is an ESA M2 mission and a milestone in the understanding of the geometry of the Universe [1]. Euclid faces two main challenges from the point of view of the data processing. Firstly, the unprecedented accuracy which must be achieved in order to meet the scientific goals. Secondly, the mission will depend heavily on the processing and reprocessing of ground-based data which will form the bulk of the stored data volume [2]. In total Euclid will produce up to 26 PB per year of observations [3].

2. EAS AND SGS

The Euclid Science Ground Segment is a distributed data processing and data storage system which is responsible for the delivery of the science-ready data to ESA [3]. The SGS is formed by 9 national Science Data Centres (SDCs) and the Euclid Science Operations Centre. Each data centre provides resources for processing, along with expertise in the coding of Euclid pipelines. The Euclid Archive System (EAS) is a core element of the SGS implementation of a data-centric approach [4] to the data processing.

In the SGS and the EAS the data is a combination of data files (images, spectra, catalogs) and metadata. The metadata includes the full data lineage, which is necessary to achieve the goals of the Euclid mission. The EAS must provide to SGS users and the subsystems of SGS the ability to trace any bit of information produced in the SGS. The EAS must also provide quality information on the produced images, spectra and catalogs. According to experience in previous missions (e.g. Astro-WISE [4] and the LOFAR Long-Term Archive [5]) the data volume of metadata will not exceed 5% of the total data volume. The clear distinction between the bulk of the data (stored in files) and extensive metadata (describing the data products in files) allows most of data mining operations to be moved to the relational database without access to the data files.

The EAS is a joint development between ESA and the Euclid Consortium and is led by the SDC of the Netherlands and the ESDC (ESAC Science Data Centre).

It should be stressed that the EAS is not an archive in the conventional sense. Instead, the EAS is a distributed scientific information system which ensures that any operation with the data is registered and can be traced back to the user and pipeline. EAS ensures that SGS can access and process hundreds of Petabytes of data.

doi: 10.2788/854791

3. EAS DESIGN

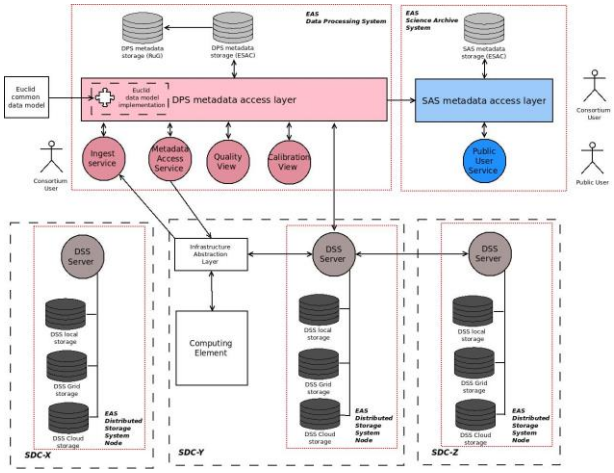

Figure 1: Overview of EAS and it's place in Euclid SGS

The EAS design is based on requirements, formed by ESA and the Euclid community, on the processing of data in a distributed environment. Due to different nature of requirements on the EAS the whole system was divided on 3 independent parts:

- The Data Processing System (DPS) supports the data processing activities within the SGS
- The Science Archive System (SAS) supports the scientific use-cases, the release delivery of data to the wide astronomical community, and long-term data preservation.
- The distributed Storage System (DSS) for the storage of data files.

Figure 1 is an overview of the EAS design showing primary users and their interaction with DPS, SAS and DSS.

The primary task of the DPS is to serve as an information system for the production of data releases. To achieve this goal, DPS keeps data lineage for each data object; tracing any operation in the system to allow data reproduction and reprocessing. DPS allows each frame, catalog or source catalog to be traced back to the original raw data, as well as the processing parameters used and the pipeline which produced it. As well it allows the current status of the processing in SGS and quality of the raw and processed data to be viewed.

The DSS serves both the DPS and SAS as a common distributed file storage solution. In this solution we utilize a distributed approach for the storage of data files and a centralized approach to the storage of the metadata. The DSS consists of a network of DSS servers which provide interfaces to the non-homogeneous data storage solutions provided by the national SDCs. At least one DSS server is installed at each SDC and stores the data files processed or created during running of a pipeline in this SDC. The design of the processing plan for each pipeline minimizes data file transfer between SDCs. To ensure zero loss of the data at least 2 copies of each data file is created in the DSS. Each copy is registered in the metadata storage.

The DPS implements the Euclid Common Data Model (ECDM) which describes both scientific data (data products generated by pipelines) and processing and operational metadata (processing and data distribution orders, location of the file in the DSS, processing plans). On the other side, SAS implements the Science Exploitation Data Model (SEDM), which describes the scientific metadata. The interface between SAS and DPS will allow the SAS to retrieve the science metadata from the DPS for an Euclid data release and put it in long-term storage in SAS. The same ECDM is used to bind pipeline data flows.

The SAS is part of the EAS and it aims to support the needs for scientific data exploration for the Euclid Consortium and the wider astronomical community. It will face the challenge of Big Data, as it will store a huge and increasing amount of scientific metadata and catalogues of 10 billion of galaxies. This will provide the worldwide astronomical community with an extremely large source of targets for future missions. Under these premises, the SAS will face the challenge of guaranteeing the long-term preservation of Euclid data while providing the scientific community with access to this data.

The SAS is being built at the ESAC Science Data Centre (ESDC), which is responsible for the development and operations of the scientific archives for the Astronomy, Planetary and Heliophysics missions of ESA. The SAS is focused on the needs of the scientific community. In this context, the SAS will provide access to the most valuable scientific metadata coming from the EAS-DPS through a set of public data releases. According to the policy of public releases defined by the Euclid Consortium, the plan is to deliver 3 data releases to the scientific community every 2 years after the nominal start of the mission [6].

The design of the EAS is fully hardware independent and partially software independent. Metadata storage for DPS and SAS based on RDBMS.. For DPS it is currently Oracle 11g, but it will be possible to switch to other RDBMS in the future. Metadata storage in the DPS and the SAS are completely independent and can be based on different solutions. A DSS server can use local or NFS-mounted filesystems, Grid SE, sftp servers, iRODS, Astro-WISE dataservers, GPFS. It can be further extended for other existing or upcoming storage cases. The ability to extend DSS storage to practically any storage solution was introduced in the design to cope with the possible changes during the lifetime of the project and to support the SGS for at least 8 years after the launch.
The design of the SAS follows the latest generation of archives being developed by the ESDC, taking full advantage of the existing knowledge, expertise and code. The SAS builds on top of the latest ESDC's common Archives Building System Infrastructure (ABSI), which defines the common components to the latest ESA Science Archives (i.e. Gaia) within a three-tier modular architecture (client, server and data layer).

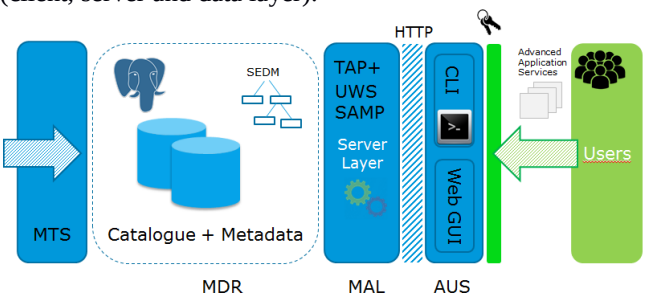

Figure 2: SAS Architecture Diagram

4. EAS INTERFACES

There are a number of interfaces and services which are implemented for the EAS. The EAS DPS will support interaction with the SDCs for the data processing, through the Infrastructure Abstract Layer (IAL). IAL queries the EAS DPS to retrieve data distribution and data processing commands. The IAL also sends metadata for newly created data products to the EAS DPS for ingestion. The basic metadata browsing service allows users to browse the content of the metadata storage and retrieve metadata for the data product in XML or other format.

The DPS will be supplied with at least 2 other services: one for the manual inspection of the quality of data products and another for the retrieval and update of calibration frames.

The SAS will provide two ways of access through the Archive User Services (AUS): a web-based portal and a command line interface for programmable access. The web-portal is based on the Google Web Toolkit technology [7] and it has been designed to be easy to use and to provide a friendly interaction.

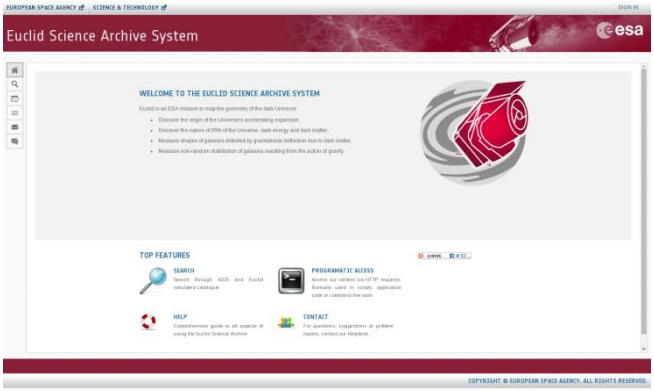

Figure 3: SAS Web User Interface

On the server side, the components are based on Java and it will integrate standard VO protocols to manage requests from the users.
At the database level, the metadata repository will support the SEDM, which describes the metadata and catalogues oriented to scientific exploration. At the time of writing, the RDBMS is PostgreSQL join with pgSphere and Q3C modules that provides spherical data types, functions and operators to PostgreSQL.

Finally, the Metadata Transfer Service (MTS) will ingest the science metadata of the data release, compliant with the SEDM, from the ECDM implemented in the DPS. This will follow the policy of public releases defined by the Euclid Consortium.

The scientific requirements on the SAS cover three main areas: parametric search for metadata and catalogues, data retrieval and the visualization of images, spectra, etc. Regarding the visualization of maps, it will be based on the technology developed for the ESA Sky [8], that allows the exploration of the astronomical resources using a useful and intuitive web interface. In addition, it will provide a set of tools to allow on-line research.

The distributed nature of the user community requires standard protocols oriented to access, exchange and store data to guarantee the information accessibility along [9]. VO protocols play an important role within the architecture of the SAS, but also support added value tools integrated into the archive. Table Access Protocol (TAP+) [10] provides an

efficient parametric search and has been developed for the Gaia Archive[11]. TAP+ is also part of the SAS infrastructure and can be accessed from the Web interface.

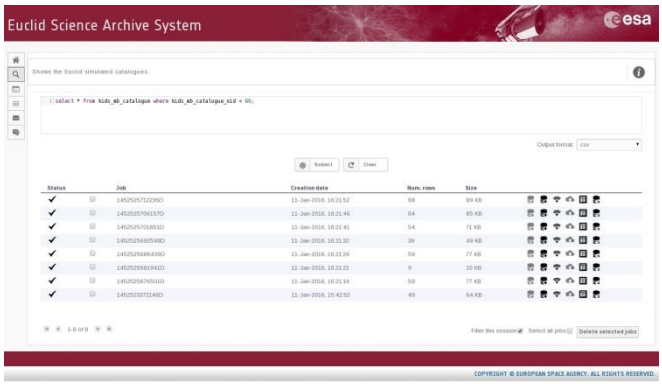

Figure 4: TAP Web Interface in SAS

VOSpace [12], the IVOA protocol for distributed data storage, is also integrated as part of the SAS infrastructure as an added value tool for the archive. It provides a storage abstraction layer and sharing capabilities transparent for the user.

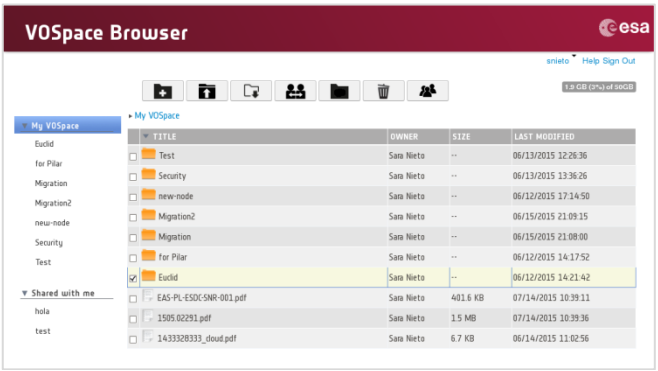

Figure 5: VOSpace Web User Interface

Other VO protocols like SAMP, will provide interoperability with astronomical analysis applications.

To conclude, SAS will provide the tools and VO interfaces to enable the Software-to-Data Paradigm and "bring the software to the data" for the Euclid science.

4. EAS STATUS

The EAS Prototype (composed by DSS and DPS) was developed and tested in 2013 and 2014. In 2015 the prototype formed the basis of the first version of EAS itself. First interfaces for the DPS and the DSS were released and tested during several IT challenges organized by Euclid Consortium in which simulated Euclid data was produced and stored. We have successfully tested massive metadata ingestion for the data objects with extensive data lineage (KIDS DR1) [13]. In 2015 EAS team tested a master-slave configuration of EAS DPS metadata storage. Metadata was transferred between the EAS DPS metadata storage mirror in ESAC and the current master site Groningen.

5. FUTURE DEVELOPMENT

In the following years prior to the start of the mission the EAS will go through a number of crucial steps in the development including the final selection of an RDBMS. We plan as well to create a system which will support a dynamic data model: accommodating changes in ECDM without re-creation of the metadata storage scheme and migrating the data between different versions of the data model.

6. REFERENCES

[1] R. Laureijs et al., Euclid Definition Study Report, arxiv1110.3193, 2011

[2] O.R.Williams et al., Data transmission, handling and dissemination of Euclid, Proceedings of NETSPACE 2014, 11, 2014

[3] F. Pasian et al., Development Plans for the Euclid Science Ground Segment, Proceedings of ADASS XXIII, ASP Conf. Ser. 485, 505, 2014

[4] K.Begeman et al., The Astro-WISE datacentric information system, Experimental Astronomy 35, 1, 2013

[5] K.Begeman et al., LOFAR Information System, Future Generation Computer Systems 27, 319, 2011

[6] Maurice P. et al., Euclid : Big Data From Dark Space, Big Data Spain, Oct. 2015

[7] Dewsbury, R., Google Web Toolkit Applications, ISBN 978-0-321-50196-7, Prentice Hall, (2007).

[8] Salgado, J., ESA Sky: A Tool to Discover Astronomical Data from ESA Missions, Big Data From Space 2016.

[9] González, J. et al., Open Archives for Open Science at the ESAC Science Data Centre (ESDC), ADASS 2015.

[10] Dowler P. et at., IVOA Recommendation: Table Access Protocol Version 1.0, 2011arXiv1110.0497D.

[11] Arviset A. et al., Big Data, Big Data Challenges and New Paradigm for the Gaia Archive, Big Data From Space 2016.

[12] Graham M. et al., IVO Working Draft : VOSpace Version 2.1, WD-VOSpace-2.1-20150601

[13] J.T.A. De Jong et al., Kilo Degree Survey, Experimental Astronomy 35, 24

CHALLENGES AND ACHIEVEMENTS IN SUPPORTING RESEARCH AND SERVICE DEVELOPMENT FOR BIG DATA

J. M. Delgado'[1,2], G. Sabatino'[1,2], R. Cuccu'[1,2], G. Rivolta'[1,2], P. G. Marchetti[3]

[1] ESA Research and Service Support, via Galileo Galilei, 1, 00044 Frascati (Italy)
[2] Progressive Systems Srl, Parco Scientifico di Tor Vergata, 00133 Roma (Italy)
[3]ESA/ESRIN European Space Agency, via Galileo Galilei, 1, 00044 Frascati (Italy)

ABSTRACT

This paper addresses the services and challenges which are faced when the objective is to support the rapid and widespread adoption of the new paradigm "bring users to data". In the Big Data era [1], started with the launch of Sentinel-1A in April 2014, data volume and processing requirements are becoming more and more challenging. Supporting the new paradigm as an alternative to traditional "bring the data to the users" approaches [2], is of paramount importance to a successful exploitation of Earth Observation (EO) Big Data as the volume, velocity and variety of data from space become an obstacle performing research or developing new services. The ESA Research and Service Support (RSS) provides services and support solutions to face the new Big Data era.

***Index Terms*— Big data exploitation, earth observation, research support service, dynamic processing resources, cloud, grid computing**

1. INTRODUCTION

The Big Data characteristics entail the availability of high capacity telecommunication networks, data processing and analysis infrastructure on the user side. We will focus here on 3 main use cases, namely when Big Data from EO is needed to actually 1- perform research activities, 2- use the data in university classes or 3- develop new downstream services and finally give value to the data. This paper addresses the requirements of a service aiming to support the new paradigm "bring users to data", provides some quantitative analyses and benchmarks on the processing, parallelization and performance requirements of Copernicus products from Sentinel-1 (SAR) and Sentinel -2 (Multi-Spectral) including the use of the dedicated toolboxes on representative use cases and describes the challenges and the achievements of the ESA Research and Service Support operational service pilot.

Finally, emphasis is given to the emerging platforms "thematic", "mission specific", "regional" and how the RSS service needs to adapt and be tailored to cover and fully support those platforms.

2. THE RSS SERVICE

The RSS service offer is composed of several elements supporting different phases of the research process flow. It includes e-collaboration environments to find and share information, reference and sample datasets, access to a huge EO data archive without the need to download on scientist or developer "own" resources, customised cloud toolboxes where scientists and developers alike can fine- tune their algorithms on selected datasets, on-demand processing environment where fine-tuned algorithms can be integrated and made available as EO applications for on-demand massive processing, and results visualization tools.

RSS users are EO Scientists, Researchers and – during the prototyping phase - Service Providers. According to users' experience resorting to the RSS service can generate significant savings in terms of time (~months), besides the savings related to processing and storage management costs. For large dataset processing these can be hundreds of CPU hours and tens of TB, respectively (see Tab. I). It is worth to notice that such savings , in turn, generate higher scientific productivity and faster time to market.

TABLE I.
DATA VOLUME AND TIME SAVINGS BY PROCESSING ON ESA RSS.

Processing campaign	Input Data Volume	N. jobs	Campaign Time (G-POD)
BIOMASAR – II Project	14.8 TB	3141	7 months
NEST-Urban Growth *	6.5 TB	7700	7 months
NEST- Iced lakes & rivers	5.5 TB	15788	8 months
S1- Iced lakes & rivers	90.9 GB	166	1.5 days

Tab. I summarizes savings in terms of storage capacity (needed only for the input data), number of jobs and campaign duration for some processing campaigns supported by RSS. The campaign time is the real processing time from the start to the end of the supported campaign.

Depending on the user needs RSS is able to support (i) the algorithm development and post-processing activity, and (ii) massive data processing. RSS makes available different solutions in correspondence of the Principal Investigator (PI) needs: i) the cloud toolbox service, used for development, analysis and processing when limited computing and data resources are needed; ii) the processing on demand service, based on a full-fledged infrastructure with (virtually) unlimited capacity.

As far as the algorithm development process is concerned, including the fine-tuning phase, the RSS CloudToolbox [3] is the basic tool offered by RSS to EO researchers (see Fig.1). Such tool is a customised virtual machine hosted on Cloud Providers with pre-installed software and additional tools and packages that can be installed on request..RAM, CPU cores and HD size are tuned around the user needs. It is possible to access this service via the RSS CloudToolbox service portal[1]

Figure 1. Cloud Toolbox Desktop view.

Once the development phase on the CloudToolbox is completed and the algorithm is deemed to be stable, this can be integrated into the RSS processing environment, thus bringing it close to data, , either if the scientist plans to run it on massive datasets (Big Data processing) or to make it available to the scientific community as a web application on the Processing on-demand website portal[2] (see Fig.2).

From the web portal, users can access the processing on-demand service, which allows to run scientific processors on selected datasets. Either single tasks or scheduled tasks can be submitted. In the latter case, particularly useful for massive processing, the "task scheduler" is configured to automatically create and submit processing tasks.

Figure 2. Processing on-demand service website

In addition, the processing on demand portal allows users to monitor the status of the tasks and also the status of each single block of the processing chain. . The different possible status are: (i) prepared, (ii) created, (iii) active, (iv) completed or (v) failed. An example is shown in Fig. 3. Finally, the processing results can be optionally downloaded via http or ftp.

Figure 3. View of ongoing task processing status: green (completed), yellow (active).

3. RSS RESOURCES

The RSS service makes available in its processing environment several applications based on EO algorithms covering different thematic areas such as Land, Marine, Atmosphere, Security and Emergency response. Such categories were defined on past 2011 in order to match the nomenclature of GMES service domains, and can be also find in the recent Copernicus services. Among the available applications, it is possible for instance to compute a SAR interferogram using the proprietary GAMMA Remote Sensing [4] processor or the open source DORIS one [5], perform multi temporal analysis using the Parallel Small BAseline Subset (P-SBAS) algorithm provided by CNR-IREA [6], apply Persistent Scatterers Interferometry using

[1] http://eogrid.esrin.esa.int/cloudtoolbox
[2] http://gpod.eo.esa.int

StaMPS [7], create MERIS global vegetation maps at regional or global levels, retrieve land surface temperatures, or use the first of the Sentinel Toolboxes , e.g. Sentinel-1 Toolbox [8], or the operational processors integrated in the environment, such as CryoSat-2 SAMOSA model processor [9]. All these tools can be freely tested and used on available datasets.

An extended list of SAR data processing services available on the processing on-demand environment is shown in the Tab. I:

TABLE II.
LIST OF AVAILABLE SERVICES SUPPORTING SAR DATA.

Service	Supported Input Data
Antartica ASAR Mapping System	ASAR GM1
ASAR-P	ERS-1/2, ASAR L0
BIOMASAR	GAMMA ASAR Tiles output
DORIS	ERS-1/2, ASAR IMS SLC
FAIRE	ERS-1/2, IMM, IMP, ASAR IMM, IMP, WSM
GAMMA ASAR Tiles	ASAR GM1, IMM, WSM
GAMMA DInSAR	ERS-1/2, ASAR IM SLC
GAMMA-L0	ERS-1/2, ASAR L0
Global Human Settlements Layer	ASAR IMM
GlobIce	ERS-1/2, ASAR IMP
NEST	ERS-1/2,ASAR, CSK, TSX, Radarsat-2, ALOS PALSAR
PF-ASAR	ASA L0
PF-ERS	ERS-1/2 L0
P-SBAS (CNR-IREA)	ERS-1/2, ASAR IM L0
Sentinel-1 Toolbox	Sentinel-1, ERS-1/2, ASAR, ALOS PALSAR, CSK, TSX, Radarsat-2
SARINvatore for Cryosat-2	SIR_SIN_FR
SARvatore for Cryosat-2	SIR_SAR1FR
STAMPS	DORIS output

The aforementioned services are integrated and made available with the functionality requested by the PIs. It is worth to notice that these services are continuously improved according to new requests or suggestions from the users.

There is a permanent ESA call for proposals[3] offering to EO scientists and service developers, the possibility to perform bulk processing and/or validation of their own algorithms resorting to the large amount of ESA EO data by means of the RSS processing on-demand environment.

In such environment, parallelisation and high-performance computing resources, exploiting the Grid and Cloud technologies based on the Globus toolkit middleware [10], provide the necessary flexibility to allow quick access to data and fast production of processing results.
The RSS physical infrastructure counts about 60 processing nodes with a total of about 370 CPU cores. This represents the RSS base capacity that is on average sufficient to satisfy users' processing requirements. When the processing requests exceed the RSS base capacity, it is possible to scale up the resources by seamlessly federating additional clusters on the Cloud (see Fig 4). The processing resources are managed by TORQUE [11] and the jobs are scheduled by MAUI [12] software.

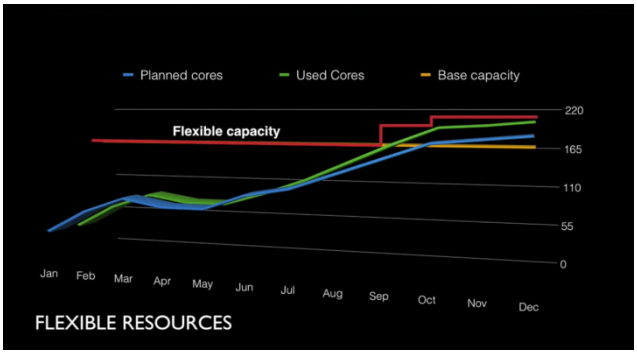

Figure 4. Flexible computing resources example.

From May 2015, RSS is federating a processing cluster on the Teide-HPC[4] Cloud infrastructure, composed of 20 worker nodes, 8 VCPUs and 32GB RAM each one, and a total of 10TB shared memory for data processing.

RSS adapted its datafarm and the data management model in a way that all the different storage components appear as a single access point (see Fig.5). The datafarm model brings several benefits such as: (i) the optimisation of the storage space utilization (ii) an easy access control and (iii) a more efficient and flexible scalability. Currently the total data volume available through RSS is over 700TB, and comprises ENVISAT, ERS, SMOS, SWARM, CRYOSAT2, and third party mission datasets. The storage capacity is divided in different facilities mainly located in the Netherlands and Italy. All the current Copernicus data is also indexed in the RSS data catalogue, which is synchronised every 6 hours with the Sentinels Scientific

[3] http://eopi.esa.int/G-POD
[4] http://teidehpc.iter.es

doi: 10.2788/854791

Data Hub, downloading the data runtime and, therefore providing as well a virtual view of Copernicus data.

The data management in RSS is maintained by a software called NRTservice, partially available in sourceforge[5], which take care of the data transfer from the remote servers to the data storage, indexing in catalogue and processing in the local servers.

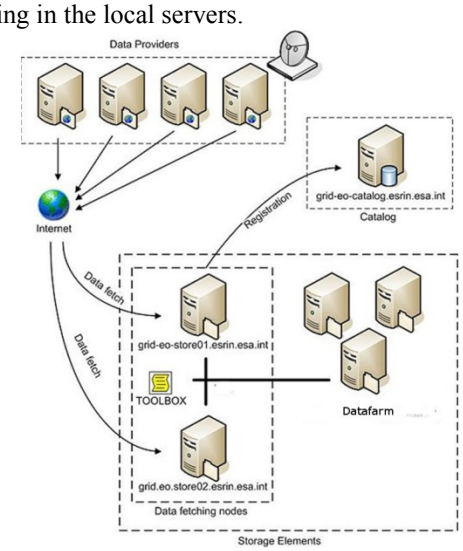

Figure 5. Datafarm schematic representation

RSS provides data processing and access to all current and past ESA EO missions (ERS, Envisat, Earth Explorers), Third Party missions, as well as to data from Copernicus missions. Besides, RSS offers a data provisioning service which can be requested by authorised users, making available on-demand specific datasets if not present yet in the RSS data catalogue. Hence, the EO scientific community and downstream service providers accessing and using RSS resources will experience even greater benefits for all the activities related to the EO research process, including algorithm development, data access, processing and analyses.

4. CONCLUSIONS

The ESA RSS service has successfully supported during the past years many PIs, researchers and service developers. The RSS services allow to the EO user community significant savings in terms of time (months), processing resources (hundreds of CPU-hours), and storage (tens of TB).

The ESA RSS processing on-demand service enables the exploitation of EO data from ESA historical archives, and is also able to support processing of ESA and TPM data, including as well Sentinel-1, Sentinel-2 and next Sentinel-3 data with own algorithms or CoTS.

5. REFERENCES

[1] Marchetti P.G., Rivolta G., D'Elia S., Farres J., Mason G. and Gobron N., "A Model for the Scientific Exploitation of Earth Observation Missions: The ESA Research and Service Support." *IEEE Geoscience and Remote Sensing Society Newsletter* , Vol. 162, pp. 10-18, March 2012.

[2] REFERENCE MODEL FOR AN OPEN ARCHIVAL INFORMATION SYSTEM (OAIS), "Recommendation for Space Data System Practices". *MAGENTA BOOK CCSDS* 650.0-M-2 June 2012

[3] Cuccu, R.; Sabatino, G.; Delgado, J.M.; Rivolta, G., "Enabling SAR data exploitation by processing on-demand", *Geoscience and Remote Sensing Symposium (IGARSS), 2015 IEEE International* , vol., no., pp.1476-1479, 26-31 July 2015

[4] C. Werner, U. Wegmüller, T. Strozzi and A. Wiesmann, "GAMMA SAR and interferometric processing software", *ERS - ENVISAT Symposium,*Gothenburg, Sweden, 16-20 Oct. 2000.

[5] B. Kampes, R. Hanssen, Z. Perski, "Radar Interferometry with Public Domain Tools", *Proceedings of FRINGE 2003*, December 1-5, Frascati, Italy, 2003

[6] S. Elefante, R. Cuccu, I. Zinno, C. De Luca, M. Manunta, G. Rivolta, R. Lanari, F. Casu, "G-POD Implementation of the P-SBAS DINSAR Algorithm to process Big Volumes of SAR data.", *Proceedings of BiDS14 (Conference on Big Data from Space 2014)*, ESA-ESRIN, Frascati Italy, pp. 158-161, November 2014.

[7] A. Hooper, "A multi-temporal InSAR method incorporating both persistent scatterer and small baseline approaches.", *Geophysical Research Letters*, Vol. 35, 2008

[8] Veci, L., Lu, J., Prats-Iraola, P., Scheiber, R., Collard, F., Fomferra, N., & Engdahl, M. "THE SENTINEL-1 TOOLBOX". *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS), 2014, pp. 1-3.*

[9] Benveniste, J., Dinardo, S.; Lucas, B. SAR "Processing on Demand Service for CryoSat-2 and Sentinel-3 at ESA G-POD." *Geophysical Research Abstracts,* Vol. 17, EGU2015-14850, 2015, EGU General Assembly 2015.

[10] Foster, I. "Globus toolkit version 4: Software for service-oriented systems." *IFIP International Conference on Network and Parallel Computing*, Springer-Verlag LNCS 3779, pp 2-13, 2006.

[11] G. Staples, Torque resource manager, in Proceedings of the 2006 *ACM/IEEE conference on Supercomputing*, *ser. SC '06*. ACM, 2006.

[12] Capit, N., Da Costa, G., Georgiou, Y., Huard, G., Martin, C., Mounie, G., Neyron, P. and Richard, O. "A batch scheduler with high level components". *Proceeding CCGRID '05 Proceedings of the Fifth IEEE International Symposium on Cluster Computing and the Grid (CCGrid'05)* - Volume 2. Pages 776-783, 2005 . IEEE Computer Society Washington, DC, USA.

[5] https://sourceforge.net/projects/nrtservice/

PREDICTION OF SYSTEM PREMATURE AGEING BASED ON FUNCTIONAL DECOMPOSITION OF TIME SERIES REPETITIVE MOTIFS AND NON-LINEAR MODELING

Loïc Boussouf[†], *André Cabarbaye*[§], *Jean François Gajewski*[†], *Celestino Esteves*[†], *Clémentine Barreyre*[†]

[†]Airbus Defence & Space, [§]Centre National d'Etudes Spatiales

ABSTRACT

In geostationary telecommunication satellites industry, modern platform can produce up to 10 000 observables every 20 seconds on a regular basis. This leads to hundreds of terabytes for a satellite manufacturer fleet. While all those observables are stored, few of them are real time monitored. This paper proposes recent advances towards a scalable methodology allowing the monitoring of ageing systems based on dynamic signal analysis. First we demonstrate on real data that scalar observation and threshold definition is not appropriate as a general methodology to reflect health of a system. Then, we propose to apply state of the art functional decomposition techniques and nonlinear surrogate modeling to observe deviation of repetitive pattern in telemetry time series and correlate it with system ageing. This method is then applied on batteries data history over 30 satellites, representing half a billion data points, equivalent to 300 years in space if compared to a mono battery satellite. Then, this raises additional questions about correlation between several observables to allow scalability of this approach.

Index Terms—data mining, time series, health monitoring, telemetry, battery, functional decomposition, surrogate modeling, ageing modeling, satellite, in orbit support

1. INTRODUCTION

CNES agency has launched a R&D initiative on Health Monitoring in 2014 that has been renewed in 2015. Objective was to propose way to update fiabilist teams expectations on satellite health based on its observables. This initiative met Airbus strategy to valorize telemetry data. Expected outcomes are multiple. First, adapt our way to operate a given platform based on updated health status: will it last longer than expected? Can we change our mission profile to make it last more time than its initial design lifetime? Secondly, a consolidated feedback on all data history is expected to produce guidelines for next platforms. Do we tend to oversize our designs? Do some equipment/suppliers differ significantly from others? To give us means to address those questions, we decided to build up a methodology over battery expected deviation over time.

Reading documented methods in health monitoring over various industries, we observe that physical knowledge is preponderant, either through expert definition of monitoring threshold levels, as in [1] and [2], or through physical and numerical simulation to predict durability of material and direct observation to observe degradation in situ as in [3]. But as suggested in [4], the observation of a functional data (such as temperature distribution over an optic fiber in [4], or such as a time series in our case) can bring much more information than punctual measurement observation. This is the approach we want to develop: how to take advantage of all measurements samples and not only measurement at a specific moment to get an accurate model representing our equipment health?

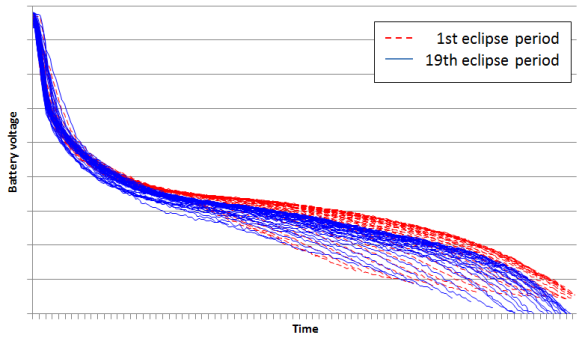

Figure 1 Battery discharge voltage over time (no scale) : first year (red), nine year later year (blue)

2. TEST CASE DESCRIPTION

Chosen test case is battery. This equipment has no observed failure on Airbus DS platforms, and is already well known from battery specialists. But it is representative of the equipment category we want to address with our methodology as it produces a set of observables every 20 seconds. In eclipse phases, battery discharges and recharges, hence presenting a repetitive dynamic behavior. We want to take advantage of this well-known test case to challenge a functional based approach, considering the whole signal dynamic.

In Fig. 1, we observe all discharge cycles during first eclipse period (dashed red) and nineteenth eclipse period (plain

blue), hence spanning over 9 years of battery lifetime (a geostationary platform knows two eclipse period a year).
We see a deviation between the two envelopes which illustrates expert knowledge: the older the battery is, the lower in voltage the discharge cycles are. This is expected as described in [7] where authors compared min cell voltage at eclipse end, for various discharge capacity on a similar platform between 2005 and 2013. Fig. 2 reproduces only the longest eclipse discharge for all 19 eclipse in satellite lifetime, zooming on central time (same conclusion applies when looking at starting values, ending values, or any other values. It shows that a punctual measurement is not enough to compare two discharge cycles, due to intrinsic measurement noise and local variation due to different battery discharge current. Noise is high enough to decorrelate any punctual measurement value from age. Correlation between sequence of values at a given time and age varies between 0.02 and 0.75 depending on considered time, with an average value of 0.56, not enough to model any relation between battery voltage and battery age.
To avoid comparison between cycles being too sensitive to intrinsic noise, we propose to rely on a way to consider whole dynamic of discharge cycle instead on single point measurement.
Concerning variations due to different payload consumption profile, it is an identified way of improvement of current paper results, and we will address this through a multi telemetry analysis, allowing comparison of discharge profile with different bus current (in a future study).

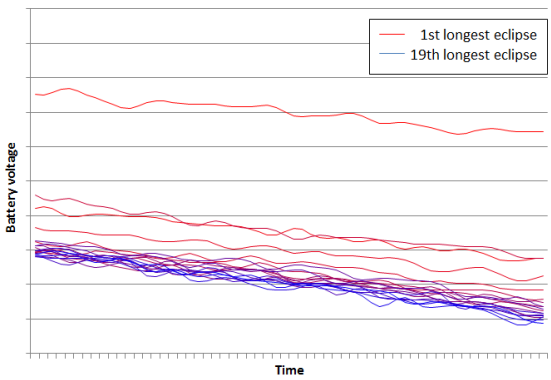

Figure 2 Illustration of intrinsic noise through 9 year of discharge cycles (spanning all shades from red to blue)

3. FUNCTIONAL DECOMPOSITION TECHNIQUE

We consider a fleet of 30 telecom geostationary satellites. Among all discharge cycle occurring in a same period of successive eclipses, we consider only the 5 longest, having similar length and depth of discharge, hence comparable. This represents 2821 discharge cycles for the whole fleet. Each discharge cycle is about 170 samples long due to eclipse duration in geostationary orbit.
Telemetry sampling may not be regular, and two cycles may not share the exact same discretization. Then all cycles are re-interpolated using the same regular discretization using a piecewise cubic spline. Final sample interval chosen is 10 seconds, where usual telemetry is interpolated every 20 seconds. To make all data comparable, we truncate all cycles to have all cycles the same length. This results in losing up to 5 samples from original data. While methods exist to allow comparison of data of different length, such as Dynamic Time Warping, it is identified as a way of improvement of current study, and will not be addressed here.
The objective is now to be able to predict age of a battery based on those discharge cycles. First approach would be to model the relation $\mathbb{R}^{170} \rightarrow \mathbb{R}$ given by $battery\_age = f(discharge\_cycles)$. But for each of those cycles, the 170 samples result from discretization of a continuous signal, hence they are highly correlated. Building a model with highly correlated inputs is discouraged. Hence, we project the whole data set, lying in $\mathbb{R}^{170}$to a smaller dim space using a dimension reduction technique.
While several other dimension reduction techniques exist, we choose PCA as it is a technique frequently found in literature as in [5] whose application is frequent as a first step to discover features in the data [8]. PCA builds a function basis based on dataset, hence very efficient to decompose similar functional entries and allowing reducing space dimension by concentrating most of input data variance in first components. This characteristic of building the functional basis from the dataset itself allows reducing expertise implication to define an appropriate function basis and is a step towards application of the exact same method to different data.
To apply PCA on functional data, we rely on the discretization of our functional data and use multivariate PCA, where each discharge cycle is one observation of 170 variables. Hence, each discharge cycle f_j sampled at x_i can be decomposed as the following:

$$f_j(x_i) = \sum_{k=1}^{p} \alpha_{jk} \phi_k(x_i)$$

As we have all f_j observations discretized, each of our observations can be represented with a finite p value, while functional PCA requires infinite decomposition to allow exact decomposition of functions. But in both techniques (multivariate and functional PCA), variance decreases along principal components, allowing dimension reduction through principal component basis truncation.

Choosing a reduced space dim of 20 principal components explain more than 97% of variance of initial dataset as seen

doi: 10.2788/854791

on Fig. 3, with orthogonal components, hence uncorrelated. After projection base has been truncated to 20 components, reconstruction of initial dataset showed 0.2% average error, measured on the difference for each sample between initial discharge cycles and reconstructed cycle. This is considered negligible.

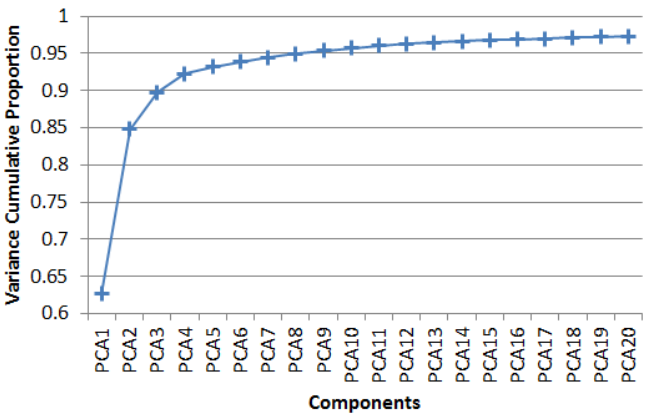

Figure 3 Variance cumulative proportion along PCA components for family 1

After reduction and before trying to model age from discharge cycles, we must validate homogeneity of initial dataset, which is made easy through reduced dimension space available through PCA decomposition on centered data. Projection representations along first dimensions show clearly 2 families of data as seen on Fig. 4.

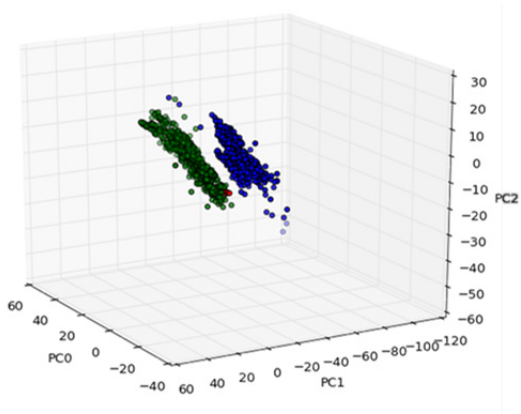

Figure 4 Projection along 3 first PCA dimensions shows easily separable clusters (NiH2 in blue, Li-Ion in green)

After exchange with battery experts, it appeared that those clusters, find out through PCA, reflects the different technologies or architecture of batteries available in the dataset: bi battery NiH2 and bi battery Li-Ion. Expert brought us additional information to separate accurately those clusters and identify an additional architecture (mono battery Li-Ion). Physical explanation behind those clusters is the differences in battery capacity allowed by those technologies and battery module number. We keep only first and third families which have enough data to allow further analysis modeling. Given the fact that remaining family (mono battery Li-Ion) was not distinct from other families (bi battery Li-Ion) in this projection, we'll assess in a future study the possibility to consider it as similar to one of the retained family despite expert knowledge considering them as different. This demonstrates in this test case the PCA ability to project functional data keeping important features, reflecting in that case battery technology difference.

4. NON LINEAR MODELING OF AGE

For each discharge cycle, we can associate eclipse period index (2 eclipse periods per year) starting from launch date. This is equivalent to battery age, and we call that value "real age" of battery. We then have a dataset of inputs lying in $\mathbb{R}^{20}$ (PCA decomposition coefficient of each discharge) and outputs lying in $\mathbb{R}$ (eclipse index). Building a surrogate model of those outputs *wrt* inputs will give use the ability to predict age for a new set of reduced inputs, hence from a new discharge cycle. Such a model is built using gaussian process method (or kriging), a method that showed very good results to catch highly nonlinear physical phenomenon as demonstrated in [6].

Kriging supposes the data as the sum of a deterministic model m (through any regression model, while low degree polynomials are preferred) and the realization of a centered gaussian process Z characterized by its two moments, mean (centered so equal to 0) and variance:

$$Y(\alpha_j) = m(\alpha_j) + Z(\alpha_j)$$

With

$$E\left(Z(\alpha_j)\right) = 0$$

$$cov\left(Z(\alpha_i), Z(\alpha_j)\right) = \sigma^2 . R(\alpha_i, \alpha_j)$$

Where σ^2 is Z variance, R is the chosen correlation function and α_j is the vector of the 20 coefficients describing PCA decomposition of the j discharge cycle. Its ability to catch non linearity is due to its variance matrix and its variance function. Those are calibrated on training dataset through the maximization of the max likelihood function. After training phase, Y can be used as a prediction model using new data. Due to its stochastic nature, it proposes a mean for a new α_j which is used as an estimation of discharge cycle index. Its variance is used to build confidence interval of those estimations. In this study, we will not use the variance of our estimations. Kriging model is an interpolator, hence going exactly through points in training data set. It is then required to keep a separate validation dataset to perform error measures. Model is built on two thirds of data set,

randomly picked, and validated on the remaining third, by the way avoiding over fitting of data.
As input dataset is not homogeneous and two families were discovered, a model is built independently for each of the two kept families, and age is predicted on validation dataset. An ideal model on ideal data (where output variance is 100% explained by input variables) would have a predicted age equal to real age on validation dataset. We represent this on Fig 5, with red envelope representing an error of ± 1 eclipse index, equivalent to ± 6 months. All results are summed up in Table 1.

	$\pm$ 6 months	$\pm$ 12 months
Family 1	67.5%	88.9%
Family 3	38.9%	67.4%

Table 1 Error distribution for each family

Methodology shows very good fidelity for family 1 (NiH2), where about 90% of estimations are in an error range of ± 12 months. This means that receiving new discharge cycle from a satellite with NiH2 batteries, we are able to compare its real age and its modeled age. Any deviation beyond 1 year between the two values should be noticed to battery expert for further investigations. Concerning family 3 (Li-Ion), it badly fits observation. If used in operations, it would raise too many false alarms.
After exchange with battery experts, it appears that only age can't explain deviation in observed discharge cycles. Bus current and mission profile can impact discharge cycles, making them looking older than they the battery really is. Identified way to improve the model is to integrate an additional telemetry which is bus current. Then modeling will be:

$$battery\_age = f(discharge\_cycles, bus\_current)$$

5. CONCLUSION AND WAY FORWARD

After this first study to model battery age with no physical assumptions and only the time series analysis of battery voltage over time, demonstration has been made that correlation can be modeled between battery age and functional telemetry. Methodology is now defined and those encouraging results confirm the possibility to use statistical tool to have an insight in satellite systems ageing, based on historical telemetry data. Lack of strong physical assumptions shows that method is scalable to other systems monitoring.
Future work will focus on considering additional telemetries to explain better the variation in modeled system age due to mission load changes. And on another hand, we will work at linking such model outputs to physical and fiabilist analysis, to make them operational as a tool for satellite system health monitoring.

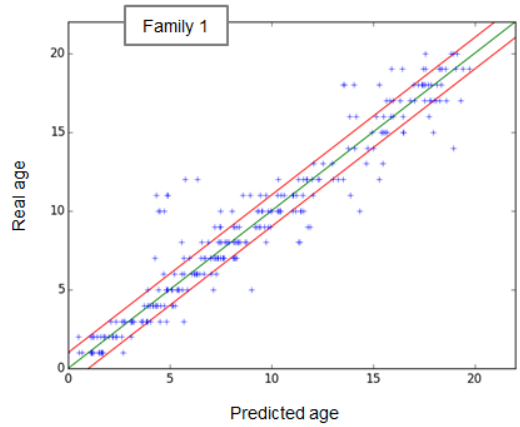

Figure 5 Predicted age vs real age for family 1 model

6. REFERENCES

[1] X. Flandrois, J. Lacaille, J.R. Masse and A. Ausloo, "Expertise transfer and automatic failure classification for the engine start capability system" *AIAA Infotech at Aerospace proceedings*, Seattle, 2009
[2] A. Getman, C.D. Cooper, G. Key, H. Zhou and N. Frankle "Detection of mobile machine damage using accelerometer data and prognostic health monitoring techniques" *Computational Intelligence in Vehicles and Vehicular Systems proceedings*, IEEE, Nashville, pp. 101-104, 2009.
[3] D. Roach and K. Rackow, "Development and Validation of Bonded Composite Doubler Repairs for Commercial Aircraft" Sandia National Laboratories report, 2007
[4] A. A. Khan, P. Cunat, Y. L. Beck, J. I. Mars, V. Vrabie, J-P. Fabre, Distributed fiber optic temperature sensors for leakage detection in hydraulic structures, *5th World Conference on Structural Control and Monitoring*, Tokyo, 2010
[5] S. Lindsay, D. Myung-kyung Woodbridge, Spacecraft State-of-health (SOH) Analysis via Data Mining, 13th International Conference on Space Operations, Pasadena, 5-9 May 2014
[6] L. Boussouf, Surrogate Based Optimization for Multidisciplinary Design, *SAE AeroTech Congress & Exhibition*, SAE International, Toulouse, 2011
[7] F. Malet, V. Leconte, Y. Borthomieu, F. Bauchais, Overview and in-orbit behaviour of the first lithium-ion batteries used onboard Eutelsat 7a and Eutelsat Hot Bird 13b geo telecommunications satellite - Update after ten years in orbit, *10th European Space Power Conference*, Noordwijkerhout, 2014
[8] J.O. Ramsay, B.W. Silverman, Functional Data Analysis, 2nd edition, Springer, p147-164

SKYQUERY: A PARALLEL DATABASE PLATFORM FOR ON-DEMAND CROSS-MATCHING

László Dobos[1], *Tamás Budavári*[2], *Evelin Bányai*[1], *István Csabai*[1] *and Alexander S. Szalay*[2]

[1]Eötvös Loránd University, Dept. of Physics of Complex Systems, Budapest, Hungary
[2]The Johns Hopkins University, Dept. of Physics & Astronomy, Baltimore, USA

ABSTRACT

SkyQuery is a scalable, parallel query engine built on top of a cluster of relational database servers to provide on-demand cross-identification of astronomical datasets based on celestial coordinates. It implements a Bayesian statistical framework to take varying positional errors into account. The Bayesian framework supports incremental matching of multiple catalogs and allows for specifying custom priors when not only the coordinates and the astrometric error but other parameters, such as the color indices, form the basis of matching. Cross-matching can be efficiently limited to an area of interest by rapid pre-filtering of large data sets when combining survey data with small area catalogs. To partition the cross-match problem and create balanced parallel jobs for execution, SkyQuery gathers statistics on-the-fly using randomly sampled, small versions of the large catalogs. SkyQuery draws on the idea of interacting with databases: cross-match problems are formulated in an extended version of the SQL language and are automatically compiled into a distributed workflow of traditional SQL queries that are executed on multiple machines, in parallel. Query results are stored in a personal database of the user called the MyDB, which is closely integrated with the SDSS SkyServer. Result-sets stored in the MyDB can be further processed using SQL or downloaded in a rich set of data formats. Users also have the possibility to upload their own coordinate lists to MyDB and cross-match them with the largest catalogs.

Index Terms— astronomical databases, cross-match, Virtual Observatory, distributed databases, query processing

1. INTRODUCTION

Cross-matching astronomical catalogs is not a one time task. While large catalogs can be matched in pairs once, multi-catalog matching poses a combinatorically exploding problem. Also, a customized approach is almost always necessary to filter out false matches, usually by taking the spectral energy distribution into account. Very often users have their own set of coordinates that need to be matched against the major astronomical catalogs. Cross-matching is a common problem the most astronomers face and one of the central services that one expects from the Virtual Observatory. Other work addressing the issue includes those of [1, 2].

The original SkyQuery [3, 4, 5] was developed more than a decade ago but failed to reach a large audience, mostly due to its lack of scalability which forced us to limit cross-matches to a mere 5,000 objects per job. We redesigned SkyQuery from scratch with scalability and flexibility in mind. The result is a scalable, parallel query engine built on top of a cluster of traditional database servers. SkyQuery provides direct access to the photometric catalogs of most public domain sky surveys and the number of available data sets is growing. Databases are co-located at out data center at the Johns Hopkins University to provide high-bandwidth access to data. The system is designed such a way that the cross-match algorithm runs right inside the database server process, reducing data movement to the minimum.

In SkyQuery, cross-match jobs are scriptable so users do not have to click through complex forms to select the necessary catalogs and specify the matching criteria. Since the SDSS SkyServer [6], scientists have become familiar with SQL to access and process astronomical data available as a relational database. We took SQL a step further and extended its syntax to accommodate the description of complex cross-match jobs, see Sec. 3.

2. BAYESIAN FRAMEWORK FOR CROSS-MATCH

In case of optical or near-optical catalogs of galaxies and stars the morphology of the objects is simple enough to characterize their celestial positions by a pair of spherical coordinates. The coordinates usually correspond to the center of the PSF of a star or the brightest pixel of a galaxy. The reliability of coordinate measurements depends on many factors but the positional error in the aforementioned bands is usually a few tenths of arc seconds. As a consequence, very good matching of catalogs can be achieved by matching based only on coordinates. We want to estimate the probability of the hypothesis H that multiple observations with coordinates $D = x_1, x_2, ..., x_n$ belong to the same celestial source with coordinates m. Rather than calculating the $P(H|D)$ posterior probability, it turns out that calculating

The research project was performed under collaboration with the NSF-funded SciServer project, Award No. ACI-1261715 and was supported from the Hungarian Scientific Research Fund grant OTKA NN 114560.

the ratio of the probability of the hypothesis, given the data, and the complementary hypothesis K (not *all* observations are from the same source), i.e. calculating the Bayes factor, is more convenient. The Bayes factor by definition is $B(H,K|D) = \frac{P(H|D)/P(H)}{P(K|D)/P(K)}$. For independent observations with positional error σ_i, the Bayes factor becomes

$$B(H,K|D) = \frac{\sinh w}{w} \prod_{i=1}^{n} \frac{w_i}{\sinh w_i},$$

where the weights $w_i = \frac{1}{\sigma_i^2}$ and $w = |\sum w_i x_i|$. Since the typical positional error is on the order of arc seconds the w_i is usually very large and the Bayes factor becomes

$$B = 2^{n-1} \frac{\prod w_i}{\sum w_i} \exp\left(-\frac{\sum_{i<j} w_i w_j \psi_{ij}^2}{2\sum w_i}\right),$$

where ψ_{ij} is the angular separation between the observations in radians. One great advantage of this formula is that it (more specifically its logarithm) can be calculated iteratively, namely by taking only the previous best match coordinates and the next catalog into account, reducing the cross-match problem to pair-wise matching. The matching criteria are established by setting a lower limit on the Bayes factor after matching all catalogs. This is more practical than setting a limit on the posterior probability for many reasons, but most importantly, cutting on the Bayes factor does not require a prior, which in non-trivial to calculate. To find the posterior probability of a match one needs to take into account at least the footprint, the flux limit and the cardinality of the catalogs. The Bayes factor, on the other hand, measures the *evidence* the data provides in favor of the hypothesis versus the complementary hypothesis. A typical value of $B > 10^3$ is considered a very strong evidence. When matching multiple catalogs incrementally, the logarithm of the Bayes factor becomes a sum of non-negative numbers. B increases monotonically, but at every step an upper limit B_{max} on the final Bayes factor of the match can be calculated. As the incremental matching progresses, B_{max} decreases monotonically, making it possible to filter out false matches early and only carry good candidates to the next step. Bayes factors of independent hypotheses can be multiplied together. For example, if one calculates the Bayes factor B_{pos} of the observations based on the positions and B_{SED} on the spectral energy distribution the combined Bayes factor becomes $B = B_{\text{pos}} B_{\text{SED}}$. This property can also be used to process multi-catalogue cross-matches in a cascading manner, which in turn will enable us in the future to distribute cross-match computation efficiently among data centres.

In practice, one starts with setting the final upper limit on B, everything else (positions, positional error) is available in the catalogs. At each step the upper limit on the angular separation is computed and the database problem becomes finding nearby objects from catalog B to each object in catalog A. SkyQuery uses an optimal solution to this problem, the so called zone approach [7].

3. BASIC CROSS-MATCH SYNTAX

To make SkyQuery "scriptable" cross-match problems are formulated in an extended version of SQL. This way the same cross-match can be rescheduled with slightly modified parameters very easily, without having to fill in complex web forms. The SQL interface also makes the system very flexible and allows for adding custom constraints and joining in further tables. The syntax also supports specifying the area of interest via the `REGION` clause. The following sample query illustrates the use of custom SQL extensions. (A few details are omitted from the query for the sake of simplicity.)

```
SELECT s.objid, s.ra, s.dec, g.objid, g.ra, g.dec,
       x.matchid, x.ra, x.dec
INTO twowayxmatch
FROM XMATCH (
  MUST EXIST IN SDSSDR7:PhotoObjAll AS s,
  MUST EXIST IN GALEX:PhotoObjAll AS g,
  LIMIT BAYESFACTOR TO 1e3 ) AS x
INNER JOIN SDSSDR7:SpecObj sp
  ON sp.bestObjID = s.objID
WHERE s.petroMag_r < 18 AND sp.z < 0.05
REGION 'CIRCLE J2000 0.0 0.0 10.0'
```

The `XMATCH` construct lists the catalogs that are to be matched. It also specifies the lower limit on the Bayes factor of accepted matches. Note that the table alias `x` is used to access the columns of the matched catalog such as the match identifier, the best match coordinates, the updated positional error, the Bayes factor, etc.

The `REGION` clause of a cross-match query limits the area of interest to a certain part of the sky. In the example above we limit the search to a circle with a radius of 10 arc min but arbitrarily complex regions are supported that can be represented by the Sperical Library of Budavári et al. [8]. Currently, the region constraint needs to be specified as a string but future versions of SkyQuery will be able to fetch region descriptions from an http URL, for example from the Virtual Observatory Footprint Service [9]. SkyQuery uses Hierarchical Triangular Mesh indexing [8] to speed up search by spatial constraints.

4. QUERY PARTITIONING AND EXECUTION

In order to parallelize query execution the cross-match job needs to be partitioned into equally sized sub-jobs. Cross-match query can contain a `REGION` constraint and an arbitrarily complex `WHERE` clause. As a consequence, jobs cannot directly be partitioned based on the sky coverage of the catalogs but also the user-specified constraints need to be taken into account when the partitions are determined. To determine the best partitioning, for every catalog referenced by a cross-match query we calculate a *conditional* histogram of the

doi: 10.2788/854791

spatial distribution of the sources. This is done by determining the most restrictive `WHERE` clause that can be applied to the catalog table without losing any objects. Of course, computing the statistics of large catalogs would be prohibitively slow, consequently, we created uniformly sampled, small versions of all data sets available in SkyQuery. Statistics queries are executed on these small catalogs rather than on the large ones to approximate the distribution of sources in the final result set. Because the zone algorithm [7] stripes the sphere by declination, we set the partition limits in right ascension. The limits are set such that the number of objects is approximately equal in each partition. This procedure is absolutely necessary for a parallel system to avoid skew among job partitions.

While the underlying SQL Server can be leveraged to execute traditional SQL queries efficiently, optimization of cross-match queries had to be implemented manually. Depending on the filter criteria and the catalog tables involved in the matching columns have to be pre-computed or even indices built on-the-fly for the system to yield optimal performance. As a simple example, we mention that Cartesian coordinates sometimes need to be computed during query execution from spherical coordinates or an HTM index needs to be built on the table before applying a region constraint. Since the database server is highly optimized for index creation, on-the-fly indexing is almost always beneficial when no readily usable index is available. Finding the optimal path for cross-match query execution is based only on the availability of indices, no statistics or computing costs are taken into account.

Query execution is done in two major steps. First the cross-match part is executed to produce a match table, then, if necessary, the match table is joined with the rest of the tables of the query. This is the point where filter constraints that could not be evaluated during the cross-match process get finally applied. Query partitions are distributed among the worker nodes at each step for parallel execution and results are gathered to be saved in MyDB as a last step.

5. SYSTEM INFRASTRUCTURE

The computing infrastructure supporting SkyQuery is illustrated in Fig. 1. A controller node is responsible for system coordination while a set of identical database servers are used to serve catalog databases and execute actual cross-match computation. The job scheduler process is running on the controller node, therefore, jobs must be designed to put minimal load on the controller machine. Jobs rather off-load computationally intensive and IO demanding tasks to the worker nodes. Typical off-loaded tasks are either SQL queries or delegated via an RPC service running on every machine.

SkyQuery is currently running on three high-performance servers with 24 CPU cores each, 256 GB of RAM and 32 TB

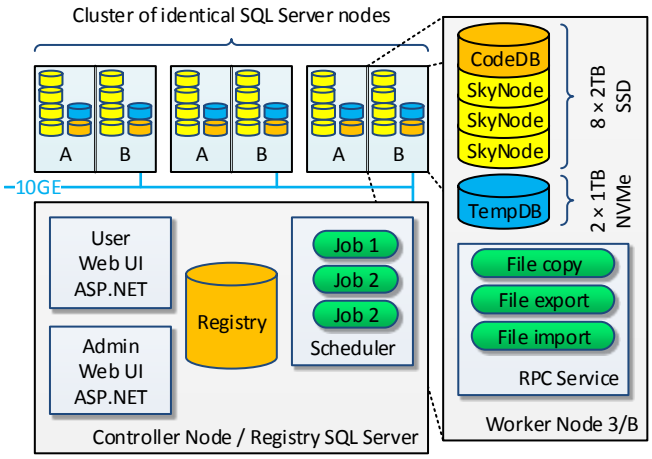

Fig. 1. Overview of the SkyQuery infrastructure. A controller node is responsible for coordinating job execution by a parallel scheduler. Worker nodes run two instances of SQL Server to serve the catalogs and execute cross-match sub-jobs. An RPC service is listening on every server to support task delegation.

consumer-grade solid state storage. The 16 Samsung 850 EVO SSDs are distributed across three Avago SAS3 3008 controllers that have a theoretical full-duplex bandwidth limit of 8 GB/s. During performance testing, we were able to achieve sequential read speeds above 4 GB/s per server. In addition to the 2 TB SSDs used as main storage, 4 PCI-e Intel NVMe storage modules (4.4 TB total) are installed in each node to store temporary data and support sorting. Network traffic is conveyed over 10G ethernet. The choice of SSDs over spinning disks, although their superior performance is indisputable, requires some justification as their degradation due to frequent writes is a well-known phenomenon. Fortunately, in case of SkyQuery, with the exception of temporary tables and MyDBs, databases are read-only, and once they are built and indexed, no further updates to them are necessary. All frequently written databases are allocated on the enterprise-grade NVMe devices.

Worker nodes hold identical sets of databases in a mirror-everything configuration. While it certainly requires a lot of disk space, one important advantage of mirroring everything is that all jobs, restricted to any small portion of the sky, can be partitioned and parallelized. This would not be possible if data were split across the servers based on, for example, sky coverage. This latter approach is prone to the problem of "sweet spots", the practice of querying the same interesting part of the sky over and over which would result in excessive load on a single server.

Since SSDs are sometimes reported to show whimsical behaviour when used with hardware RAID controllers, we abandoned the RAID idea entirely and opted for file-level parallellization and mirroring. SkyQuery nodes are config-

ured to run two database server instances in parallel, 8 SSDs associated with each. Database file groups are configured such that they have a file on each of the 8 volumes. With this setup scan operations get parallellized automatically by SQL Server. Also, since the databases are mirrored to both serves instances, in case of a SSD failure, exact copies of the database files are readily available within the same box to restore the failing volume.

6. USER INTERFACE, REST API AND SCISERVER

SkyQuery is accessible via a web interface, very similar to CasJobs of the SDSS SkyServer [10]. After registration, users can browse the schema of the datasets, write and submit queries and keep track of the submitted jobs. The web-based user interface also provides basic data import and export functionality. As most astronomers prefer direct access to the data from script instead of using a separate user interface, SkyQuery provides a set of RESTful web services for easy integration of the cross-match functionality into user workflows. The *schema service* gives programmatic access to the database schema of the MyDB and the catalogs and also provides rich meta data. MyDB is accessible via the *data surface* which allows uploading and downloading entire tables in various data formats including FITS and VOTable. New cross-match jobs can be scheduled and existing jobs be monitored via the *jobs service*. The standard RESTful protocol and support of the XML and JSON formats make implementing client libraries for SkyQuery relatively simple.

SciServer is a collaborative data analysis platform being developed at the Johns Hopkins University. It aims at integrating database warehousing technologies with non-relational data, OpenStack-based virtualization and distributed data analytics. It provides centralized services for user authentication, cloud storage of scientific data (SciDrive), access to relational database and batch SQL execution (CasJobs), MyDB for user-owned relational data and remote access via the web-based scripting platform Jupyter. SkyQuery is already integrated fully into the SciServer infrastructure.

7. REFERENCES

[1] F.-X. Pineau, T. Boch, and S. Derriere, "Efficient and Scalable Cross-Matching of (Very) Large Catalogs," in *Astronomical Data Analysis Software and Systems XX*, I. N. Evans, A. Accomazzi, D. J. Mink, and A. H. Rots, Eds., July 2011, vol. 442 of *Astronomical Society of the Pacific Conference Series*, p. 85.

[2] F. Pineau, T. Boch, S. Derriere, and Arches Consortium, "Towards a Next-Generation Catalogue Cross-Match Service," in *Astronomical Data Analysis Software an Systems XXIV (ADASS XXIV)*, A. R. Taylor and E. Rosolowsky, Eds., Sept. 2015, vol. 495 of *Astronomical Society of the Pacific Conference Series*, p. 61.

[3] T. Budavári, T. Malik, A. S. Szalay, A. R. Thakar, and J. Gray, "SkyQuery – A Prototype Distributed Query Web Service for the Virtual Observatory," in *Astronomical Data Analysis Software and Systems XII*, H. E. Payne, R. I. Jedrzejewski, and R. N. Hook, Eds., 2003, vol. 295 of *Astronomical Society of the Pacific Conference Series*, p. 31.

[4] T. Budavári, A. S. Szalay, J. Gray, W. O'Mullane, R. Williams, A. Thakar, T. Malik, N. Yasuda, and R. Mann, "Open SkyQuery – VO Compliant Dynamic Federation of Astronomical Archives," in *Astronomical Data Analysis Software and Systems (ADASS) XIII*, F. Ochsenbein, M. G. Allen, and D. Egret, Eds., July 2004, vol. 314 of *Astronomical Society of the Pacific Conference Series*, p. 177.

[5] W. O'Mullane, T. Budavári, N. Li, T. Malik, M. A. Nieto-Santisteban, A. S. Szalay, and A. R. Thakar, "OpenSkyQuery and OpenSkyNode - the VO Framework to Federate Astronomy Archives," in *Astronomical Data Analysis Software and Systems XIV*, P. Shopbell, M. Britton, and R. Ebert, Eds., Dec. 2005, vol. 347 of *Astronomical Society of the Pacific Conference Series*, p. 341.

[6] A. S. Szalay, J. Gray, A. R. Thakar, P. Z. Kunszt, T. Malik, J. Raddick, C. Stoughton, and J. vandenBerg, "The sdss skyserver: public access to the sloan digital sky server data," in *Proceedings of the 2002 ACM SIGMOD international conference on Management of data*. ACM, 2002, pp. 570–581.

[7] J. Gray, M. A. Nieto-Santisteban, and A. S. Szalay, "The zones algorithm for finding points-near-a-point or cross-matching spatial datasets," *arXiv preprint cs/0701171*, 2007.

[8] T. Budavári, A. S. Szalay, and G. Fekete, "Searchable sky coverage of astronomical observations: Footprints and exposures," *Publications of the Astronomical Society of the Pacific*, vol. 122, no. 897, pp. 1375–1388, 2010.

[9] T. Budavári, L. Dobos, A. S. Szalay, G. Greene, J. Gray, and A. H. Rots, "Footprint services for everyone," in *Astronomical Data Analysis Software and Systems XVI*, 2007, vol. 376, p. 559.

[10] N. Li and A. R. Thakar, "CasJobs and MyDB: A Batch Query Workbench," *Computing in Science and Engineering*, vol. 10, pp. 18–29, Jan. 2008.

THE EARTH OBSERVATION IMAGE LIBRARIAN (EOLIB): THE DATA MINING COMPONENT OF THE TERRASAR-X PAYLOAD GROUND SEGMENT

D. Espinoza-Molina[1], *V. Manilici*[2], *O. Dumitru*[1], *C. Reck*[2],
S. Cui[1], *H. Rotzoll*[2], *M. Hofmann*[2],*G. Schwarz*[1], *M. Datcu*[1]
[1]German Aerospace Center(DLR), Remote Sensing Technology Institute (IMF), Munich, Germany
[2]German Aerospace Center (DLR), German Remote Sensing Data Center (DFD),Munich, Germany

ABSTRACT

In this paper we present the Earth Observation Image Librarian (called EOLib) as a new generation of Image Information Mining Systems. EOLib is operated in the Payload Ground Segment of TerraSAR-X. The main goal of EOLib is to provide semantic annotations of satellite image content and offer to the end user a semantic catalogue via a web user interface. Moreover, EOLib has more functionality such as searches based on image metadata and semantics, visual exploration of the image archives, metadata extraction, etc. The system consists of components such as a query engine, knowledge discovery in databases, visual data mining, epitome generation, and user services. EOLib is able to ingest a TerraSAR-X scene with 8000×8000 pixels in about three minutes. The EOLib workflow starts with the ingestion of a scene, it continues with the semantic annotation of the image content based on machine learning methods, and it ends with publishing the semantic catalogue and enabling the search by metadata and semantic image descriptions.

***Index Terms*—** Software Architecture, Image Information Mining, Payload Ground Segment, Data Mining, Earth Observation.

1. INTRODUCTION

Over more than 15 years, major efforts have been made to introduce Image Information Mining (IIM) to Earth Observation (EO) data holdings in order to allow various applications to retrieve all desired information. These efforts have resulted in a major scientific progress in the understanding remote sensing specific principles of IIM, by developing and validating algorithms, methods and systems for EO knowledge discovery and data mining, by evaluating appropriate methods and technologies [1], by analysing the user needs and expectations, and finally, by successfully demonstrating the capabilities of IIM on EO data. These efforts resulted in several IIM system implementations as - for example - the Knowledge-driven content-based Information Mining system (KIM) [2] that presented the concept of image retrieval based on content using a Bayesian approach for the computation of similarities between images. Later, GeoIRIS (Geospatial Information Retrieval and Indexing System-Content Mining) [3] proposed a primitive feature extraction based on patches instead of pixels and semantic modelling. However, later, the problems of matching the image content with semantic definitions adopted by humans became more and more evident, causing the so-called semantic gap [4]. The semantic gap demonstrated the necessity of semantic definitions to be included in the image retrieval. In an attempt to reduce the semantic gap, more systems including labelling or the definition of the image content by semantic names were introduced. For instance, [5] demonstrated that the semantic representation has an intrinsic benefit for image retrieval by introducing the concept of query by semantic example (semantics and content). Later, in addition to the semantic definitions, the combination of several types of data in order to improve the image retrieval was presented in [6].

The European Space Agency (ESA) has been contributing to the development of IIM systems by funding several projects in this field. Among various system implementations, we have to mention the Knowledge-centred Earth Observation (KEO) system prototype. This system permitted users to interactively extract relevant features and information from EO data, and to provide outputs as, for example, valuable information extracted from data, in easily accessible formats. A second example is the KLAUS project, the main goal of which was an overall improvement of the KEO prototype system with a focus on having models for land use management. Currently, in this context and contributing to the big data era, where new mining techniques are necessary due to the volume, variability, and velocity of such data [7], we present EOLib (Earth Observation Image Librarian) as a next-generation of an Image Information Mining system implementing novel techniques for image content exploration and knowledge discovery in databases. EOLib produces information about the content of EO products which is usually hidden in raster data and metadata. EOLib is being operated in the Payload Ground Segment (PGS) of TerraSAR-X (TSX) and will aim at enlarging the IIM scope for a more complete exploration of the EO data sources by establishing large scale information mining functions within the multi-mission PGS operating missions like TerraSAR-X/TanDEM-X, Sentinel-1/2 and similar high resolution Synthetic Aperture Radar (SAR) and optical imaging missions. The EOLib system will allow users to find EO products of interest for their specific applications with semantic concepts. The EOLib system offers functions such as data model generation by means of tiling and the extraction of primitive feature from EO products, visual data mining to browse the image archives, knowledge discovery in databases to define a semantic annotation of the image content, queries based on different parameters via user services, and an epitome production functionality. Thus, EOLib will consist of a set of tools for sustainable long term and efficient utilization of EO data content. Our previous publication [8] discussed the planned EOLib architecture prior to implementation. We had the opportunity to review the architecture in ways that were not foreseeable before implementation. Interfaces were refined (both the connection graph and the interface semantics). Finally, we addressed several critical performance problems.

The rest of the paper is organized as follows. Section 2 presents an overview of the EOLib system architecture. Section 3 describes some examples and Section 4 concludes the paper.

2. EOLIB ARCHITECTURE OVERVIEW

The EOLib baseline architecture is depicted in Fig. 1. EOLib consists of several independent systems that communicate with each other via established interfaces in a service-oriented architecture. Each of these systems may integrate one or several components which provide specific functionality to the system. The novelty in EOLib is that there are several user-oriented components that enlarge the functionality of the Payload Ground Segment (PGS) such as Visual Data Mining (VDM), Knowledge Discovery in Databases (KDD), Query Engine (QE) and the User Services (US). The current version of the EOLib implementation operates with TerraSAR-X L1b products. The rest of this section briefly describes the PGS components represented in blue and the new EOLib components marked in orange.

2.1. Payload Ground Segment Components

The PGS of the German Aerospace Center (DLR) consists of several subsystems such as data acquisition, data processing, data and information management, librarian and archiving. EOLib is designed to be integrated mainly with the Data and Information Management System (DIMS) [9].

In the baseline architecture presented in Fig. 1, the PGS components are shown in blue and the data items in green. All DIMS components can be managed through a single interface, the *Operating Tool*. The *Processing* component coordinates the repeated operation of the underlying EOLib data model generation on batches of EO data. The *Long-Term Archive* (LTA) provides persistent storage and access to EOLib data items. The LTA is based on the DIMS Product Library (PL) and is the master repository for EOLib. An additional high-speed database is required for performance reasons as it shall be possible to rebuild the database from LTA data. Data from the processing workflow is uploaded via the *Ingestion Interface* to the *Data Mining Database* (DMDB) and EOLib data items are transferred and stored in the Long-Term Archive together with the standard L1b products. A simple mechanism that ensures data consistency (LTA against DMDB) is provided. The *Online User Services*, based on the PGS librarians' Earth Observation on the Web (EOWEB), EOWEB GeoPortal (EGP) and Geospatial Data Access System (GDAS), is enhanced by additional EOLib components for browsing a specific part of the DMDB like the semantic catalogue and by offering queries based on metadata and semantic annotations.

2.2. EOLib System Components

The new EOLib system components contain most of the novel and innovative EOLib functionality. They are shown in the baseline architecture in orange and the data items in yellow. Here it can be seen that DMG and DMDB are internal components managed by a tailored operator while the QE, KDD, and VDM components provide front-end functionality to the user and operator. The EOLib system modules are either independent components or functionality integrated into existing PGS components.

Fig. 1 shows that the *Data Model Generation* (DMG) is controlled by the *Processing* component. The data model generation component is a processing chain that produces EOLib data items from a TSX product and its metadata. The main DMG functionality is metadata extraction, image tiling with multiple resolutions, basic feature extraction, and high resolution quick-look generation. The output of DMG consists of metadata, tiles, high resolution quick-looks, and extracted features. They are saved into the LTA. These data are later transferred to the data mining database via the *Ingestion Interface*. Currently, DMG counts on three feature extraction methods, namely, Gabor Linear Moments (GLM) [10], Gabor Logarithmic Cumulants (GLC) [11], and Weber Local Descriptors (WLD) [12]. These methods are tile-based and extract mainly texture as image descriptors. The *Data Mining Database* (DMDB) provides high-speed storage and some data mining functionality whose processing and retrieval performance requires a database-close implementation. It is based on the relational database MonetDB [13]. The DMDB component manages data handling, storage, administration and some of the processing for the entire EOLib components.

The *Query Engine* (QE) allows the user to search for EOLib relevant data. The following types of queries are supported: 1) Querying based on metadata, where the user can query the image archive using standard metadata as, for example, coordinate systems, type or product, acquisition time, etc. 2) Querying based on semantic annotations; here the user can select a semantic label from the available labels in the semantic catalogue to perform the query. It is worth to mention that those labels are pre-defined labels previously obtained as results of the semantic annotation by using the Knowledge Discovery tool. The *Knowledge Discovery in Databases* (KDD) component adds semantic annotations to EO products. It consists of a Graphical User Interface (GUI) which interacts with the user and receives the user's input; it is based on relevance feedback methods and the KDD core which accepts the user's requirements from the GUI and uses them to get the data from the DMDB. The KDD core component includes a Support Vector Machine (SVM) as its machine learning method in order to classify the image content and to define semantic labels. In the semantic definition step, the operator's feedback is passed as training data to the machine learning method and it performs the prediction of the results [14]. After the training, the whole data can be annotated. The *Visual Data Mining* (VDM) component is a structured browsing facility for large amounts of image data. It is composed of a GUI which allows the user to navigate in the image archive. It gets the data from the DMDB and adapts them by dimensionality reduction to present them to the user through the GUI [15]. The *Epitome Generation* (EG) is integrated with the DMDB. The epitome is a summary of the EO product information content (i.e. metadata, high resolution quick-looks, basic features and all the annotations) as presented as actionable information for information mining in individual EO products. The epitome is a result of the data model generation and semantic annotation and may be delivered with the standard EO product or as a distinct product component. It is intended to be used for the individual product content inspection. Epitome access is done offline, on the user's PC/notebook, using the Epitome Browser (a separate program).

3. EXPERIMENTAL RESULTS

The DMG of the EOLib system starts the ingestion of the TerraSAR-X product. It sets the input parameters (i.e., product path, patch size, levels of resolution, etc.). Later, during the metadata extraction, the XML annotation file is read and it extracts the relevant metadata entries as, for example, the four corner coordinates, acquisition angles, resolution, pixel spacing, number of bands, acquisition time, etc. Further, the TSX image is tiled into several patches generating a grid of multi-size patches together with their high resolution quick-looks. Later, the primitive features are extracted from each generated patch by the selected methods. Finally, all the generated information is written into an XML file called the data model and it is transferred to the long term archive. In a further step, the ingestion interface uploads the generated information into the DMDB,

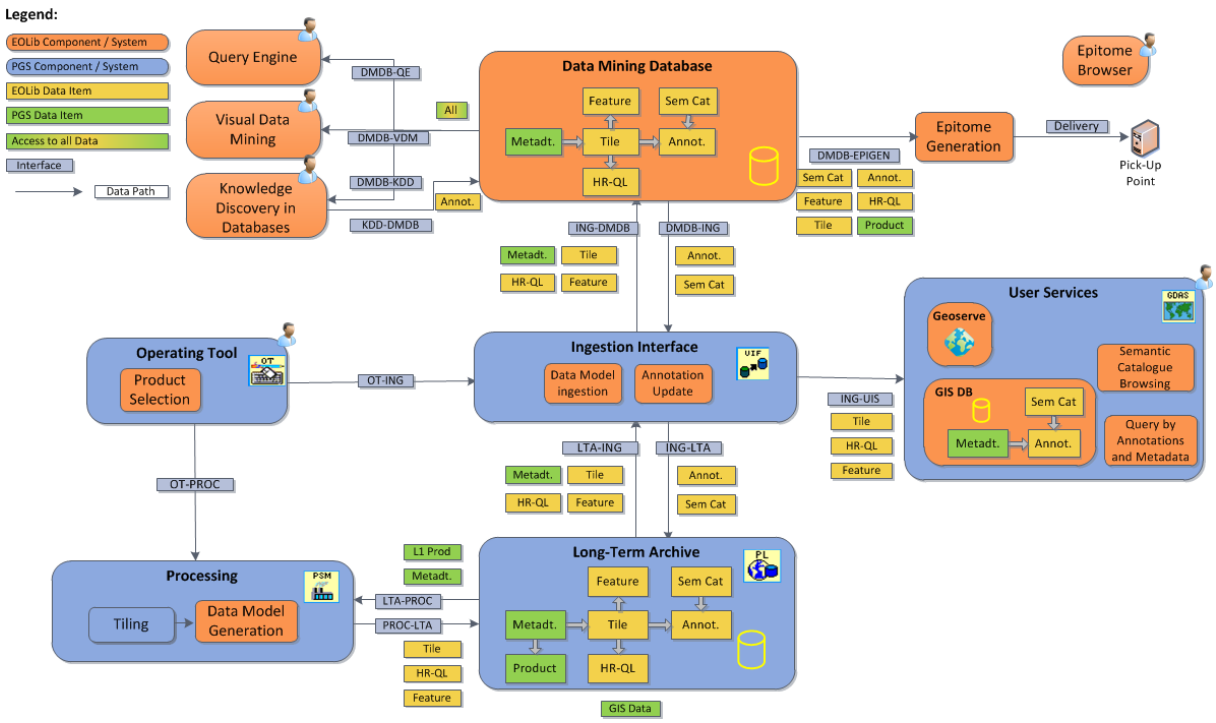

Fig. 1: EOLib Baseline Architecture. Systems/components are shown as rectangles with rounded corners (blue for the current PGS and orange for new EOLib components). Interfaces are displayed as arrows pointing to the direction of the information flow. Components with user interfaces have an icon user shown in the upper right corner. Information items transmitted via an interface are displayed near the interface arrow as rectangles (yellow for EOLib and green for PGS).

thus enabling the remainder of the components (i.e., KDD, VDM, QE). The generation of the data model using a TerraSAR-X scene of 8000×8000 pixels takes approximately less than three minutes, which is a reasonable computing time in the big data era. The use of metadata enriches the data model by adding more parameters that can be later used in advanced queries. The data model will be completed by adding semantic annotations of the image content provided by active learning methods. The next step is to provide semantic definitions to the image content. This function is performed by using the KDD component, which allows us to load one or several scenes and to annotate the image content. These annotations are stored back into the DMDB and will be transferred later to the user services via the ingestion interface. The US will allow queries based on metadata and semantic descriptions via a graphical web user interface.

The EOLib system has been processing mainly TerraSAR-X scenes. Currently, the DMDB comprises about 1200 scenes taken from around the world. These scenes were tiled on two grid levels with sizes of 256×256 and 128×128 pixels resulting in approximately ten million tiles. From each tile its primitive features were extracted using Gabor filters (GLM and GLC) with 4 scales and 6 orientations as input parameters, and a Weber local descriptor with 18 excitation levels, and 8 orientations as input parameters. Both Gabor methods yield a primitive feature vector of 48 dimensions, and a vector with 144 dimensions in the case of Weber descriptors. As examples of semantic definitions we annotated several scenes with land use and land cover categories. The categories were taken from the EO Taxonomy presented in [16], which is hierarchically organized in two levels. The first level includes 8 main land cover land use classes like like urban areas, transport, industrial areas, agri-

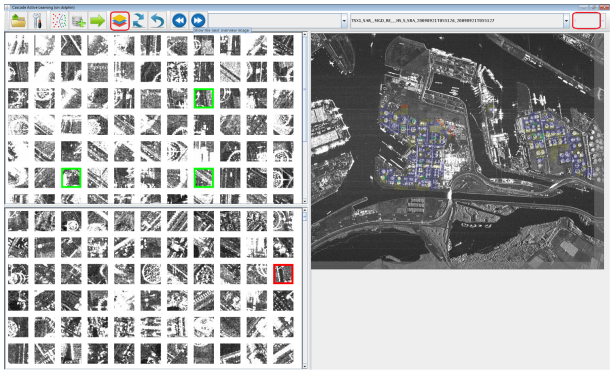

Fig. 2: Example of the Knowledge Discovery GUI component. The tiles marked in blue represent oil containers that were found using machine leaning methods implemented in this component.

culture, military facilities, bare ground, water bodies, and natural vegetation, while the second level contains about 10 subcategories for each main category. In urban area, for example, we can find semantic categories like high buildings, high density residential areas, informal settlements, low density residential areas, skyscrapers, etc. Fig. 2 shows an example of the GUI used for semantic definition and Fig. 3 displays the user web services interface.

In the example shown in Fig. 2, a TSX scene is displayed on the right part and a list of its tiles is shown on the left part. We can observe that the scene has categories like oil containers, water bodies,

doi: 10.2788/854791

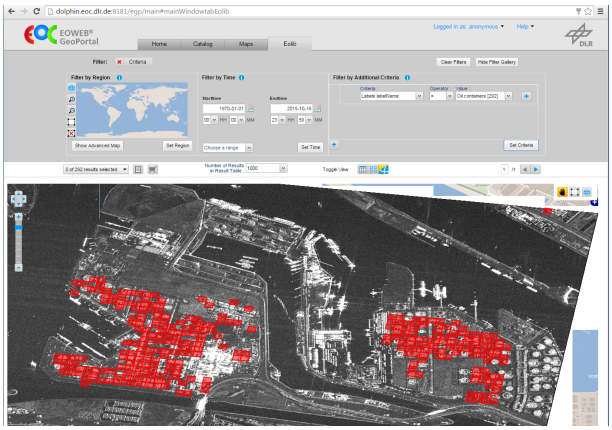

Fig. 3: Example of the User Web Services component. The tiles marked in red represent oil containers that were found using machine learning methods.

industrial areas, etc. In this example, the operator is looking for oil containers, so the tiles marked in green are the positive examples, while the negative examples are marked in red. Those examples are passed to the SVM. The SVM performs the prediction of the desired label and returns the classification result, which is presented in blue. When the classification is satisfactory, the annotation of the tiles with semantic categories is stored in the DMDB and further transmitted to the LTA and then this information will be available for searches in the US. The US interface is shown in Fig. 3; here it can be seen that a search for "oil containers" has been made. The results are highlighted in red.

4. CONCLUSIONS

In this paper we introduced the Earth Observation Image Librarian (EOLib), the data mining component of TerraSAR-X being installed in the TerraSAR-X Payload Ground Segment. The implemented system is installed in the TerraSAR-X Payload Ground Segment. The architecture is presented as a modular system integrating several components with well-defined functionality allowing the improvement or replacement of each component without affecting the whole system. It provides user services such as searches based on different parameters as, for example, metadata and semantic annotations of the image content. The ingestion of a new scene takes on average three minutes.

5. ACKNOWLEDGMENT

We would like to thank ESA for funding the EOLib project.

6. REFERENCES

[1] O. Dumitru and M. Datcu, "Information content of very high resolution SAR images: Study of feature extraction and imaging parameters," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 51, no. 8, pp. 4591–4610, Aug. 2013.

[2] M. Datcu, H. Daschiel, A. Pelizzari, M. Quartulli, A. Galoppo, A. Colapicchioni, M. Pastori, K. Seidel, P.G. Marchetti, and S. D'Elia, "Information mining in remote sensing image archives: system concepts," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, no. 12, pp. 2923 – 2936, Dec. 2003.

[3] C.-R. Shyu, M. Klaric, G.J. Scott, A.S Barb, C.H. Davis, and K. Palaniappan, "GeoIRIS: Geospatial Information Retrieval and Indexing System-Content Mining, Semantics Modeling, and Complex Queries," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 4, pp. 839–852, 2007.

[4] A.W.M. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain, "Content-based image retrieval at the end of the early years," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 12, pp. 1349–1380, Dec. 2000.

[5] N. Rasiwasia, P.J. Moreno, and N. Vasconcelos, "Bridging the Gap: Query by Semantic Example," *IEEE Transactions on Multimedia*, vol. 9, no. 5, pp. 923 –938, Aug. 2007.

[6] D. Espinoza-Molina and M. Datcu, "Earth-Observation Image Retrieval Based on Content, Semantics, and Metadata," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 51, no. 11, pp. 5145–5159, Nov. 2013.

[7] A. Bifet, "Mining big data in real time," *Informatica (Slovenia)*, vol. 37, no. 1, pp. 15–20, 2013.

[8] V. Manilici, S. Kiemle, C. Reck, and M. Winkler, "Architecture concept for an information mining system for earth observation data," in *Proc. PV*, 2013.

[9] M. Wolfmüller, D. Dietrich, E. Sireteanu, S. Kiemle, E. Mikusch, and M. Böttcher, "Data Flow and Workflow Organization - The Data Management for the TerraSAR-X Payload Ground Segment," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 1, pp. 44 –50, Jan. 2008.

[10] B. S. Manjunath and W. Y. Ma, "Texture features for browsing and retrieval of image data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 8, pp. 837–842, Aug. 1996.

[11] J. Singh and M. Datcu, "SAR Image Categorization With Log Cumulants of the Fractional Fourier Transform Coefficients," *IEEE Transaction on Geoscience and Remote Sensing*, vol. 51, no. 12, pp. 5273–5282, 2013.

[12] S. Cui, C.O. Dumitru, and M. Datcu, "Ratio-Detector-Based Feature Extraction for Very High Resolution SAR Image Patch Indexing," *IEEE Geoscience and Remote Sensing Letters*, vol. 10, no. 5, pp. 1175–1179, Sep. 2013.

[13] CWI, *Monetdb Copyright (c)*, Mar 2008-2011, Available from: http://www.monetdb.org/Home.

[14] P. Blanchart, M. Ferecatu, S. Cui, and M. Datcu, "Pattern retrieval in large image databases using multiscale coarse-to-fine cascaded active learning," *Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of*, vol. 7, no. 4, pp. 1127–1141, April 2014.

[15] D. Espinoza-Molina, M. Datcu, D. Teleaga, and C. Balint, "Application of visual data mining for earth observation use cases," in *Proc. ESA-EUSC-JRC 2014 - 9th Conference on Image Information Mining Conference: The Sentinels Era*, 2014, pp. 111–114.

[16] C. O. Dumitru, S. Cui, G. Schwarz, and M Datcu, "A Taxonomy for High Resolution SAR Images," in *Proc. ESA-EUSC-JRC 2014 - 9th Conference on Image Information Mining Conference: The Sentinels Era*, 2014, pp. 89–92.

SENSYF EXPERIENCE ON INTEGRATION OF EO SERVICES IN A GENERIC, CLOUD-BASED EO EXPLOITATION PLATFORM

Nuno Almeida[1], Nuno Catarino[1], Antonio Gutiérrez[1], Nuno Grosso[1], João Andrade[1], Hervé Caumont[2], Pedro Gonçalves[2], Guillermo Villa[3], Antoine Mangin[4], Romain Serra[4], Harald Johnsen[5], Tom Grydeland[5], Stephen Emsley[6], Eduardo Jauch[7], Jose Moreno[8], Antonio Ruiz[8]

[1] DEIMOS Engenharia, Portugal; [2] Terradue, Italy; [3] IGN-CNIG, Spain; [4] ACRI-ST, France; [5] Norut, Norway; [6] ARGANS, UK; [7] IST, Portugal; [8] UVEG, Spain

ABSTRACT

SenSyF (Sentinels Synergy Framework) is a cloud-based data processing framework for EO-based services. It has been pioneer in addressing Big Data issues from the Earth Observation point of view, and is a precursor of several of the technologies and methodologies that will be deployed in ESA's Thematic Exploitation Platforms and other related systems.

The SenSyF system focuses on developing fully automated data management, together with access to a processing and exploitation framework, including Earth Observation specific tools. SenSyF is both a development and validation platform for data intensive applications using Earth Observation data. With SenSyF, scientific, institutional or commercial institutions developing EO-based applications and services can take advantage of distributed computational and storage resources, tailored for applications dependent on big Earth Observation data, and without resorting to deep infrastructure and technological investments.

This paper describes the integration process and the experience gathered from different EO Service providers during the project.

Index Terms **–** Copernicus, Synergy, Cloud, EO Services, Integration

1. INTRODUCTION

In addition to the SenSyF Cloud Platform, leveraging the HADOOP Streaming programming model, SenSyF links to the Sentinels Scientific Data Hub and will be connected to the future ESA ngEO catalogue. For the initial integration phase of any incoming Service, a Cloud Sandbox model furnishes an environment very similar to the space agencies operational environments, where the services are proof-tested against large EO dataset series. Built upon the processing infrastructure, a complete set of tools (SenSyF SDK - Service Development Kit) for data pre-processing is included, comprising the basic operations of data acquisition and dissemination, retrieval, atomic manipulation tools (re-gridding, image composition, etc.) as well as high level common processing algorithms (e.g. vegetation indexes calculation). This common tools are provided "ready to use" as part of the Cloud Sandbox.

Throughout the last three years, the SenSyF team has tested the concept on seven beta testing services, covering a wide spectrum of use cases and application areas within EO domains. The services currently deployed range from the monitoring water quality and quantity, and monitoring of agriculture irrigation processes, to mapping the growing season in arctic and alpine areas.

During this period, the development and integration of the services in the platform has provided a considerable amount of valuable information and insight on the advantages and disadvantages of a generic exploitation platform. A large effort is being devoted to assessing this issues and identifying future development paths, including the need to mitigate the technology and knowledge gaps. A detailed analysis is performed on the applicability and benefits, to the service developers, of the adoption of an EO-focused cloud-based platform.

This assessment and lessons learned by the SenSyF consortium will provide valuable inputs for future exploitation platforms, designed to generate Value Added Services using Earth Observation data.

2. MAIN ACTORS

The development, deployment and exploration of a service on SenSyF involve three different actors. They are the developers, the operators and the end users. Each one has specific tasks during the process.

The service developers are the users that design and implement services using EO Data. They need a proper development environment (tools, libraries, etc.).

The service operators/administrators are responsible for the operations and maintenance. This kind of user needs a proper interface to monitor and control the service to be able to fulfill its SLA.

The end users are the ones that receive the output of the services (scientists, agricultures, research institutes).

doi: 10.2788/854791

3. DEVELOP & DEPLOYMENT STEPS

The development and deployment phases are entirely performed by the service developers. In fact, there is a sequence of steps to develop, test and deploy an EO application. In a general way, before the deployment on a public cloud, the developer needs to design the service, develop its jobs, define the workflow and validate the application on the development environment. The next sections will explain further each one of these steps and the main difficulties faced by the service developers.

4. DEVELOPMENT ENVIRONMENT

The development environment is basically a sandbox service. A remote Virtual Machine (VM) which allows the developers to deploy their scientific algorithms, develop the service workflow (Hadoop Map/Reduce system [6]) and also test the service before the deployment on the cloud environment, as shown in figure 1.

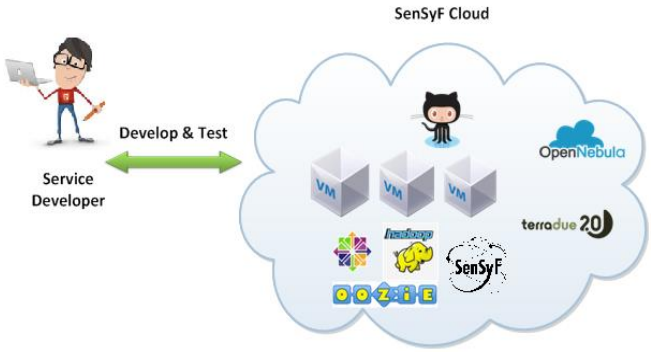

Figure 1. Representation of Development Environment

This VM represents a single node of a bigger cluster and contains two CPUs to allow the parallelization tests.

The development environment, in the first version of the framework, was perceived as somewhat confusing to the service developers who reported a considerable learning curve since the first login until they felt comfortable deploying their applications there. On the final version, thanks to the feedback of these early adopters, there is much more documentation explaining the structure, the file systems and how to easily develop in this sandbox environment.

5. SERVICE DEVELOPMENT KIT

The sandbox environment includes a Service Development Kit (SDK) which is a set of tools and third party libraries that can be integrated by the developers, on their algorithms and scripts (developed in different programming languages), to access and process Data and disseminate results.

The Data Access tools available are used to search and get the catalogue Data. These operations are performed thanks to a library called CIOP [10] (Cloud Interoperability Operational Pilot) that contains commands to copy Data between different file systems and perform queries on the catalogue.

Pre-processing and processing operations over the Data (tile, orthorectification, crop, reprojection, merge, etc.) can be performed with the SenSyF Tools and the Third Party libraries (GDAL [7], BEAM [9], Sentinel Toolboxes [1][2][3], etc.).

The dissemination of the results can be done in many different ways (email notifications, reports, webportal display). Thanks to some SenSyF Tools developed with this goal the end user will receive the outputs of the application.

In order to have some feedback about the SDK Tools available and about how easy was to use and integrate them on the algorithms and scripts, several surveys were distributed by the service developers. An analysis to these surveys allows concluding that the use of the SDK Tools was a simple process with a small learning curve. Almost all of the demonstrative services used the SDK tools and the developers found them easy to install and integrate on their workflows. Since all the tools are available through a repository on the sandbox, they can be automatically installed if the developer writes a simple dependencies list.

To clarify the usage of each tool, an SDK space was created containing manual pages and tutorials for all the tools. The service developers also confirmed the utility of this informative space and all the tutorials available.

6. WORKFLOW DESIGN

In SenSyF, the EO Service Developer has first to design the service as a sequence of jobs (workflow).

Figure 2. Sequence of jobs (workflow)

Each job of the workflow can run in parallel or be a single instance. The parallelization is done job by job by the inputs and managed automatically by Hadoop [6].

This was a concept that service developers took some time to assimilate. In some of the cases, the main difficulty faced by the developers was to reach the best DAG (directed acyclic graph) for their applications and represent their services through a workflow (figure 2).

In the Hadoop system [6] the workflows are orchestrated with an application descriptor file (XML file describing the jobs, inputs, workflow sequence). At the beginning, this was a constraint for the service developers, who were not interested in creating/editing XML files by themselves. To solve this issue and avoid many hours to

doi: 10.2788/854791

create a simple workflow, a Graphical User Interface (GUI) was created – a SenSyF Workflow Designer.

The SenSyF Workflow Designer came to simplify the process of creating a workflow. With this tool, the service developer is able to design and orchestrate a workflow in a user friendly way. The GUI also drives the developer through some steps which simplify the entire process. The service developer only has to enter some inputs on the GUI (path to the binaries, query parameters, etc.) and create the workflow dragging and dropping boxes. Figure 3 illustrates the main window of the GUI.

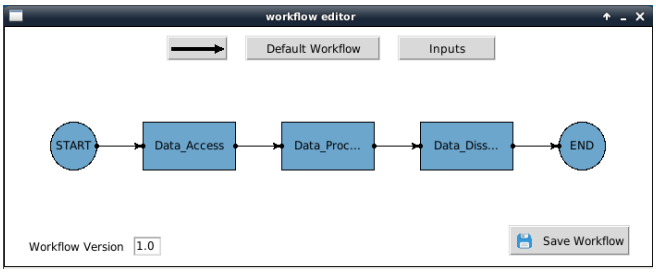

Figure 3. Workflow Manager GUI

7. SERVICE VALIDATION

After the development phase, the most important step is to test the service in the sandbox environment before deploying it in a big cluster on a public cloud. This phase was easy for the service developers since, from the first version of the framework, they had access to specific testing tools targeting the workflow and parts of it (single jobs). During the development phase, the developers found it was a good practice to perform tests constantly. This procedure simplified the service validation.

The final validation can also be performed through a user friendly dashboard where the tester only inserts the inputs (time, coordinates, etc.) and clicks a "Run button".

During the test run, the developer is also able to follow the workflow in a debug area and easily locate the errors in case they occur.

8. FROM THE SANDBOX TO THE CLUSTER

Once the service is validated, the final step involves packaging the application and deploying it on a public cloud (Amazon, Interoute, etc.). To package the application, we selected the GitHub concept since it is a strong tool to save the code and manage different releases of the workflow. After the packaging, the user gets an RPM file containing the application structure, files and dependencies (tools required). Then, it is very straightforward to install the application on a clean sandbox or in the master node of a big cluster, in order to be replicated on the slave nodes.

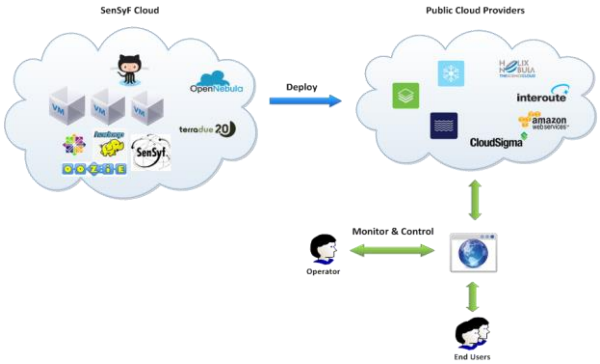

Figure 4. Cloud Deployment Schema

In figure 4, it is also possible to check the interaction of the operator through a dashboard or a webportal. The operator can start, stop and monitor the service. Then the end user will receive the results (webportal, email, report, etc.).

Before the deployment of the service on a public cloud, it is necessary to open a cluster, with the intended number of nodes. To do this, the developer has a cloud dashboard to select the number of virtual machines needed. It is important to note that each VM (node) will be a copy of the sandbox environment that allows running the workflow in parallel. The developers are also able to monitor the VMs on the dashboard. In figure 5, it can be seen a cluster with a master node and fifty slave nodes.

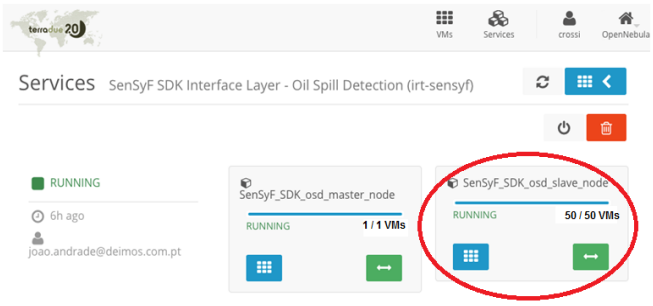

Figure 5. Cloud Monitoring Dashboard

9. SENSYF DEMONSTRATION USE CASES

In order to prove the SenSyF concept and test the development, integration and deployment of applications on the SenSyF Framework, demonstration services were challenged to test the platform (development environment).

Service developers were asked to bring their applications to the SenSyF Framework and take profit from the Big Data access and the parallelization power.

The main idea was to let the developers of the services face and overcome the obstacles during the integration process and, at the same time, give their feedback about how

doi: 10.2788/854791

to improve the framework. This way, seven services, with different study areas, tried to develop their algorithms using the tools available on the SDK, tested the service on the development environment and, finally, almost all of them were able to run their workflows on bigger clusters in the public cloud. In fact, those Services, related with water monitoring, cold areas growing season and soil freezing, agriculture, land cover classification, change detection and cloud removal, were able to integrate themselves on the Framework. However, at the end, some of them concluded that the kind of application they managed doesn't need a processing power that big to compensate the hours lost during the integration at the SenSyF platform. This is a very important conclusion. Services that cover a small area of interest or use small rates of EO Data do not need the cloud to process the inputs and get results. They can do it locally.

To better illustrate the described integration process, test the framework, prove the concept and, at the same time, drive the other developers during their roadway, a pilot use case was developed and deployed by the SDK Team (Deimos Engenharia). The implementation of this use case was a very enriching process since it allowed facing the difficulties together with the demonstrative services, improve the tools and compile a set of lessons learned that make the platform more powerful and user friendly.

10. USERS CONSULTATION & FEEDBACK

The development of the SenSyF Framework had two different phases. In the end of both versions, in order to understand how the developers were dealing with their integration, consultations were performed. This users' consultation process allowed us to take conclusions about the strongest and weakest points of SenSyF and understand what should be improved to attract more developers. Conclusions related with the scalability, performance and integration were also taken from the services' feedback.

The services deployed on SenSyF successfully proved the scalability and, concerning performance, it is strongly related with the service workflow design. This way, the workflow design shall be deeply analysed by the developers in order to choose the best option to their applications. This step and the integration of algorithms into the sandbox environment were exactly the phases where the developers faced the main difficulties. They all admitted that the learning curve was one of the weakest points, specially at the first version of the SenSyF that they considered very immature and demanding despite the continuously support given by the Framework and SDK teams. The concept and the technological barrier to adapt the applications were the main obstacle found by everyone, once many service developers are not programming specialists. However, thanks to the improvements implemented (GUIs, dashboards, new documentation, support site) the time needed to be familiar with the development environment was reduced significantly.

From the developers' feedback about the second version of the platform, it is possible to conclude that, during the project life, the integration was becoming faster and simpler and, once the developers are familiar with the procedure and the concept, they can easily deploy many applications and take profit of the cloud processing and the Big EO Data access

In general, almost all the developers, operators and end users considered a good experience and a challenging exercise to bring their applications to the SenSyF Cloud Platform, and await the complete release of Sentinel data to fully exploit their services.

11. REFERENCES

[1] https://sentinel.esa.int/web/sentinel/toolboxes/sentinel-1

[2] https://sentinel.esa.int/web/sentinel/toolboxes/sentinel-2

[3] https://sentinel.esa.int/web/sentinel/toolboxes/sentinel-3

[4] http://www.eo-sensyf.com, SenSyF new website

[5] http://www.sensyf.eu, SenSyF project website

[6] http://hadoop.apache.org/

[7] http://www.gdal.org/

[8] https://earth.esa.int/web/nest/home/

[9] http://www.brockmann-consult.de/cms/web/beam/

[10] http://wiki.services.eoportal.org/tiki-index.php?page=CIOP

[11] Poster "SenSyF: Processing Framework for Sentinel Data", SENTINEL-2 for Science Workshop, Frascati, Italy, 20-22 May 2014,

[12] Poster "SenSyF: Processing Framework for Sentinel Data", Conference on Big Data from Space (BiDS'14), Frascati, Italy, 12-14 November 2014

[13] Paper "SenSyF: Processing Framework for Sentinel Data", Conference on Big Data from Space (BiDS'14), Frascati, Italy, 12-14 November 2014

[14] Poster "SENtinel Synergy Framework", 9th GEO European Projects Workshop 2015, Copenhagen, Denmark, 15-16 June 2015

[15] Promotional Poster "Big Data Solution for EU Copernicus", EO Open Science 2.0, Frascati, Italy, 12-14 October 2015

AUTOMATED SERVICE BUILDER FOR SEMANTIC SERVICE ORIENTED ARCHITECTURES – ASB

Bernard Valentin, Matthieu Melcot, Leslie Gale

Space Applications Services

ABSTRACT

We propose a novel approach for the execution of EO Products Processing Chains. The processors are deployed into a Cloud Environment that brings scalability and flexibility. In addition, each processor is interfaced through OGC Web Processing Services (WPS), leading to a modular multi-mission processing platform.

Index Terms— IPF, Processing, Cluster, Container, Cloud, OGC WPS, Semantics

1. INTRODUCTION

The framework used for the delivery of EO Services has evolved from Server-based to Grid-based processing environments (e.g. G-POD[1]), and onward to Cloud Computing. ESA at ESRIN within the context of several projects has been investigating and applying OGC standards, and more specifically the Web Processing Service (WPS[2]) Interface Standard that provides rules for standardizing inputs and output for geospatial processing services. These standards are adopted in various domains such as Industry, Government and Academic, but still underused in EO Services.

ASB will establish a generic concept and apply it to Instrument Processing Facilities (IPF's). IPF's are responsible for transforming instrument raw data into higher-level products by means of processors forming end-to-end Processing Chains. Each processor applies algorithms to correct, transform, merge data products or extract specific features from them. General opinion is that IPF's are very mission dependent and not suited to an automated build process. ASB will demonstrate current capabilities and make recommendations for further automation.

In the last five years, the enthusiasm of the IT sector for Big Data fostered the development of Cluster-based computing platforms such as Hadoop and Spark, and Cluster Management frameworks such as Mesos. Additionally, Cloud technology led major EO actors such as NASA JPL[3] to experiment the distribution of image processing into the Cloud.

However, there are still many operational IPFs and similar implementations of Processing Chains used at prototyping level that were designed and deployed prior to the Cloud era that do not take advantage of the distribution of the processing power on a Cluster. They are typically deployed into static architectures[4], where updates are tedious, and scalability is hardly possible. ASB[5] adopts the latest technologies promoting the work of ESRIN to overcome some of the persistent problems with existing IPFs by proposing a scalable and dynamically re-configurable platform for the execution of IPF's, or of any resources demanding processing facilities. Also, and for the intended users of ASB of great importance, ASB will focus on the (automated) building of workflows from user specifications making the deployment of the processors in the cloud transparent to users, but under their control.

2. OBJECTIVES

The main goal of the platform is to provide a dynamic, scalable multi-mission (automated) processing environment. To that end, the following objectives are addressed:

Scalability – EO products are huge and processing them requires great amounts of computing power, network and storage. The infrastructure shall be able to scale and adapt to the users needs in term of products generation and get rid of the classical infrastructure limitations.

Automated Generation of Workflows – Although endeavors have been made to standardize the interfaces among the processors and the processes — see the Generic IPF Interface Specifications[6] — disparities remain between processor interfaces of different missions. To automate the generation of workflows — that is, the organization of processes with defined inputs and outputs, implementing a processor — a generic interface for all the processors is needed to provide modularity and make it possible to share processes between different processing chains and missions.

Generic Orchestration – The orchestration of the processes is closely related to their related missions. A generic mechanism, adaptable to the specific orchestration baselines, is essential to run workflows from different missions into the same platform.

IPF Evolution – Traditional IPF's are static and upgrading the workflow definitions or the processes are not easy tasks. A simplified means to deploy new (versions of) processes and edit the workflows accordingly is therefore essential.

Monitoring and Control – Today's IPF's do not provide fine-grained control over the workflows execution. The solution shall provide advanced monitoring and control capabilities to overcome this limitation.

3. CONCEPTS

Cloud Environments – A scalable infrastructure where new Virtual Machines (VM's) can be deployed, started and stopped, automatically or on-demand to cope with changing computing needs. A Cloud environment also provides an elastic storage capacity which allows adapting the volumes seamlessly during the operations.

To benefit fully from parallel programming models such as MapReduce with Hadoop, the processors must be designed taking this paradigm into account. This is obviously not the case for the vast majority of the processors used today, including the ones running in the IPFs. Moreover, it is not necessary to adopt this way of developing and operating processors when input parameters are already fragmented in units of manageable size. This is in particular the case of EO products that are traditionally packed in tiles with limited surface coverage. To cover a bigger area, the contained and the intersecting tiles can be processed separately and the individual results merged to generate the final product.

In its initial version, ASB will put the focus on the dynamic deployment, the load balancing, and the orchestration of atomic processes. In order to anticipate evolution in the technologies and platform offerings, attention will be paid to integrate abstraction layers, allowing, e.g. switching seamlessly between cloud providers or adopting alternate programming models. These capabilities are made possible by OpenSource tools such as Apache Mesos (resource management and scheduling), Apache Zookeeper (cluster management) and Apache LibCloud (unified API for public and private clouds, including OpenStack).

Flexible Orchestrator – A fully customizable workflow engine provides the necessary flexibility to adapt to the multiple workflow patterns used in the IPFs, such as Parallel Split, Synchronization, and Exclusive Choice. When executed, workflows are transformed from a static to a dynamic representation in order to gain the ability to reconfigure themselves. This allows conditional branching and iterative operations needed in Products such as Calibration Products.

Data Types Semantics – Care must be taken that the processes chained in workflows are compatible and in particular that outputs of a given process may be used as inputs for the subsequent process. In a platform that allows the editing workflow definitions, it is highly recommended to include a mechanism that verifies the consistency of the dataflow between processes. An ontology containing the data types available is used to provide clear inputs and outputs signatures.

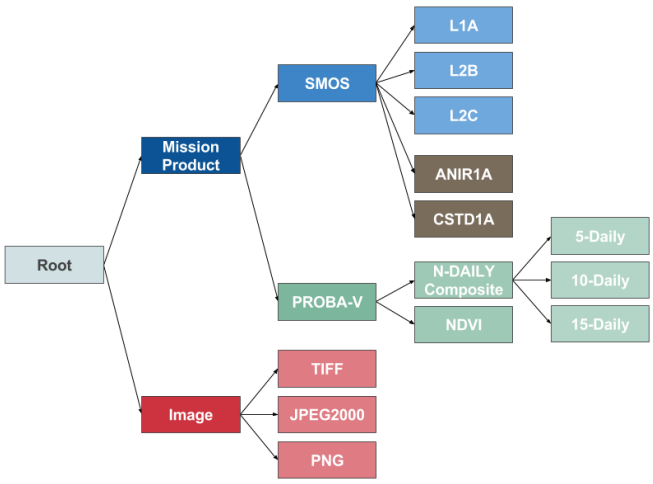

Fig. 1. Data Types Ontology (Fragment)

The data types ontology is used to verify the compatibility of input and output process parameters. This mechanism must have the knowledge of the broader, narrower and equivalent data types. For example, JPG and PNG are both Image types. A JPG output may be provided to a process that accepts an Image as input but not the other way around. The ontology can be enriched at any time with new data types, still paying attention that the modifications do not invalidate existing workflow definitions.

Processes-Data Locality Paradigm – Processes size is usually more than one order of magnitude smaller than the Data size. Furthermore the same Process is often applied on a large list of collocated data. The classical approach of processing remote data by first fetching the data is cumbersome.

A more convenient strategy leads to bringing the processes close to the data to be processed. This method, facilitated by the use of a Cloud Environment, highly reduces the data transfers that are not only time but also money (in the case of public clouds) consuming. Also, placement constraints on the processes deployment enforces the minimization of the data transfers between processes, especially in the case of use of huge intermediate or auxiliary data, such as the 8 Gigabytes SMOS L1B G-Matrix. The dynamic placement of processes is enabled by the Linux Containers technology.

Editable IPF's – A Workflow Editor provides the operators with a graphical means to create and edit workflows from a list of available tasks (the internal representation of an executable process). This is the first step towards an automated build and will greatly simply the effort required by the user in the prototyping and deployment of new IFPs.

doi: 10.2788/854791

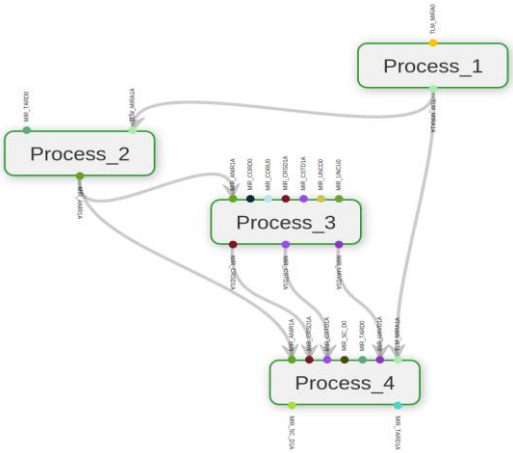

Fig. 2. Workflow Editor

OGC WPS – This gives the ability to interact with a specific process through an ad-hoc interface inside a standard Web Service. Processes are executed through a standard HTTP request, where URL's representing the input data files are provided as inputs. Process outputs may also be fetched through URL's.

4. ARCHITECTURE

4.1. Static Architecture

The architecture of ASB relies heavily on a collaboration between web-services where each one assumes clear responsibilities.

Figure 3 depicts the architecture of ASB. Operators and end-users interact with the ASB Core Components via their user interfaces. These components may be co-located or distributed. Their resource needs are relatively low. The lower part of the figure represents the processing components, meant to be deployed in a Cloud Environment.

The core services communicate with each other through the REST interface they are exposing. This brings the possibility to share the resources between the services and to update specific components individually without impacting the rest of the platform.

Product generation requests are received by the Service Builder core component. The Service Builder obtains workflows, processes and data types definitions from the Knowledge Base and constructs a specific workflow execution order. The order is issued to the Orchestrator component that is in charge of organizing and managing the execution of the processes via the Task Manager.

The Distributed Task Queue ensures the distribution of the process invocations over the Cloud where processes stemming from different workflows are wisely distributed according to the availability of the nodes and the current location of the input data.

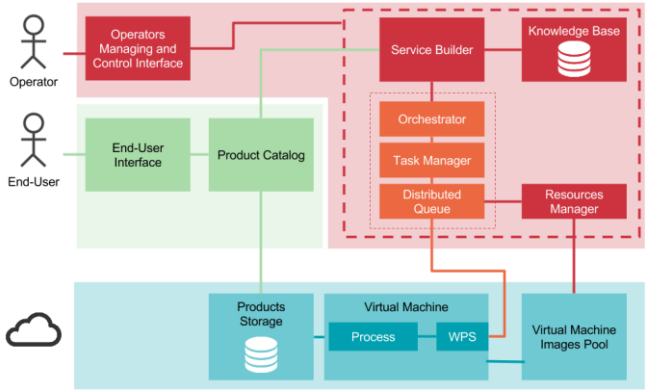

Fig. 3. ASB Architecture

The Resources Manager deploys the required processes from a pool of Virtual Machines (VM's) and container images. The processes sit behind WPS interfaces allowing them to be addressed in a seamless manner.

ASB Core Components are implemented in Python. The Product Catalog stores its data in a PostgreSQL/ PostGIS database. The Service Builder, the Orchestrator, the Task Manager, and the Distributed Queue are asynchronous components that use polling mechanisms. Containers are used to replicate process specific environments in the VM's, with the benefit of sandboxing the resources.

4.2. Dynamic Architecture

The open-source cluster manager Mesos provides the scalability of the platform. Mesos allows abstracting the computing resources and making them available to the applications built on top of it.

This means that the same cluster can be used to execute the processes from the ASB workflow, but also those of an Hadoop (or Spark) environment in cases some ASB processes need it internally. The balancing of the resource allocation between ASB and the Hadoop (or Spark) environment is configurable at the cluster manager level.

5. BEHAVIOUR

There are two key ASB mechanisms. The first one is used to configure processes and workflows platform. The second mechanism takes place at the time a product generation request is issued.

5.1. Workflow Editing

The constitution of a new workflow follows the steps presented in figure 4.

1. The operator provides the URL and name of a new OGC WPS Process. ASB automatically fetches the process description and generates an initial version of the process definition.

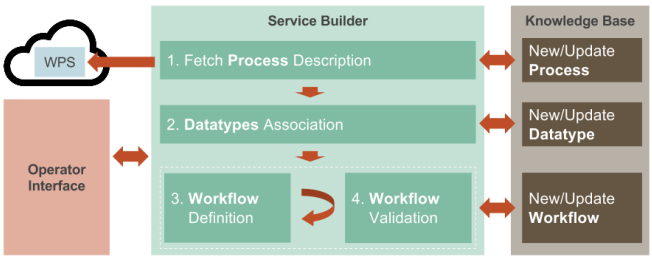

Fig. 4. Workflow Editing Steps

2. The operator augments the initial process definition, e.g. by associating the input and output parameters with data types. New data types are defined if necessary.
3. The operator graphically combines processes to define a workflow (see figure 2).
4. ASB automatically validates the associations by e.g. verifying the compatibility of the connected output and input parameters. At each step, the definitions are persisted in the Knowledge Base.

5.2. Workflow Execution

Building a Product requires the Service Builder to execute the sequence of steps depicted in figure 5.

1. Knowledge Elements, including the workflow definition and the related process definitions, are retrieved from the Knowledge Base.
2. The necessary resources are deployed in the Cloud Environment.
3. A workflow execution order is built and handed over to the Orchestrator.
4. When the requested product is ready, this is registered in the Product Catalog and a notification is issued.

5. CONCLUSIONS AND FUTURE DEVELOPMENTS

Following ESA's perspective to bring user's processors to the huge amount of exploitable EO data, ASB will provide an environment to host custom processors and allow testing and sharing of new workflows and algorithms.

ASB will implement technology and methods for distributed EO Processing Chains offering full support to users for building workflows incorporating processes and (big) distributed databases. This will significantly simplify the creation and management of processors that form an integral part of the IPF. It will increase the reliability of the operational IPF and reduce costs through standardization of the interfacing whilst allowing a transparent interface to the commercially driven processing environments currently led by Hadoop and Cloud processing technologies.

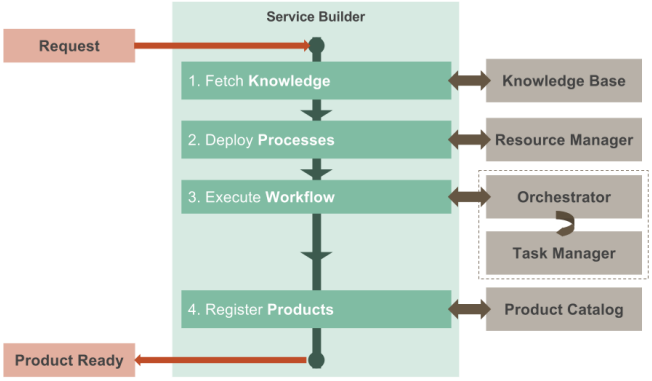

Fig. 5. Workflow Execution Steps

ASB scalability means that it can be used from the prototyping of retrieval algorithms by the scientific community, where execution could be on a local computer, through to the instantiation of a complex operations IPF by the IT community taking advantage of cloud processing capabilities.

The concepts and techniques used in ASB are not limited to the processing of EO data. In addition to enhancing the platform, future developments will also address the applicability in other domains.

The structure of the ASB platform leaves the way open to future extensions where users are closely involved in the creation and sharing of new processes and processors.

A programmatic interface such as iPython Notebook or those provided by Google Earth Engine in addition to the Workflow Graphical Editor would represent a valuable enhancement in order to meet the scientific expectations in terms of customization, sharing and reuse of the algorithms to process data.

The ASB concepts will be validated and compared using existing processors from SMOS and PROBA-V.

6. REFERENCES

[1] ESA Earth Observation G-POD home page, http://gpod.eo.esa.int/

[2] "Web processing service 1.0.0," OGC 05-007r7, OGC, June 2007.

[3] AWS Case Study: NASA/JPL's MER and CARVE Missions, https://aws.amazon.com/fr/solutions/case-studies/nasa-jpl/

[4] "Sentinel-2 Instrument Processing Facility Technical Specification," SRS S2-PDGS-TAS-DI-BPDP-CCTS-IPF, April 2012.

[5] ASB project page on ESA Research & Service Support, https://wiki.services.eoportal.org/tiki-index.php?page=ASB

[6] "Generic IPF Interface Specifications," ICD MMFI-GSEG-EOPG-TN-07-0003 , ESA, August 2009.

doi: 10.2788/854791

A DATASCOPE FOR SATELLITE DATA

*Romulo Goncalves *, Martin Kersten **, Kostis Kyzirakos ***

NLeSC* and CWI**, The Netherlands

ABSTRACT

The Copernicus program will produce a sizable stream of remote sensor data, which when combined with ground sensor systems provides a fertile soil to explain the past and predict the future of Planet Earth's health. The sheer size and data complexity calls for innovative technology to provide efficient and effective data exploration.

We present a DataScope for Satellite data, a framework geared at interactive derivation of knowledge from various remote sensing data sources. It can handle their voluminous data sizes and geo-distribution. Moreover, it simplifies the task for the data scientist who wants to adapt his search parameters to steer his scientific quest. At the same time being able to browse the data stored in the Cloud, to execute complex algorithms, and correlate it with data stored at a different location.

Index Terms— Satellite data, Data Management, Data processing, DataScope

1. INTRODUCTION

The Copernicus program uses remote sensing to monitor the health of our planet. Remote sensing provides large area coverage for regional surveys, data acquisition at different scales and resolution, repetitive coverage, monitoring dynamic themes like water, forests, agriculture, etc. It shifts traditional field work into a digital lab work setting, much like most astronomers nowadays look at computer screens then through the lens of a telescope.

Remote sensing data is more than just spectral, spatial and temporal information. Its size, complexity and diversity will increase by orders of magnitude in the coming decade as more sentinels are launched to space and combined with ground surface sensors data. Furthermore, it is designed to serve different consumers: scientists, municipalities, companies, or home users. Data exploration will occur at different resolutions with an user collecting an insight or submitting a job for a precise answer.

The pivotal technologies to achieve results are Database Management Systems (DBMSs), they are used on science work-flows on institutional compute clusters or within the public cloud. Unfortunately, most DBMSs considered for Earth Observation (EO) share three major drawbacks: static storage model, data ownership, and lack of support for step wise data exploration [1]. One might argue that the NoSQL class of databases, such as Cassandra and HBase or MapReduce frameworks [2] are less restrictive to the change of focus by the Data Scientists. The caveat, however, is that they call for strong programming experience to harvest the potential performance. The model is mostly effective for single scan, cleanup, and summarization tasks. The trend to align NoSQL and SQL systems is proven by product offerings such as BigInsights [1] and Pivotal [2].

A traditional approach is to rely on specialized libraries which are mostly file based systems, rich in functionality and created by domain experts. However, they suffer from data isolation and applications become dependent on data formats and humans software skills. They are known for their poor scalability, e.g. on multi-core computers, and data interchange between different libraries is tedious. Standardization of e.g. HDF5 leaves many options open for a domain expert to encode his information. It is this freedom that often make re-use of the data by others a daunting programming task.

A leapfrogging experience will come from fusing data from different sources into a single data exploration framework. Such a framework should be resilient over time, adopting and evolving with the data received, and stimulate research insights for a broad spectrum user community. Hence, EO data science requires a scientific instrument capable of seeing what the data management tools and specialized libraries cannot. In Science and Industrial literature such an instrument is referenced as DataScope [3, 4].

With a DataScope, raw data from our environment is extracted and organized before being massaged by a set of tools to extract information [4]. It provides fast data ingestion and a guided dynamic step wise exploration in those "heterogeneous mountains of data", i.e., the bottom layer of new data science instruments for data exploration.

In short paper we will present our DataScope for remote sensing data. It creates a true symbiosis between the best of bread open-source columnar database system, specialized libraries, and programming languages (R, Python, Java, C).

[1] http://www-01.ibm.com/software/data/infosphere/biginsights/
[2] http://blog.gopivotal.com/tag/hawq

doi: 10.2788/854791

2. A DATASCOPE ECOSYSTEM

With a DataScope raw data from our environment is extracted and organized before being massaged by a set of tools to extract knowledge.Its novelty is visible in three key-areas.

2.1. Data access

The DataScope is designed to efficiently combine heterogeneous data sets while providing *in-situ* data access, i.e., it keeps data in its original format while scalable and distributed processing functionality is offered through a DBMS. It provides transparent access to all data kept in the file repository through a tabular or array-based interface abstraction [5].

The *in-situ* data access is possible due to the large amounts of metadata (data of data) existent on file formats such as NetCDF, LAS/LAZ, MSEEDS, HDF5, etc. Such metadata is used for effective data skipping, but also to collect data insights, e.g. summaries and samples, without having to process the entire data set. It is treated as rough sets which allows us to use machine learning techniques for clustering and approximated answers [6]. Combined with statistical models and equations they are a compressed form for data representation.

Such an approach gives the user the opportunity to continue performing data curation activities since the main data archive is the file-based repository, i.e., the raw data is kept outside of the DBMS. The data imported into the DBMS is easily invalidated in case of updates. For a new data format version the catalog tables are easily updatable and a new data loader is provided.

The dynamic data loading comprises of three phases: the attachment of a file, the import of the file's content and the collection of statistics to boost query optimization. During the attachment, the file's metadata is loaded into a special DBMS catalog. At query time, such a catalog is inspected to decide whether the file has information relevant to the query or not. In such a case the file's content is imported into the database.

The data import happens in two flavors, if the file format has each attribute sequentially stored then the import memory maps each attribute as a column, otherwise, the data is converted and loaded into the database as temporary data. In the latter case, cache policies, such as Least Recently Used (LRU), are used for data eviction.

To improve data parallelism, during the import one or more relational tables are created. These tables are seen as partitions and they are glued using a *merge table* to hide the data partition from the user (c.f., Figure 1). The following SQL code shows the attachment of files under a sub-directory to each partition, the creation of a merge table and the addition of partitions to it.

```
call attachDir('tab1', <dir_path>');
...
call attachDir('tabN', <dir_path>');
```

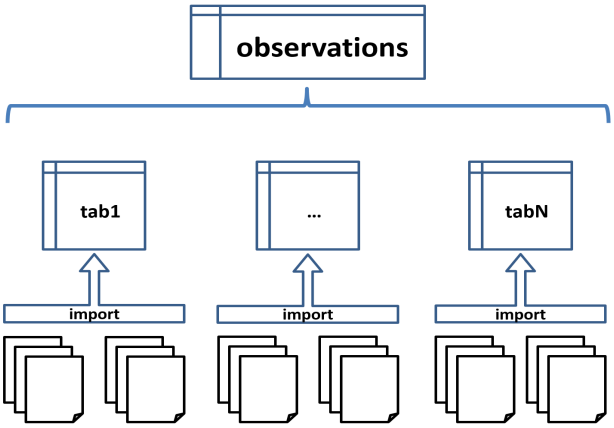

Fig. 1. Merge Table

```
create merge table observations
        (id bigint, x real, y real, z real);

1) alter table observation add table tab1;
...
N) alter table observation add table tabN;
```

2.2. Data processing

At the heart of our DataScope there is a modern column-store, MonetDB [7], which steps away from traditional relational database management systems (RDBMS). Through vertical partitioning of relational tables column-store significantly reduce data access. In our case, vertical partitioning is exploited to reduce the number of columns to be imported. Such data organization improves data compression, simplifies data skipping strategies and it suits well vector processing.

On top, a column-store provides efficient secondary indices for in-memory filtering. In the filtering step, the majority of the queries are range selections. Such type of filtering is speed up using secondary indexes such as column imprints [8]. Column-imprints resembles bitmaps that index ranges of values in each cache line of each column. This makes them very efficient in range queries since they allow skipping cache lines that do not contain data for a desired range. Hence, an imprint is used during query evaluation to limit data access, and thus minimize memory traffic. The compression of imprints is CPU friendly and exploits the empirical observation that data often exhibits local clustering or partial ordering as a side effect of the construction process such as in the EO data collection process.

A coarser grain filtering is achieved through partition elimination which happens at query compilation time. During query compilation, the predicates are inspected and compared with statistical information collected for each partition. Such information is obtained from the file's header, by sampling, or simply by running an analyze call. For the partitions shown

in Figure 1, the following SQL code is used to collect the minimum and maximum value for columns *x, y, and z* of each partition.

```
analyze sys.tab1 (x, y, z) minmax;
...
analyze sys.tabN (x, y, z) minmax;
```

For example, a range selection such as *x between 51.5 and 54.5 and y between 3.75 and 7.75* the query plan will only request the scan of partitions within the 2D bounding box.

2.3. Data distilling

Our DataScope provides a powerful, seamless integration for handling declarative data (SQL), statistics and learning (R), image/timeseries analysis (SciQL). As a last resort, domain specific libraries can be easily linked with the system. R is used to express statistical analysis, simple geographic summaries, and for visualization. N-dimensional array based functions are expressed through SciQL. SQL as declarative language is ideal to express complex ad-hoc queries and data flows, it has been extended with support for the objects and functions defined in the specification of the Simple Feature Specification of the Open Geo-spatial Consortium(OGC) [3].

2.3.1. New functionality

The architecture is designed following a modular approach where new functionality is easily incorporated through new operators or with a seamless integration of specialized libraries. The integration of new functionality follows three steps with each of them related with each MonetDB layer, i.e., front-end layer, intermediate layer (optimization layer), and kernel.

The first step is to expose the new functionality to the user using a SQL procedure or function. The second step is to define the operator for the intermediate query plan represented in MonetDB Assembly Language (MAL). Such operator is interpreted by the DBMS optimizers as an internal operator, and thus considered in all the optimization strategies. The final step happens at the Kernel where a C function is defined and which is linked to the MAL operator at runtime.

```
SQL:
 create procedure listDir(dirPath string)
         external name eo.listdir;
MAL:
 module eo;
 pattern listdir(dirname:str):void
 address eoListDir

C:
 void eoListDir(
         Client cntxt, MalBlkPtr mb,
         MalStkPtr stk, InstrPtr pci);
```

[3] http://www.opengeospatial.org/

2.3.2. External environments

Our solution exploits the fact that deep integration of external environments, such as R and Python, are easy to implement in MonetDB. In the case of R environment, used for statistical computing [9], statistical analyses are integrated through a MonetDB.R client program [10].

Through the R dplyr package MonetDB.R [10] filtering, grouping, and aggregation operations are pushed into the database for execution and only the result is moved to the R environment. This way, large amounts of data can be inspected and filtered out before transferring it to the R front-end to identify, for instance, correlations.

Load and filter a DBMS table from R

```
01: /* load libraries */
02: library(dplyr)
03: library(MonetDB.R)
04:
05: /* Connect to database netcdf */
06: src <- src_monetdb("earth_observation")
07:
08: /* Load table */
09: tbl1 <- tbl(src, "cum_prec")
10:
11: /* Apply filters */
12: tbl2 <- filter(tbl1, lat > 51.5, lat < 54.5)
13: tbl3 <- filter(tbl2, lon > 3.75, lon 7.75)
14:
15: /* Sort data */
16: tbl4 <- filter(tbl3, desc(prec))
17:
18: /* Push down the query and collect result */
19: cumprecdf <- collect(tbl4)
```

Fig. 2. R code

The R code in Figure 2 shows how to connect to a DBMS instance hosting *earth_observation* database and retrieve all cumulative precipitation for a 2D bounding box defined between latitude 51.5 and 54.5 and longitude between 3.75 and 7.75. Using dplyr, the R statements on *line 09* specify the table be accessed, the bounding box range selection is done with statements on *line 12 and 13* while the sorting of the result is done on *line 16*. Once the collect statement on *line 19* is issued, these statements are pushed down to the DBMS, optimized, and executed.

Load and filter a DBMS table from SQL

```
01: select * from cum_prectbl1
02: where
03: lat between 51 and 50.0
04: and lon between 3.75 and 7.75
05: order by prec DESC
```

Fig. 3. SQL code

Such statements are equivalent to the SQL query in Figure 3. In Figure 3 the table identification is done on *line 01*, the range selections on *line 03 and 04* while the sorting is covered by *line 05*.

2.4. A vision in action

Our DataScope ecosystem has emerged through a decade of support for scientific data management and adheres to the vision expressed in [1, 11]. It stands on the shoulders of previous experience on building solutions for remote sensing data processing. A snippet of the history in two domains:

Astronomy We developed our solution by providing an alternative implementation for the *The Sloan Digital Sky Survey (SDSS)*, which aims to map one third for the sky to obtain observations of 100 million objects. Subsequently, the technique was included in the Dutch LOFAR radio telescope software stack, and is currently researched in the context of an Amazon supported project for the Square Kilometer Array telescope.

Remote sensing Our DataScope efficiency was shown with several use cases, ranging heat islands detection for urban planning [11] using LIDAR, cadastral data and ground based sensor data as the data substrate, to satellite image analysis for forest fire detection and precision farming [4]. It facilitated the integration of different heterogeneous data sets for exploration in four dimensions, 3D space and time. It provided flexible and efficient analysis of spatiotemporal data.

In [11] large NetCDF file repositories were attached as well as a large LiDAR file repository. Using MonetDBs R front-end the data sets were explored in search of heat islands at Rotterdam city center. The focus was on the monthly temperature average difference for February 2015 at 4PM. With our DataScope in-situ data access spatiotemporal data referring to Rotterdam city center and February temperature measurements were imported from the attached file repositories. The imported data was then aggregated and sorted before being passed to the R environment. See Figure 2.4.

Step-wise exploration is thus at the distance of a click. A simple zoom-out interaction by the user on his/her web-browser will trigger a chained reaction which would propagate from the R interface all the way down to a data import from the file repositories attached to the DBMS.

3. SUMMARY

In our experience, knowledge and deployment of modern scientific database technology and software technology is a *sine qua non* to push the data science envelop of EO. We shortly introduced a DataScope for remote sensing which provides the means to access and combine geographical data through interlinked data hubs (snippets as data representation), integration and fusion of sensor data, and metadata management to improve data processing.

Our solution combines the best of both file-based and database management systems (open-source) technologies.

[4] EU Projects, e.g. TELEIOS http://www.earthobservatory.eu http://www.linkedeodata.eu/

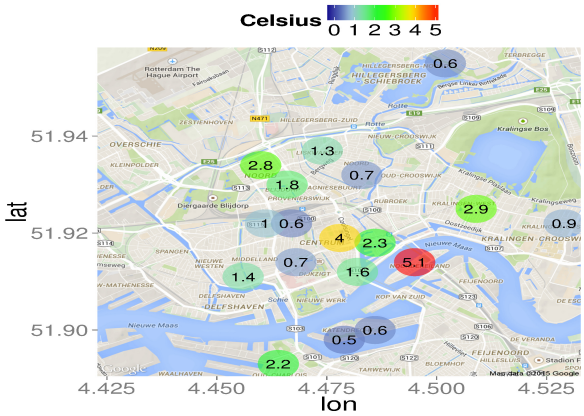

Fig. 4. Rotterdam city center [11]

For data access it efficiently combines heterogeneous data sets while providing in-situ data access. For data processing it exploits external accelerators (HPC and HTC) to accelerate computations. For data exploration its modular architecture allows the integration of specialized functionality without downgrading data processing efficiency.

The results are presented through interactive front-ends and provided as service for downstream data dissemination. Each layer is coupled in seamless way to have shorter times between data acquisition and data presentation.

4. REFERENCES

[1] Martin L. Kersten and et al, "The researcher's guide to the data deluge: Querying a scientific database in just a few seconds," *PVLDB*, pp. 1474–1477, 2011.

[2] Jeffrey Dean and et al, "Mapreduce: Simplified data processing on large clusters," *Commun. ACM*, 2008.

[3] "Microsoft excel-datascope," url-http://research.microsoft.com/enus/projects/exceldatascope/.

[4] "The datascope," http://releases.jhu.edu/2010/11/01/new-jhu-computer-to-enabledata-analysis-not-possible-today/.

[5] Milena Ivanova and et al, "Data Vaults: A Database Welcome to Scientific File Repositories," SSDBM'13.

[6] Prerna Mahajan and et al, "Article: Rough set approach in machine learning: A review," *International Journal of Computer Applications*, pp. 1–13, 2012.

[7] Stratos Idreos and et al, "Monetdb: Two decades of research in column-oriented database architectures," *IEEE Data Engineering Bulletin*, pp. 40–45, 2012.

[8] Lefteris Sidirourgos and et al, "Column imprints: a secondary index structure," in *Proceedings of the ACM SIGMOD*, 2013.

[9] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing'14, 2014.

[10] Hannes Mühleisen and Thomas Lumley, "Best of both worlds: Relational databases and statistics," in *SSDBM'13*.

[11] R. Goncalves and et al, "A round table for multi-disciplinary research on geospatial and climate data," *IEEE e-Science*, 2015.

doi: 10.2788/854791

THE NANSEN-CLOUD: A SCIENTIFIC PLATFORM AS A SERVICE FOR THE NANSEN GROUP

Morten Wergeland Hansen, Anton Korosov, Aleksander Vines, Tor Olaussen, Torill Hamre

Nansen Environmental and Remote Sensing Center
Thormøhlens gate 47
NO-5006 Bergen
Norway

ABSTRACT

We introduce a concept that we call a Scientific Platform as a Service (SPaaS) to aid the management and synergistic use of EO data within the Nansen Group. We present a prototype implementation of the SPaaS (the Nansen-Cloud), which handles data from satellite remote sensing but also in-situ and model data. The Nansen-Cloud can be understood as a system providing the integration of scientific tools, algorithms, data stored at various locations, and various data catalogs, via an application programming interface (API). This should allow users to work on their local desktops, using local CPU with integration to cloud systems to analyze the EO data. In this manner, the SPaaS helps users to focus on the scientific research without bothering where the data is stored or its format, nor the maintenance of the infrastructure or the software.

***Index Terms*—** Earth observation, cloud computing, scientific service

1. INTRODUCTION

The Nansen group consists of environmental research centers in Norway, Russia, China, Bangladesh, India and South Africa. The centers employ Earth Observation (EO), numerical modeling and data assimilation in their research. Generation of common research tools and data access is a challenge. Actual observation data comes from many different satellites and in-situ measurements, and are downloaded locally on a daily basis for operational processing of, e.g., wind retrievals from synthetic aperture radars (SARs) and harmful algae bloom monitoring employing optical instruments (e.g., Fig. (1)). Since the local computing power and storage capacity is limited, coordination between the centers in processing, archiving, accessing and visualizing the available data is a key goal.

In this paper, we present a pilot version of the Nansen-Cloud SPaaS. The planned logical design is presented in section 2, the architectural design in section 3, the physical design in section 4 and some example use cases in section 5. The summary and conclusion follow in section 6.

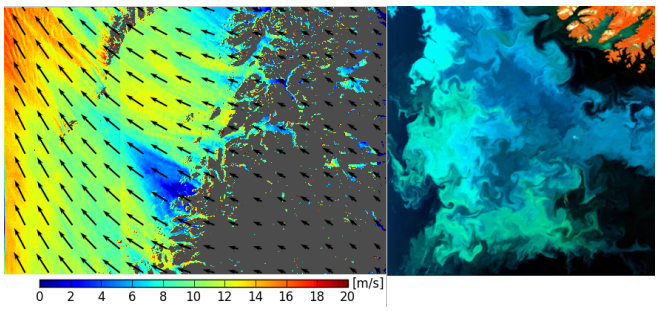

Fig. 1. High resolution wind retrieved from combination of a Sentinel-1a SAR image and an NCEP model wind forecast on 3 October 2014 (left) and an algae bloom in the Barents Sea (right). The Norwegian coast is seen to the right and in the upper right corner in the respective images.

2. LOGICAL DESIGN

2.1. Overall logical structure

The SPaaS concept is inspired by standard cloud computing service models, i.e., Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) [1]. These models typically serve to improve software development speed, allowing the user to focus on the application itself, and to reduce costs by outsourcing hardware and software maintenance to the service provider [2]. As such, The SPaaS proposed for the Nansen-Cloud shall help reducing the time and effort that scientists have to spend searching for data, and to develop their own tools for processing and analysis.

The SPaaS is structured as illustrated in Figure 2. It integrates existing tools and data repositories into a distributed system, by providing the necessary "glue" to make them interoperable and inter-connected. The software layers supporting this interoperability are implemented as modular open source components, allowing the whole system to be open, extensible and scalable. The design objectives will be accomplished by integrating already developed tools with an application program interface (API) to allow easy access to the data in

doi: 10.2788/854791

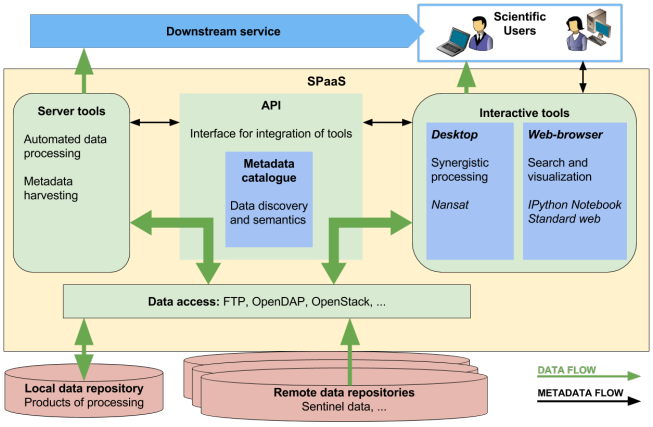

Fig. 2. Technical structure and main data and information flow in the proposed Scientific Platform as a Service (SPaaS).

a single platform. The user experience, i.e., in employing the interactive tools for efficient prototyping, testing and operationalization of multi-sensor synergistic algorithms, is thus of major importance.

The central part of the SPaaS is the metadata catalog. This contains granular metadata describing the structure, location and content of available satellite, modeling and in-situ datasets. Information is stored in a relational database.

The tools (installed on desktops or servers) can be used to search, visualize and perform in-depth analysis of satellite EO, and other relevant datasets, and are connected to the data provisioning services via the API and the metadata catalog. The web-browser tools are used to enable quick and simple search and retrieval of data via the browser, whereas the desktop tools can be used advanced data analysis.

In addition to the desktop clients used in algorithm development and research, the back-end server tools are used for regular and automated processing. Typical tasks performed here are the download, processing and preparation of previews of incoming satellite or model data using qualified algorithms, and ingestion of metadata to the catalog. The server tools can also be used to perform automated indexing of data at other online repositories.

The data access protocols provide seamless access to data from local and remote repositories, viewed as a single virtual distributed file system integrated with the tools, so that any SPaaS member's local data repository is virtually merged into a single shared distributed data archive. This contains satellite data, as well as model (e.g., from the Copernicus Marine Services) and in-situ data (e.g., from the NOAA National Data Buoy Center).

2.2. Logical structure of the metadata catalog

A central table called Dataset contains information about the data provided in a single physical file or data stream, e.g., OpenDAP. Data can be defined on grids, trajectories or points, and the datasets may contain several snapshots of one or more geophysical variables from several time steps. The datasets are described by

- the spatial and temporal coverage (start/stop date, geographical reference, resolution, etc)
- the physical location of the file or data stream (e.g. local file system, FTP, OpenDAP [3])
- the source of the data (platform and instrument)

The metadata follows the Committee on Earth Observation Satellites (CEOS) International Directory Network (IDN) Directory Interchange Format (DIF) format, including the NASA Global Change Master Directory (GCMD) science keywords. Metadata about the geophysical variables contained in the dataset are stored in a separate table.

3. ARCHITECTURAL DESIGN

3.1. Hardware

The SPaaS concept is based on the use of existing hardware such as personal computers and more powerful servers for storage and bulk processing, and for providing the centralized metadata catalog. This is accomplished with standardized and pre-configured VMs.

3.2. Software

The Python/Django framework (`https://www.djangoproject.com`) is used for

- performing temporal, spatial and contextual data search
- providing the high level API
- creation and customization of web-pages

The structure of the database (tables and fields) and common data operations (adding or extracting data from the database) are defined in Django-models. The API is defined in Django-views, and the appearance of web-pages is defined in the Django template language. External metadata catalogs and data repositories such as NORMAP (`http://normap.nersc.no`) and NMDC (`http://nmdc.no`) can be searched directly based on provided API's integrated in Django model managers. Data search and discovery can be performed both via a simple web interface developed in Django, and via the IPython command line or notebook interfaces.

Nansat [4, 5] is included in the VM configurations and is used to read, process, analyze, reproject, and export files and data streams.

For analysis and processing of non-raster geophysical data, scientific python (scipy) and OSGEO GDAL/OGR is

doi: 10.2788/854791

available. A separate package that gives meaning to the data by adding metadata, similar to the Nansat functionality (see [5]), may also be a useful tool to be developed and added.

Spatial and Geographic objects for PostgreSQL (PostGIS) is used on the server-side to provide data access in the production environment, i.e., the stable service type of platform containing release versions of the processing software and data.

SpatiaLite is used for development and testing in local pre-release architectures.

Virtual machines can be deployed either on a local user PC or in a cluster of servers in a cloud. In local PCs, Oracle VM VirtualBox is used as virtualization software. Vagrant (`https://www.vagrantup.com/`) is used as the orchestrator - a manager which helps to create, provision and run virtual machines. In the case of external servers or IaaS, any virtualization software can be used (e.g., VMWare [`http://www.vmware.com/`]). Provisioning of the virtual machines is performed using the Ansible provisioning tool (`http://www.ansible.com/`). Ansible automatically installs and configures software and data access on the machine. The system is platform independent and, once set up, can be changed and redistributed easily. Through the use of virtualization, the host operating system (OS) is untouched and the virtual OS is integrated to allow the use of any development environment. The guest OS provides the SSH-, HTTP-, and X-protocols to communicate with the host.

Distributed version control with GIT lets users and developers clone the central repository to maintain local version control. The central repository is served via GitHub (`http://github.com`).

Automatic and continuous collection, compilation and testing of the software is accomplished with TravisCI (`https://travis-ci.org/`), which integrates with GitHub.

We plan to use OpenStack (`https://www.openstack.org`) for controlling the server-side compute, storage and networking resources (IaaS) for Nansen-Cloud.

We plan to use OpenShift (`https://www.openshift.com`) for server-side deployment of the Nansen-Cloud platform (PaaS). A definition file tells the PaaS the requirements of the Nansen-Cloud applications (e.g., a specific processing chain), and allows for automation of the configuration through scripts, and scalability to automatically create additional instances of the application.

4. DATA DESIGN

The data storage framework is aligned with international interoperability efforts as described through the Infrastructure for Spatial Information in the European Community (INSPIRE), World Meteorological Organization (WMO) Information System (WIS) and GEOSS specifications. This ensures easy and open data access for the research communities.

5. EXAMPLE USE CASES

5.1. Installation and synchronization of common programming environment

In order to start using Nansen-Cloud or any of its components a user needs to install four open source cross-platform packages available for Windows, Mac and Linux: GIT, Vagrant, Ansible and VirtualBox. These packages are available for free, and their installation is system specific. After this step a user needs to open a terminal window and type two commands:

```
# Get VM configurations
git clone https://github.com/
      nansencenter/nersc-vagrant.git
# Start course VM
vagrant up course
```

This will download and install a minimal version of Linux Ubuntu 14.04 in a virtual machine, provision it with all required third party libraries, and install the Nansat and Nansen-Cloud software. In order to update the code or any components of the programming environment a user needs to run the following commands:

```
# Update VM configurations
git pull
# Update course VM
vagrant provision course
```

5.2. Search and collocate datasets

For batch processing of many satellite images or for searching collocated datasets, a user needs to first add metadata to the Nansen-Cloud catalog using the following command:

```
# Ingest dataset metadata to catalog
python manage.py ingest <path>
```

Next, after opening the IPython terminal the following commands can be used:

```
# Fetch all datasets from the database
allDs = Dataset.objects.all()

# Fetch only MODIS images
modisDs = Dataset.objects.filter(
  source__instrument__short_name
      ='MODIS')

# Fetch ASAR images overlapping with
# the first MODIS image
asarDatasets = Dataset.objects.filter(
  source__instrument__short_name='ASAR'
).filter(
  datalocation__geometry__intersects=
  modisDs[0].datalocation__geometry
```


doi: 10.2788/854791

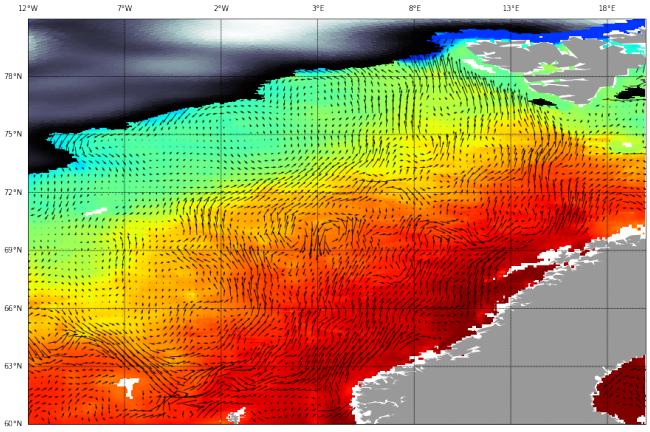

Fig. 3. An example illustration of collocated sea surface temperature, sea ice concentration, sea surface geostrophic current on 1 July 2013 from NORMAP.

)

When one or more datasets are discovered, they can be opened with Nansat (or other software installed on the VM) for inter-comparison, visualization, and further analysis. An example image containing collocated sea surface temperature (SST), ice concentration, and sea surface geostrophic served via OpenDAP in the NORMAP project is shown in Figure 3. The code for creating this image with Nansat, together with many other examples, is provided in the Nansat lecture notes (`https://github.com/nansencenter/nansat-lectures`).

6. STATUS AND SUMMARY

In a pilot version of a SPaaS (the Nansen-Cloud), we have prepared virtual machine (VM) configurations that enable users to install all the software components needed to do Earth observation data analysis with Nansat and scientific Python on local laptops and desktops. The present configurations of Nansen-Cloud VMs are open-source and provide access to data on local servers (presently, data access is restricted to the NERSC network), e.g., SAR data from Envisat ASAR, Radarsat-2, and Sentinel-1, Landsat, MODIS, MERIS, and other sources. The VM configurations may be easily modified to fit other networks. The data may be ingested in a local metadata catalog for search and retrieval, including automatic collocation. The catalog is available in a website, enabling browser search in a map at given time intervals. In addition to this browser data discovery engine, it is also possible to search the catalog directly in the IPython command line tool. Most importantly, the browser and command line data discovery tools enable the user to easily locate and extract the actual observational data for further analysis, e.g. using Nansat. Advantages of Nansen-Cloud are exemplified in use cases (section 5).

7. REFERENCES

[1] Q. Zhang, L. Cheng, and R. Boutaba, "Cloud computing: state-of-the-art and research challenges," *Journal of internet services and applications*, 2010.

[2] National Institute of Standards and Technologies, "The NIST Definition of Cloud Computing," Tech. Rep., National Institute of Standards and Technologies, 2011.

[3] Peter Cornillon, James Gallagher, and Tom Sgorous, "OPeNDAP: Access data in a distributed, heterogeneous environment," *Data Science Journal*, 2003.

[4] Anton Korosov, Morten W. Hansen, Asuka Yamakawa, Knut-Frode Dagestad, Aleksander Vines, Maik Riechert, and Alexander Myasoedov, "nansat: Nansat-0.6.8," `http://dx.doi.org/10.5281/zenodo.45703`, Feb. 2016.

[5] Aleksander Vines, Morten Wergeland Hansen, Asuka Yamakawa, Tor Olaussen, Torill Hamre, and Anton Korosov, "NANSAT - a scientist friendly Python toolbox for processing 2D satellite EO and model raster data," in *Proceedings of BiDS'16*, 2016.

THE ITALIAN DATA PROCESSING CENTRE FOR THE GAIA MISSION

R. Messineo[1], *D. Busonero*[2], *A, Mulone*[1], *R. Morbidelli*[2], *L. Bramante*[1], *A. Vecchiato*[2], *F. Solitro*[1]
M. Lattanzi[2], *M. Martino*[1], *A. Riva*[2], M. Montironi[1]

[1]ALTEC S.p.A Corso Marche 79, 10141, Torino Italy http://www.altecspace.it/

[2]INAF-OATo Strada Osservatorio 20, 10025, Pino Torinese (To) http://www.oato.inaf.it/

ABSTRACT

The DPCT is the Italian Data Processing Centre (DPCT), built by ALTEC in collaboration with OATo (the Astrophysical Observatory of Turin) and funded by ASI, to support the Italian participation to the Gaia mission data processing tasks. DPCT operations started about one month after the Gaia satellite launch and since then data have been received and processed without interruption with a huge numbers of executed workflows and jobs and databases size of tens of terabytes. The paper focuses on the description of the DPCT software and hardware architecture built to manage both big data volumes and complex operations scenarios. The DPCT system has been designed as an integrated system with the capability to manage all the phases of the data processing pipeline: data receiving, data processing, data extraction, data archiving and data sending. In addition, the DPCT system includes also data access and analysis tools allowing Italian scientists to be active users during operations.

Index Terms— Gaia. DPCT, Processing, VLDB, Data Analysis, Architecture, Space Mission and Operations.

1. INTRODUCTION

The Italian Data Processing Centre (also called DPCT, where T stands for Turin), has to be conceived as the set of products, people and processes necessary to the Astrometric Verification Unit (AVU) and GAREQ data processing as planned by the Data Processing and Analysis Consortium (DPAC) [1][2]. In addition to these processing functions, the DPCT will support other CU3 activities, in particular the hosting and operations of the Initial GAIA Source List (IGSL) database. The DPCT is designed, implemented and hosted at the ALTEC center in Turin – Italy.

In Gaia, the CU3 AVU unit is responsible for the development and maintenance of the following three CU3 software systems: the Base angle monitoring system (BAM) [3], the Astrometric Instrument Model (AIM) [4], and the Global Sphere Reconstruction (GSR) [6]. AIM is the software in charge of processing the Astro data telemetry in order to monitor and analyse the Astro instrument response over the mission lifetime. AVU/BAM is the software in charge of processing the BAM telemetry in order to monitor and analyse the BAM behaviour over time. GSR is the mathematical and numerical framework that shall be used to perform an astrometric sphere reconstruction and to compare it with that produced by AGIS. The GAia Relativistic Experiment (on) Quadrupole (GAREQ) is a scientific experiment to evaluate the light deflection due to the quadrupole components of the masses of some selected giant planets observed by Gaia.

The scientific components hosted at the DPCT are different in term of data volumes, amount of processing, algorithm complexity, and operations. The DPCT architecture has been designed considering all functional and non-functional specifications as required by scientific software and trying to find common elements to reduce the architecture complexity. The basic principles behind the DPCT architecture are: to maximize the usage of general purpose hardware, the reuse of software components and the adoption of widespread technologies. The resulting architecture is described in the next document sections as well as the results from the first two years of mission operations.

2. ARCHITECTURE

The overall architecture has been designed to reach the following goals: scalability, performance, high-availability, reliability, modularity and security. The resulting architecture is an integrated system of software and hardware designed to implement the DPCT requirements and managed by a team having skills and competence allowing to maximize the hardware infrastructure through the coding of software running on the selected hardware components.

2.1 SOFTWARE

The software is a heterogeneous and complex software system, composed by many subsystems structured in an open architecture in which each subsystem is a peer responsible for a different class of services. The DPCT infrastructure software is coded using the Java object language as it was

chosen by the Gaia system architecture group and many software framework and library were already available. This section provides the subsystem decomposition of the DPCT infrastructure software. A subsystem is characterized by the service it provides to other subsystems. The implementation of the subsystems is targeted to minimize the coupling among subsystems and to maximize cohesion among classes within a subsystem. The high level subsystem decomposition, described in the next UML component diagram, is derived by functional requirements and it is shown in Figure 1.

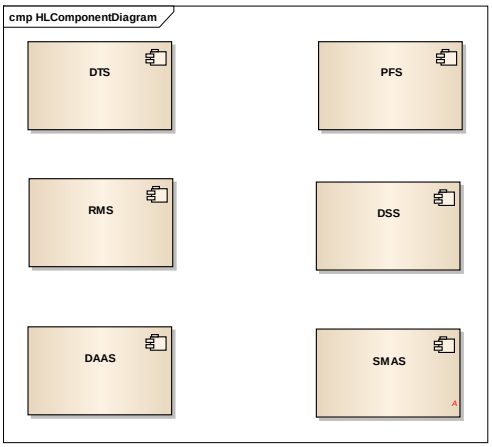

Figure 1 High Level subsystems decomposition.

The RMS provide the control logic to coordinate and manage the software execution infrastructure provided by the PFS, allowing the definition and scheduling of all jobs in charge of performing the required scientific operations based on the received data via a workflow. The main feature provided by the PFS is running and controlling the execution of all defined jobs based on the input provided by the RMS. It is a distributed software which manages all hardware resource provided by the processing cluster. The Data Storage Subsystem is a heterogeneous subsystem including databases, cluster file system, data access layer software and application dedicated software to perform data ingestion and data extraction in the Gaia mission file format. Persistent data management is the main critical point of the overall DPCT.
The most important elements of the DSS are the databases built with the Oracle RAC technology, which are expected to reach the size of about 350 TB. The Data Transfer Subsystem manages the regular data transfer with the ESAC DPC (DPCE) and other external computing centre like CINECA that will provide support to run complex scientific algorithms that cannot run on the local hardware resource. The System Monitoring and Administration Subsystem (SMAS) provides the functions necessary to manage the whole system: configuration, monitoring, reporting, etc. It provides all the application components used by the DPCT operator to manage and follow operations for infrastructural aspects. The DAAS provides the scientists with all the information needed by the science team to supervise all the processing activity and to perform the specific data analysis in support of the Gaia mission assigned to DPCT [9].

2.1 HARDWARE

From the hardware point of view the infrastructure is assembled with general purpose hardware integrated through standards interfaces. The core of the infrastructure is the storage tier built with two HP P7400 storage arrays providing in the current configuration about 900 TB of raw space disk of different typologies: high speed disks and high capacity disks. The storage size is expected to increase before the end of the mission as the DPCT will maintain the input and the output of the pipeline online also after the mission end. The computation tier is built with HP general purpose servers providing a total of 500 CPU cores and 4TB of RAM.
In addition, the DPCT can use the computation power of the CINECA supercomputer to run part of the GSR pipeline which needs HPC resources. The Network tier provides all network services of DPCT; both LAN, WAN. There are several internal networks, high performance network low latency 10 Gbps are used in the infrastructure core while standard performance network 1 Gbps interconnect core infrastructure to client workstation in the operations room. Internet is the baseline for the data exchange among all DPCs. Current installed network has a bandwidth of 1Gbps Mbps but data transfers are executed maximum at 500 Mbps due to Aspera COTS license limitation. The ALTEC Internet connectivity is provided via the GARR Italian Research network. Finally, the infrastructure has a dedicated backup system allowing backup on two levels: disk and tape.
The backup system works continuously in order to have updated backup to be stored also in external sites dedicated to maintain data in case of disaster.
Remote connectivity is allowed to scientists and operators in order to use DPCT services from everywhere via VPN.

3. DAILY PIPELINES

DPCT runs three daily pipelines: ingestion pipeline, AIM processing and BAM processing pipeline. Figure 2 shows the Daily processing pipelines and the overall dataflow. Daily pipelines have been designed to be automatic as they run unattended in all days of the year.

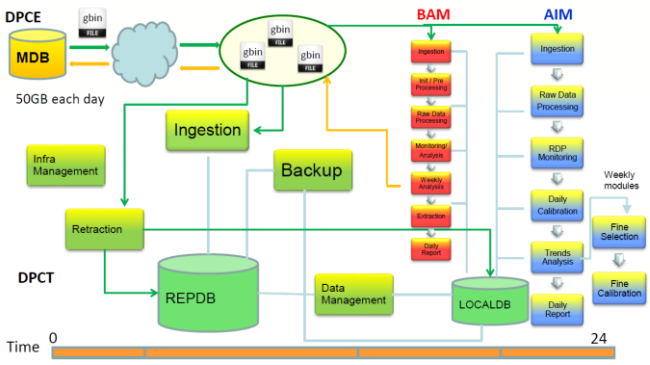

Figure 2 Daily Pipelines

In the ingestion pipeline data received from the DPCE (Data Processing Center ESAC) are checked, filtered if needed and stored into the processing database and into the repository database. In the repository database all received data (housekeeping to L1 data products) are stored for long term archiving. Both the AIM and the BAM processing pipelines process raw observations collected on 24 hours of satellite acquisition and perform analysis in order to have a report as final product to be distributed to scientists and payload experts.

The AIM processing pipeline is the most complex on daily basis in terms of number of astrometric observations received as input, it processes from 2.5 million to 11 million of objects, and of resources needed to perform the processing as the AIM raw data processing algorithms are CPU demanding. The AIM system executes modules performing raw data processing, image parameter determination, LSF/PSF modelization and calibration and AF monitoring and diagnostics.

The BAM processing pipelines is similar to the AIM one but it processes less objects. The observation acquired in the BAM payload are about 8 thousand. This means that the resulting BAM processing lasts much less than the AIM processing, one hour vs. 16 hours. The required processing time impacts the reprocessing capability that is requested to the DPCs to deal with direct or indirect errors in processing. A dedicated environment is configured on the operations platform to manage data reprocessing if needed.

4. DRC PIPELINES

DPCT runs two data reduction pipelines: the data management pipeline and the GSR processing pipeline. Figure 3 shows the DRC processing pipelines and the overall dataflow.

The DRC data reduction pipeline has the purpose of consolidating the dataset received on a daily basis by means of data consistency checks, data retraction and data organization. Over the Gaia mission lifetime the data volume expected as input to the DRC data management pipelines will increase over cycles hundreds of billions of objects are foreseen at the end of the mission.

After that the cyclic dataset is consolidated the pre-processing phase is executed to select input data for the GSR processing and implement simple transformations about data model. After the dataset for a variable number of astrometric sources, from few million to one hundred million, has been prepared, the GSR processing pipeline starts the execution of seven modules, called in sequence by the infrastructure software, which provides the overall workflow and the database interaction functionalities.

The total number of observations to be processed by GSR at the end of the mission, with runs of 100 million of sources, is about 8000 million. The execution of this processing will run on HP DL580 general purpose servers and it will last days.

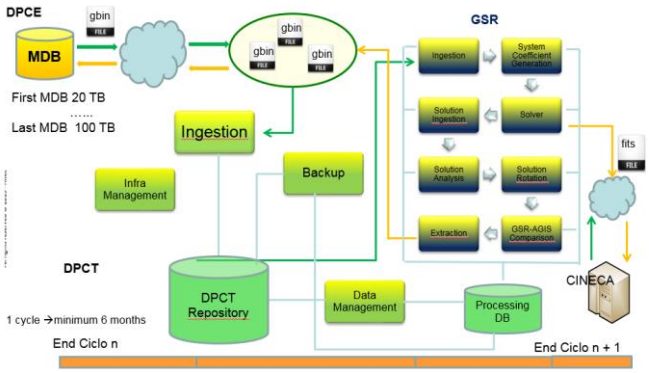

Figure 3 DRC Pipelines

The software is mainly developed using the Java language, with the exception of the Solver module.

All the software needs to run with a high degree of parallelism. This is achieved in different ways for the Java and for the C part. The activity performed by the former usually can be split into several independent jobs which are managed by a coordinator: basically a queue manager that can publish and dispatch them to the available computing resources through a whiteboard system.

The Solver module, implementing a parallelized version of the PC-LSQR algorithm, is executed on the CINECA tier-0 super computer as it needs an HPC infrastructure [7]. In order to use computing services of the super computer its implementation is based on hybrid MPI/OpenMP technologies that exploits both the huge number of CPUs of the massive parallel supercomputers, and the multithreading capabilities of the single computing nodes.

doi: 10.2788/854791

5. OPERATIONS

DPCT operations started immediately after the launch activating its operations team and procedures. The DPCT is operated by an integrated team composed by ALTEC personnel, responsible for the DPCT overall infrastructure, and INAF-OATo personnel, responsible for all scientific aspects. The DPCT Operations Team is composed by two collaborating teams: DPCT Infra Operations Team and DPCT Scientific Operations Team. In general the DPCT operator carries out the current operations of the DPCT, following operational procedures. He is able to solve accidents described in the operational procedures and investigates other accidents. During the commissioning phase the two main DPCT objectives were: to support the AVU payload experts [2], [5] and to setup the pipelines, teams, procedures and software in order to be ready for the nominal phase. The early operations started on the 25th of July 2014, with the end of commissioning phase and the start of data segment 0 acquisition. The figures of processing and data collected from the start of the mission are very huge and allow us to day that the Gaia-DPCT project is a big data project though the system was designed before spread of big data technologies.

The DPCT has processed at time to write this paper 750 runs of AIM and BAM. The number of workflow executed is about 90000 and the number of jobs processed is about 8.9 million. The size of received input data is about 50 TB while the largest Oracle database ad DPCT has a size of 100 TB. Finally, at the beginning of the 2016 the cyclic processing system GSR will move to operation and the activities at DPCT will increase in terms of processing and data size.

The Gaia mission lifecycle is about five years and one year of extension is under evaluation. DPCT will continue processing activities beyond the mission end as at least two year of data reduction processing are scheduled to complete processing on consolidate data and release the final mission data products. The processing complexity will increase over time and it will be managed through several actions: additional procurement steps allowing to exceed the PetaByte of raw disk space, software improvement focused on scalability and strengthening of the data management system including the backup part.

6. CONCLUSION

In this paper we described the infrastructure, the pipelines complexity and the operations of the DPCT project highlighting the big data elements. The project definition phase was before the birth of the “big data” area but the DPCT can be seen an example of big data project. Moreover in the future the data collected at DPCT could be used in other projects dedicated to data exploitation also integrating data coming from several space mission [10].

7. ACKNOWLEDGMENTS

The authors are members of the Gaia Data Processing and Analysis Consortium (DPAC) and this work has been supported by the ASI contracts 2013/030/I.0 and I/058/10/0 (Prof. M. Lattanzi, INAF-OATo, is the Italian Principal Investigator).

7. REFERENCES

[1] Messineo, R., Morbidelli, R., Martino, M., Pigozzi, E., Mulone, A. F., and Vecchiato, A., “Italian dpc infrastructure and operations to support of the italian participation to the gaia data processing” in [Space Telescopes and Instrumentation 2012], Proc. SPIE 8451-13 (2012).

[2] M.G. Lattanzi, R. Morbidelli, R. Drimmel, M. Sarasso, D. Busonero, M.T. Crosta, D. Gardiol, A. Vecchiato INAF - Astronomical Observatory of Torino; M. Martino, A. Ciampolini, R. Messineo, A. Mulone, E. Pigozzi, V.Icardi, F. Solitro - ALTEC S.p.A; M. M. Castronuovo ASI – Agenzia Spaziale Italiana “Gaia DPCT the Italian contribution to Gaia data processing and archiving” IAU Symposium 276, The Astrophysics of Planetary Systems: Formation, Structure, and Dynamical Evolution, Torino - Italy, 11-15 October 2010.

[3] Busonero, D., Lattanzi, M., Gai, M., Licata, E., & Messineo, “Running AIM: initial data treatment and marcsec level calibration procedures for Gaia within the astrometric verification unit”, in [Space Telescopes and Instrumentation 2014], Proc. SPIE, Vol. 9150, 91500K (2014)

[4] Busonero, Deborah, “Astrometric instrument modeling in the context of Gaia astrometric verification: tasks and activities”, in [Space Telescopes and Instrumentation 2012], Proc. SPIE, Vol. 8449, 84490F (2012)

[5] Riva, A., Gai, M., Lattanzi, M. G., Russo, F., & Buzzi, “BAM: A metrology device for a high precision astrometric mission”, RMXAC, Vol. 45, 35 (2014)

[6] Riva, Alberto; Lattanzi, Mario G.; Drimmel, Ronald; Gai, Mario; Busonero, Deborah; Buzzi, Raffaella; Pecoraro, Marco; Russo, Federico; Messineo, Rosario; “The Basic Angle Monitoring (BAM) software tool in the context of Gaia's astrometric verification”, in [Space Telescopes and Instrumentation 2014], Proc. SPIE, Vol. 9150, 91501Z (2014)

[7] Vecchiato Alberto et al.; “The global sphere reconstruction for the Gaia mission in the Astrometric Verification Unit” in [Space Telescopes and Instrumentation 2012], Proc. SPIE, Vol. 8451, 84513C (2012)

[8] Bandieramonte, M., U. Becciani, A. Vecchiato, Mario G. Lattanzi, and B. Bucciarelli. “AVU-GSR Gaia Mission: A Hybrid Solution for HPC and Grid-MPI Infrastructures”. In: Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE), 2012 IEEE 21st International Workshop on. Ed. by Sumitra Reddy and Khalil Drira, pp. 167–172

[9] A. Mulone, R. Morbidelli, R. Messineo, R. De March, F. Filippi, M. Vaschetto, F. Sella, S. Uzzi, C. Manetta. “Data Access Services at DPCT to support scientists’ online and offline data analysis”, to appear in Proc. Big Data from Space (BiDS16). IEEE, 2016.

[10] R. De March et al., “TEmporal Characterization of the remote SEnsors response to radiation damage in L2”, to appear in Proc. Big Data from Space (BiDS16). IEEE, 2016.

INTEGRATED ENVIRONMENTAL MONITORING IN MOUNTAINS WITHIN THE SENTINEL ALPINE OBSERVATORY AT EURAC RESEARCH

Roberto Monsorno, Marcello Petitta, Alexander Jacob, Armin Costa, Bartolomeo Ventura, Claudia Notarnicola, Marc Zebisch

Institute for Applied Remote Sensing, EURAC, European Academy, Bolzano, Italy

ABSTRACT

The Sentinel Alpine Observatory is a new initiative by EURAC to foster Sentinel based remote sensing applications for mountains and, specifically, for the European Alps. It comprises a long-term archive for Sentinel data over the Alps, processing approaches adapted for mountains and the development of tailored applications for users from mountain regions on snow, vegetation, soil moisture and other relevant variables. Big Data both from space and from in-situ measurements are analyzed organizing them into a specific facility composed by a Software Defined Storage, a Processing Cluster and a Data Base Cluster, implementing the DataCube multi-dimensional array.

Index Terms— Sentinel, Environmental Monitoring, Alps, remote sensing, data integration.

1. INTRODUCTION

Remote sensing data represent one of the most important source of spatially explicit environmental information currently available. In EURAC, within the Institute for Applied Remote Sensing, we aim at exploiting remote sensing data for environmental analysis mainly in the Alpine area and in other mountain regions. A novelty approach to integrate remote sensing, in-situ and model data has been developed and is currently used for different assessments, from vulnerability to resilience in alpine territory.

The region in which the Institute is based, South Tyrol – Italy, has a large historical archive of environmental data collected by several agencies. Meteorological, ecological, land use, land cover, administrative information are collected, stored and used in several regional projects. Moreover, with the effort of a considerable number of projects and the associated instrumentation acquired, we consider South Tyrol and the associated collected data, a *virtual laboratory* for Alpine Environment.

Nowadays the availability of Sentinel data creates the necessity to manage a large volume of heterogeneous data and datasets. Moreover, all those data have to be available for research and administrative use in real-time.

During the last years, a new, modern infrastructure has been developed based on high-volume storage, computational power, network speed and timely availability using state-of-the-art technology, resulting in a scalable and performant structure for research and real-time data management, production and analysis.

2. INFRASTRUCTURE

An integrated approach to monitor the environment requires dealing with heterogeneous environmental data, namely:

- Earth Observation (EO) data form several satellite missions. Since March 2009, NASA missions Aqua and Terra have been transmitting images of Europe to the EURAC Ground Segment, and since February 2012, the NASA mission Suomi-NPP has been added to the daily mission scheduler.
- In-situ data collected with a distributed network of ground stations, which include meteorological instruments, soil moisture sensors and other instruments which acquire at daily level several environmental variables related to ecology, meteorology and phenology.
- UAV (Unmanned Aerial Vehicle) acquisitions. Since 2014, EURAC collects data using a UAV equipped with several optical sensors such as a hyperspectral camera, a modified IR (Infrared) camera and a standard RGB camera.
- Model data. Beside observations, EURAC researchers produce large amount of model data related to environmental research, such as atmospheric energy balance, solar incoming radiation, pollutant concentration and other relevant output.

In order to acquire, manage, compute, distribute, and use this heterogeneous and large amount of data it was necessary to adapt EURAC infrastructure expanding and designing 3 crucial systems: i) the network layer, ii) the storage layer and iii) the processing layer.

A scheme of the proposed infrastructure is presented in Figure 1 where the link with the data providers (ESA, ASI, antenna and other confederated partners) is shown. In the center of the figure, the infrastructure scheme is depicted,

doi: 10.2788/854791

including the information about the single components such as the network, the storage and the processing cluster.

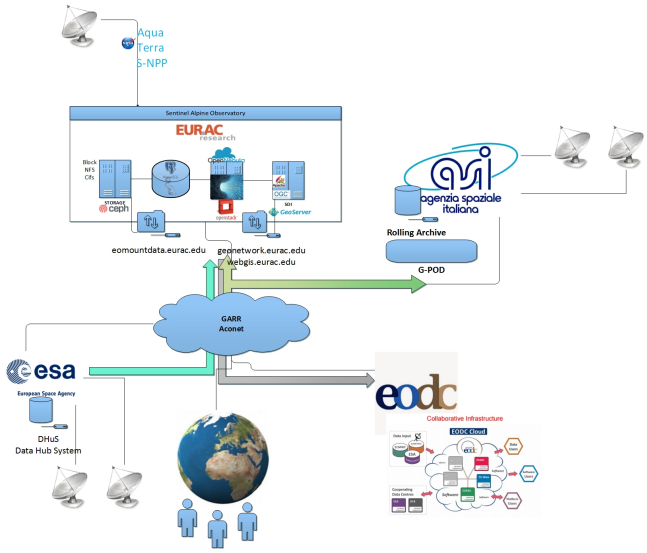

Figure 1: The overview design of the EURAC computational infrastructure is presented here. In the upper part the "Sentinel Alpine Observatory" which is the central node where data from different sources, such the EURAC Antenna, ESA, ASI (Italian Space Agency) and other Collaborative Infrastructure are collected, transformed and distributed

2.1. Network

With the launch of the new ESA Sentinel satellites and the availability of the associated datasets, the needs to transfer very large amount of data, in form of images, is increased. ESA as the Sentinel Data provider offers automatic download of data of interest for the research, but in order to have an efficient use of the available information, a strong network connection and a reliable bandwidth are compulsory requirements.

EURAC, in collaboration with LUB (Free University of Bozen – Libera Università di Bolzano) established a connection with 10 Gbps to the Austrian (ACOnet) and Italian (GARR) scientific and education National networks, while EURAC and the Bolzano University are directly connected at 40 Gbps. This bandwidth ensures a fast acquisition of data, allowing EURAC to process and distribute new products derived from Sentinel satellites timely.

2.2. Storage

EURAC scientific work is based on the analysis of EO and environmental data and the estimated annual need of storage is of 100 TB. In 2014 we had distributed storage servers and distributed data. The need was to have a solution where we could aggregate all data, for archiving, for processing and for dissemination. Therefore, it was necessary to have a scalable solution at the PB (Petabyte) scale.

For the foreseen storage platform, the following features have been realized:

- on-disk storage at low cost,
- one unique file system, remotely accessible using NFS-like protocols (seen as a single disk on client side)
- horizontally scalable for storage from a few terabytes to peta-scale, by adding new storage nodes,
- scalability for performances, to allow massive distributed processing,
- easy to manage,
- fault tolerant,
- open source file system working on linux. After a benchmarking analysis, we adopted a CEPH Software Defined Storage solution [1].

2.3. Processing Cluster

On top of the storage layer and with the need to have an integration with it, we need a cluster, which enables processing and dissemination of the collected data. In order to control and manage the networking, the storage and the processing systems we introduced a cloud operating system, open source and with strong and active community development for the API. The solution adopted is OpenNebula [2]. We are testing the solution on a small cluster, while we are deploying new hardware, in terms of performant blade servers, where we will deploy our working environment. Since the OpenNebula solution is a cloud computing platform for private and public clouds, it will enable us for a possible infrastructure federation and interoperability.

2.4. Analysis DB

An efficient data organization for our research and operational products at EURAC is crucial, since the volume and the variety of the data is growing by the day. We are developing different products for the alpine area such as snow cover, soil moisture and leaf area index maps. These products have to be stored in an easily accessible way for further analysis and creation of maps and analysis services. Currently we have setup a multi-dimensional array database, Rasdaman, [4], [5], [6] which is a middleware on top of postgresql & postgis for efficient querying and management of raster data.

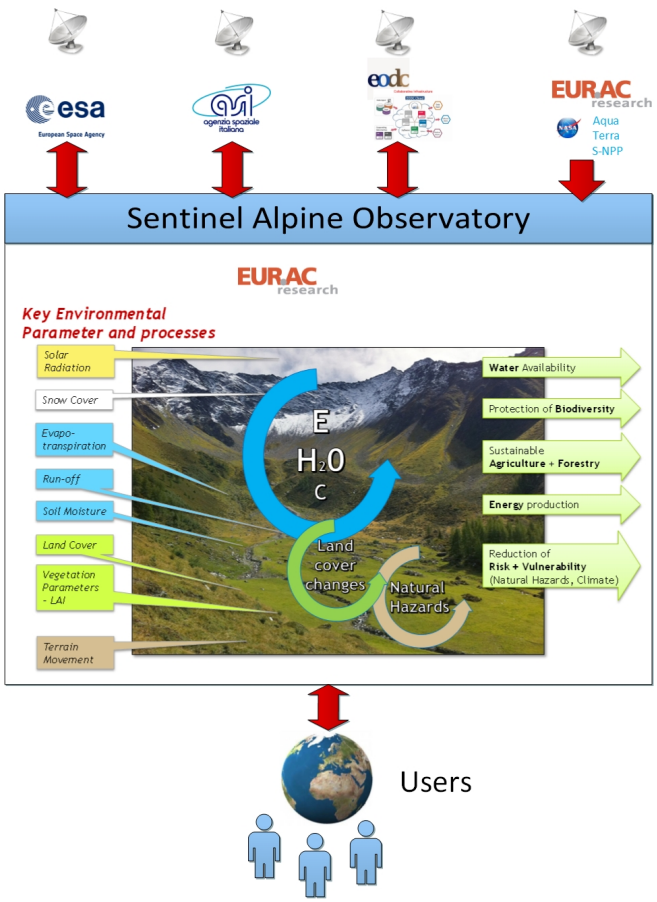

Figure 2: The conceptual design of the Sentinel Alpine Observatory in EURAC is presented. The Observatory exchange data with the different providers (ESA, ASI, Collaborative Segments and the Antenna). Different algorithms and models are applied to the data in order to retrieve several parameters such as Solar Radiation, Snow cover, Evapotranspiration, Run-off, soil moisture, Land cove, Vegetation parameters. All those variables are then used to provide services associated to water availability, Energy production, etc. to end users and stakeholders

The computational requirements are rather high, since from one side we have to produce EO data with in-house developed algorithms, from the other side we have complex models which use several datasets.
A typical example is the evaluation of the solar radiation for the entire region, South Tyrol, with a resolution of 0.5 m [3]. This calculation is rather intense for quasi-real-time use. For example to cover 1500 km² at 0.5m of resolution of the South Tyrol Region, we split the region in 125 squared tiles, each one of a size 16 km. To estimate the monthly sum of the solar radiation for each pixel within a tile, it took up to 10 hours, by using a small cluster, with 32 CPU and 192 GB of RAM. Based on previous considerations, the processing power which covers all the above mentioned needs can be quantified in 280 CPUs (cores), 448 GPU and 4 TB RAM, which is the current structure of the system.

3. ALPINE APPLICATIONS AND SERVICES

The strategy under the Satellite Alpine Observatory (SAO) is to provide services ready-to-use in quasi-real-time, tailored for the needs of the specific user.
The SAO, as shown in Figure 1, can be considered as a layer between the data providers such as ESA, ASI, EURAC and the final users. SAO aims at providing not only environmental variables retrieved from data and models, but also to deliver services associated to environmental analysis. More general, the role of such a complex structure serves and responds to clear needs from two classes of end-users. On the one hand the researchers who need to access quickly large amount and different kind of data in order to produce reliable and meaningful results; on the other hand, stakeholders which use the data and the surrogated products to take decisions. Users comprise decision makers, energy providers, citizens and tourists.
Some of the most relevant services produced by the SAO from Sentinel platforms are:

- Wet snow maps at high resolution with Sentinel 1 data over the Alps
- snow maps at high resolution with Sentinel 2 data over the Alps (viewable snow)
- Grassland LAI maps with Sentinel 2 and Sentinel 3
- Change detection maps in forested areas with Sentinel 1 and Sentinel 2
- Soil moisture over the Alps with Sentinel 1.

In addiction, SAO is currently producing other products at lower resolution derived from MODIS:

- Air Quality PM10: The PM10 (Particulate Matter) maps are based on MODIS Aerosol Optical Depth (AOD) data with a resolution of 10 km. The input data are the MOD04 products, the Boundary Layer Height (BLH) and the PM data from a set of ground stations. The processing steps include the estimation of the relationship between the PM, the AOD and the BLH for historical time series in each of the available stations. Then, a fitted function is used to estimate the PM concentration for the entire region of interest. The method, the validation and the error analysis are published in several papers [7], [8], [9]. The data are distributed to regional environmental agencies and researcher.

doi: 10.2788/854791

- MODIS daily snow cover: The snow cover map based on MODIS data has a resolution of 250 m. The input data are the atmospherically-corrected reflectances of MODIS MOD09GQ, MOD09GA and the MODIS level 1 product MOD021KM. The processing steps include converting the map projection (Sinusoidal to UTM WGS 84) and a topographic correction based on a digital elevation model (DEM). Then snow and clouds are classified using a decision-tree algorithm. Snow under the forest canopy is detected by comparing the Normalized Difference Vegetation Index (NDVI) of a winter image and a reference image from summer season. To enhance the cloud classification a cloud cover derived from MODIS level 1 product, MOD021KM is also included [10].
- Surface Solar Radiation in South Tyrol: EURAC developed a freely accessible, geo-referenced solar potential database, which will allow private users and public bodies to estimate the solar potential on a roof level, and to give concrete recommendations to public authorities and politicians about the relevance of solar energy and the potential to further develop solar energy in Tyrol and the Province of Bolzano [11].

4. VOLUME VARIETY VERACITY

Considering that EURAC is planning to archive Sentinel 1, 2 and 3 data over the Alps, a data volume of 40 TB/year is expected as minimum. Considering a life of 5 years only raw data will rise up to hundreds of TB in few years.

On top of this the in-situ data collection as well as flight campaigns with high resolution multispectral, RGB and thermal cameras produce additional data.

The variety of data types is a big data issue. The best solution is to have a multiple DB infrastructure, which allows handling several heterogeneous data types.

For in-situ data, the quality of the data is very important: this implies to automatize a basic quality check on raw and level-1 data.

All subsystems presented in the former paragraph are orchestrated by a Data Exchange Server (DES) a java based multithread application which is monitoring and organizing interaction between the systems.

5. OUTLOOK TO FEDERATED INFRASTRUCTURE

The medium- and long-term perspectives focus on the following objectives:

- Store and maintain historical data which can be used for Climate Change identification and study, as well as exploited for delivering sensitive information to decision makers and stakeholders.
- Establish processes to facilitate the access and the use of large dataset to scientist, researcher, public entities, SMEs and common citizens, in order to target, when possible, the EU request of open data
- The collaboration with other institutions becomes central for the long-term perspective and for a reliable and accurate data management.

The resources at EURAC have been sized on the specific needs of the research on mountain environment. The future work for us is to open the Sentinel Alpine Observatory to a possible federated scientific infrastructure on mountain areas, within the framework of the Sentinel Collaborative Ground Segment.

6. REFERENCES

[1] http://ceph.com

[2] http://opennebula.org/

[3] http://webgis.eurac.edu/solartirol/?lang=ita

[4] Peter Baumann, Paula Furtado, Roland Ritsch, and Norbert Widmann, "The rasdaman approach to multidimensional database management," in *Proceedings of he 1997 ACM symposium on Applied computing.* ACM, 1997, pp. 166–173.

[5] P. Baumann: Array Databases and Raster Data Management. In*: T. Özsu, L. Liu (eds.): Encyclopedia of Database Systems,* Springer, Heidelberg, 2009

[6] D. Misev, P. Baumann: A Database Language More Suitable for the Earth System Sciences. *G. Lohmann et al (eds.): Towards an Interdisciplinary Approach in Earth System Science.* Springer 2015, doi:10.1007/978-3-319-13865-7

[7] E. Emili, C. Popp, M. Petitta, M. Riffler, S. Wunderle, and M. Zebisch*, "PM10 remote sensing from geostationary SEVIRI and polar-orbiting MODIS sensors over the complex terrain of the European Alpine region,"* Remote Sens. Environ., vol. 114, no. 11, pp. 2485–2499, 2010.

[8] E. Emili, C. Popp, S. Wunderle, M. Zebisch, and M. Petitta, *"Mapping particulate matter in alpine regions with satellite and ground-based measurements: An exploratory study for data assimilation,"* Atmos. Environ., vol. 45, no. 26, pp. 4344–4353, 2011.

[9] E. Emili, a. Lyapustin, Y. Wang, C. Popp, S. Korkin, M. Zebisch, S. Wunderle, and M. Petitta, *"High spatial resolution aerosol retrieval with MAIAC: Application to mountain regions*," J. Geophys. Res. Atmos., vol. 116, no. 23, pp. 1–12, 2011.

[10] Notarnicola, C.; Duguay, M.; Moelg, N.; Schellenberger, T.; Tetzlaff, A.; Monsorno, R.; Costa, A.; Steurer, C.; Zebisch, M. *Snow Cover Maps from MODIS Images at 250 m Resolution, Part 1: Algorithm Description.* Remote Sens. 2013, 5, 110-126.

[11] WagnerJ. E., Petitta M., *Atmospheric radiative transfer parametrization for solar energy yield calculations on buildings.* http://arxiv.org/abs/1506.06059

TOWARDS A FLEXIBLE INFRASTUCTURE FOR BIG DATA

C. González-Vila, A. Cabrera*, G. González-Felipe**

Research and Renewable Energies Institute of Tenerife (ITER)
Santa Cruz de Tenerife, Spain

ABSTRACT

Teide-HPC supercomputer is a high performance computing (HPC) infrastructure dedicated to research and development in various areas such as weather forecasting, bioinformatics, computing fluid dynamics, astrophysics with the Institute of Astrophysics of the Canary Islands[1] among others. The main use of the supercomputer is the standard computing cluster, launching batch jobs using a queue among a resource manager, which optimizes the usage of the infrastructure and its shared quotas. With the goal of providing new services by adding an abstraction layer, we have implemented a cloud solution based in the OpenNebula software. This solution allows to simplify the usage of our HPC infrastructure and gives more flexibility to the users of the infrastructure.

Index Terms— Cloud computing, IaaS, Storage Platform, Computing Platform

1. INTRODUCTION

TeideHPC has necessities of interoperability for different services using a high performance computing infrastructure. To deal with this problem, we have opted to use a cloud computing solution based on a free and open-source software solution, simplifying the requirements for users allowing them to have private clusters to use the resources we currently have. In this poster we present our view of the whole scenario, starting with the infrastructure specifications (section II) and the contribution of the cloud platform (section III) and we present the orchestration module, a self-provided management API that allows the transition between the computer profiles in the system (section IV).

2. INFRASTRUCTURE

TeideHPC supercomputer is composed of 1100 Fujitsu computer servers, with a total of 17800 computing cores and 36 TB of memory, a high-performance network and a parallel system of NetApp storage. The whole infrastructure is located in the D-ALiX datacenter [2].

2.1. Datacenter

D-ALiX is the datacenter where TeideHPC supercomputer is installed. It is located in the south coast of Tenerife (Canary Islands, Spain). The datacenter was built in 2012 with the highest levels of security, cooling and electrical availability, with TIER IV qualification in the electrical instrastructure and TIER III in cooling systems.

The datacenter has 2000 m² of technical ground, and it is designed to be modular so it can offer housing services at different levels: space in racks for servers, ground footprint, private cages and rooms.

D-ALiX is set up inside a reinforced industrial building, within the concept of “bunker inside bunker”. This approach, apart from reducing the necessary investments maintaining the demanded levels of quality, entails in a larger modularity as the datacenter has been designed to assist in a horizontal expansion rather than floor by floor. A full expansion to double the datacenter capacity would take only 6 months to be functional.

2.2. Conectivity

D-Alix datacenter is an internet Neutral Access Point (NAP), and the arriving point of submarine cables to the island, increasing the connectivity options.

Among the companies that form ITER group, CanaLink was established as a submarine cable operator, to give services to some of the most important communications submarine cables. Some examples are the ACE (Africa Coast to Europe), MST (Mainstreet Technologies de Main One Cable Company), SAT-3 (South Atlantic-3), Atlantis-2 or SAT-2 (South Atlantic-2). At the same time, it has been involved in some new projects, like the WACS (West Africa Cable System). All these infrastructures support communication connectivity between points as far as South Africa, Argentina, Brasil, Portugal or the United Kingdom. [2]

CanaLink is also the owner of the Canalink Cable (1800 km) that connects La Palma, Tenerife and Gran Canaria islands (Canary Islands) with Morocco and the European continent. This cable is formed by 4 optical fiber pairs, with a maximum aggregated bandwidth of 5.12 Tbps. The maximum available bandwidth depends on the number of lambdas equipped on every fiber pair, with a maximum of 128 · 10 Gbps [3].

Additionally, the communication infrastructure allows ITER to connect to the Spanish academic and research network that provides advanced communication services to the scientific community and national universities (RedIRIS), such as the observatory located in Roque de los Muchachos in La Palma. RedIRIS, at the same time, is connected to the pan-European data network for the research and education community (GÉANT). The availability of this connectivity, facilitated a cooperative project between the Supercomputer Center of Galicia (CESGA)[4] and ITER, some tests in the interoperability field were done, sharing computation and storage resources through a dedicated connection. The tests involved access to remote storage, computing containers and cloud bursting. Some of this studies are still running with other entities like the Port d'Informació Científica (PIC)[5] for the use of computation and storage resources.

2.3. Storage Housing

The great computing resources accessible within the TeideHPC supercomputer is completed with the availability of housing plans in the datacenter, which permits high performance access between both infrastructures, and the rest of the world through the network connections mentioned before.

A cluster storage may be installed in the datacenter using the available space, so it can be accessed directly from the computing resources connecting this storage to the cluster storage network with one or more optic fiber cables. This connection can reach higher aggregated bandwidths and low latencies due to the direct connection to the storage network switch.

2.4. Green Power Supply

The high performance computing environment is affected by a power consumption variability which is higher than the standard server environments. Due to the different characteristics of the computation jobs, the total power consumption for the whole infrastructure will be directly proportional to the running processes and their resource usage in terms of CPU, memory, network and I/O.

The energy storage systems, such as batteries, can offer economical benefits to infrastructures where the energy rates depend on the maximum power specified in the contract between the provider and datacenter itself [6]. Our case has a mixed clause, where the price depends on both the consumed energy and the maximum peak power reached, which has a strong economic penalization when surpassed.

In this sense, ITER, in their area of renewable energy development, has experience in the implementation of wind and photovoltaic (PV) generation systems, including self-developed high-power inverters (up to 150 KW) ([7], [8]). Furthermore, ITER is developing studies in the field of Smart Grids, implementing a project for an smart energy solution in small-size villages with photovoltaic distributed generation [9].

Connecting the Datacenter to an Smart Grid will allow the use of heterogeneous, distributed and discontinuous sources of energy, and therefore reducing the additional costs for exceeding the contracted power limit.

Currently, a part of the photovoltaic park of ITER is located in the ceiling surface of the datacenter, with an installed capacity of 400 kWp. This allows the Datacenter to meet 300 KW of base power demand during the day and avoid penalties for over consumption peaks during heavy work loads of HPC tasks, which can exceed 550 KW, 50 KW over the 500 KW contracted. Having local energy storage, such as banks of batteries, and an intelligent energy management [9], will let us to deal with such penalties, using the stored energy only when is necessary.

3. IAAS

Teide-HPC is the second biggest computational infrastructure taking account of computing power in Spain[10]. This capabilities allow the execution of different jobs at the same time, involving various knowledge areas. The access to the available resources can be done at different levels, giving specific functionalities to each user. The most low level usage to the user allows him to tweak the operating system, that runs over a virtualized environment, while in other options, like the high throughput executions over the batch queue, give the user a layer of abstraction.

As said before, Teide-HPC cloud service has allowed to provide access to resources in a Infrastructure as a Service

doi: 10.2788/854791

(IaaS) model. Some use cases that used this tool would be the implementation of a render farm for cinema production or the resources assignation for the Grid Processing on Demand project of the European Space Agency. It's been possible to develop improvements in the high performance computing services already running, allowing users to postprocess the output of their executions to choose which outputs should be transferred, saving the user time invested in long data transfers.

3.1. Cloud Platform Management with OpenNebula

3.1.1. Cloud Management Platform

The Cloud Management Platforms (CMP) available in the market has a wide variety of software, both open source and privative. Between this options you can find some widely used solutions, as OpenStack or CloudStack, with great functionality levels and highly configurable, but with a low learning curve. As an example, there are more than 15 components in an OpenStack based cloud, and the core module OpenStack Nova has around 600 configuration parameters. This usually leads to the increase of integration costs, due to the support requirements for the IT team.

From the studied solutions, Eucalyptus[11] was bought by Hewlett Packard (HP) to create the product known as HPE Helion Eucalyptus. The principal advantage is its compatibility with the AWS API, which accelerates the start-up time for experienced users in the Amazon Services.

On the other hand OpenNebula resulted a simple and powerful solution, with low requirements for the set up. The initial installation of the software has only two components, the core and the web interface called SunStone. This configuration includes all the necessary components to define the storage for the virtual machines (VMs), which can be shared through NFS, a parallel filesystem like Ceph or other kind of storage, the network configuration, that can be based in OpenVSwitch, 802.1Q VLANs, VxLAN, and others and some other specifications. OpenNebula also includes features to interoperate with external cloud interfaces and has no requirement for a concrete hypervisor, currently supporting KVM, Xen, and VMWare.

The unified OpenNebula web portal, also helps managing storage, network and virtualization, and provides monitorization and security options to help managing multiple independent clusters, according to the specified configuration and policies.

3.1.2. Cloud service in TeideHPC

Teide-HPC cloud service is based on the open source software OpenNebula [12]. OpenNebula is a free and open-source cloud computing platform for managing heterogeneous distributed data center infrastructures. It allows to manage a virtual infrastructure to build clouds of IAAS.

The computational model in the cloud differs to the standard batch jobs using queues on its orientation. Batch queues are meant to maximize the usage of an infrastructure, while cloud services are oriented to improve the scalability of any service.
Deployed services are prepared to grow in number, depending on the requirements, achieving better performance, and to shrink when the peaks of demand are over.

Teide-HPC cloud service has allowed to provide access to resources in a Infrastructure as a Service (IaaS) model. Some use cases that used this tool would be the implementation of a render farm for cinema production or the resources assignation for the Grid Processing on Demand project of the European Space Agency. It's been possible to develop improvements in the high performance computing services already running, allowing users to postprocess the output of their executions to choose which outputs should be transferred, saving the user time invested in long data transfers.

The European Space Agency Grid Processing on Demand (G-POD) for Earth Observation Applications [13] is a grid environment that allows its users to execute compute intensive jobs in the resources assigned to the grid. Between this resources, there are assigned some resources from TeideHPC supercomputing, adding a total of 160 CPU cores and 640GB of RAM memory.
The previously mentioned render farm use case was carried out in an animation film (Caputer the flag) and it allowed in a simple way to use 400 virtual machines, instantiating a client provided virtual machine. The result cluster had assigned 12800 computing cores, 12.8 TB of RAM memory and more than 20 TB of hard drive storage.

One of the main advantages of the duality of TeideHPC, by using batch and cloud computing paradigms simultaneously is the capacity to deploy cloud services that interact with the traditional computing environment. This is the case of the

doi: 10.2788/854791

on-site post-processing tasks for some softwares. As an example, we have some appliances available to post-processing the results of computing fluid dynamics (CFD) jobs with XFlow or visualization tasks with Paraview.

3.2. Batch System

The batch queue system allows TeideHPC to execute high throughput jobs, but also can be adapted to the Big Data field. As previous work in the field demonstrated [14], it is possible to set up a processing engine like Spark or Hadoop in the queue system, suffering some negligible overhead, compared to the global computing resources utilized. This configuration allows the optimization of the available infrastructure, assigning priorities and limitations and avoiding job starvation.

4. API CONTROL - COMPUTING BURSTING ON DEMAND

One of the main characteristics of the cloud computing paradigm is the services flexibility and scalability, that allows a service to increase the number of instances when necessary. This functionality binds the number of cloud resources with the load of the system itself. To avoid underuse of this resources when the load decrease it is necessary to have tools that allow changing the profile of the infrastructure, assigning resources where they are needed, and doing it in a flexible and automatic manner.

With that target in mind, we developed an API library that facilitates the deployment of the infrastructure, over different profiles, doing iteratively all the configuration tasks needed to be running in the new profile, beginning with network configurations (for the managmente, data and Infiniband network), configuration management system (chef), authorization and the installation of the operating system..

The deployment of a HPC computing node guided by this API allows, in a matter of minutes, to get the system up and running as a cloud worker node. The use of configuration management tools has in mind security and privacy to manage encryption and authentication.

5. CONCLUSIONS

We are going towards a sustainable, scalable and secure datacenter in order to be ready for the future challenges of big data and high performance computation. To achieve so, we are unifying every aspect of the infrastructure, minimizing the energy costs by improving the storage and the usage of the energy produced using renewable sources. To adapt our workload in a more flexible way, we support our work using an IAAS approach, using cloud computing and cloud bursting when necessary. When cloud computing is not possible, we can perform a standard approach using a batch queue system for big data software with the appropriate priorities and limitations.

6. REFERENCES

[1] Institute of Astrophysics of the Canary Islands. http://www.iac.es.
[2] D-Alix Datacenter. http://www.d-alix.com
[3] Canalink cable press note. https://www.alcatel-lucent.com/press/2010/002064
[4] Fundación Pública Galega Centro Tecnolóxico de Supercomputación de Galicia (CESGA). http://cesga.es.
[5] Port d'Informació Científica (PIC). http://www.pic.es.
[6] Darshan S. Palasamudram, Ramesh K. Sitaraman, Bhuvan Urgaonkar, Rahul Urgaonkar. Using Batteries to Reduce the Power Costs of Internet-scale Distributed Networks.
[7] A. Linares, J. Fernández, M. Friend, E. Llarena, C. Montes, G. Moncho, N. Losada, O. González, D. Molina, A. Pío, M. Cendagorta. Voltage Dips: Verification, Validation and Certification Procedure for PV Installations
[8] E. Llarena, C. Montes, A. Linares, D. Molina, A. Pío, O. González, L. López-Manzanares, J. Fernández, J. Rodríguez, M. Friend, M. Cendagorta. A localized weighted efficiency for inverters
[9] L. Marechal, J. Fernández, G. Moncho, C. Montes, M. Delgado, J. Rodríguez, M. Friend, M. Cendagorta. Building a PV based microgrid as a smart energy solution for a small-sized touristic village.
[10] Top500 list. http://top500.org
[11] The HP implementation of Eucalyptus. http://www8.hp.com/us/en/cloud/helion-eucalyptus-overview.html
[12] OpenNebula. http://opennebula.org.
[13] ESA Grid Processing On Demand. https://gpod.eo.esa.int/
[14] Ruben Tous, Anastasios Gounaris, Carlos Tripiana, Jordi Torres, Eduard Ayguadé, Jesús Labarta, Yolanda Becerra, David Carrera, Mateo Valero. Spark Deployment and Performance Evaluation on the MareNostrum Supercomputer

A FOOTPRINT DATABASE OF THE HERSCHEL SPACE OBSERVATORY

Erika Varga-Verebélyi[1], *László Dobos*[2], *Eva Verdugo*[3], *David Teyssier*[3], *Katrina Exter*[3], *Ivan Valtchanov*[3], *Tamás Budavári*[4] *and Csaba Kiss*[1]

[1]Konkoly Observatory, Research Centre for Astronomy and Earth Sciences, Budapest, Hungary
[2]Eötvös Loránd University, Dept. of Physics of Complex Systems, Budapest, Hungary
[3]European Space Astronomy Centre, European Space Agency
[4]The Johns Hopkins University, Department of Applied Mathematics and Statistics, Baltimore, USA

ABSTRACT

We present a searchable database of the exact geometric description of footprints of observations made by ESA's far-infrared satellite, the Herschel Space Observatory [1]. Herschel operated in numerous observational modes including single pointed observations, raster maps and scan maps. Maps have been reconstructed from basic pointing information and, as opposed to pixellated maps, are represented in a simple geometric format. Footprints are indexed using Hierarchical Triangular Mesh (HTM) [2], which supports finding Herschel observations by many ways: simple point coverage, cone search and intersection search. The footprint database is available via a RESTful web service as well as via a traditional web site at `http://herschel.vo.elte.hu/`.

Index Terms— sky coverage, footprint, spatial indexing, astronomical databases, Virtual Observatory

1. INTRODUCTION

Knowing the exact sky coverage of systematic sky surveys is of high importance for statistical reasons. Non-survey telescopes, on the other hand, often cover an area comparable to large surveys during their years of operation. By uniformly reducing the individual observations, source catalogs are being created, e.g. the Hubble Source Catalog [3]. The footprints of non-survey catalogs are sometimes very complex compared to survey footprints and require correct handling of all instrumental observational modes, detector footprints and telescope pointing modes.

In this work we present our method of reconstructing the footprints of more than 47,000 Herschel Space Observatory (HSO) observations. The Herschel Space Observatory, an ESA cornerstone mission, flew between 2009 and 2013. The 3.5 m telescope was deployed at L2 and made observations in the far-infrared and submillimeter bands with its three instruments: the HIFI heterodyne spectrometer [4], and the PACS and SPIRE imaging photometers/integral field spectrometers [5, 6]. During Herschel's four years of operation it successfully observed more than 20,000 square degrees of the sky with varying spatial resolution and depth. Herschel data are available via the Herschel Science Archive (HSA) which is being complemented with exact sky coverage information by this work.

The research was supported by the Hungarian Scientific Research Fund OTKA NN 103244 and OTKA NN 114560 grants and PECS contract 4000109997/13/NL/KML of the ESA. *Herschel* is an ESA space observatory with science instruments provided by European-led Principal Investigator consortia and with important participation from NASA.

2. SKY COVERAGE OF HERSCHEL OBSERVATIONS

The sky coverage (footprint) of the observations was reconstructed from raw pointing information extracted from the HSA via the Herschel Interactive Processing Environment (HIPE). We built a parallel tool to pre-process and ingest pointing data into a database and all subsequent process-

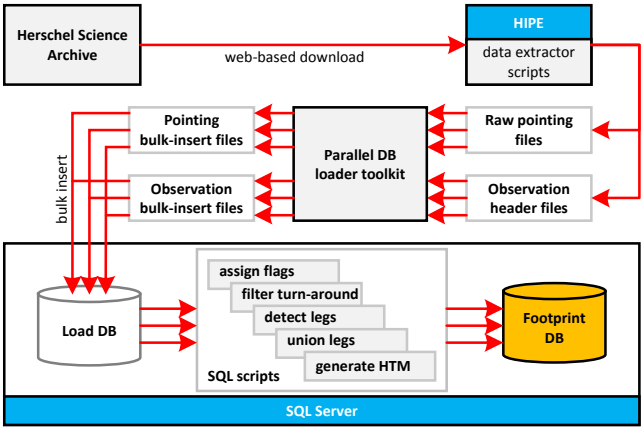

Fig. 1. The footprint generation process. First raw pointing data is downloaded from the HSA using HIPE, then it is ingested into a relational database. All subsequent processing is done inside the database using user-defined functions. Multiple arrows represent parallel processing.

doi: 10.2788/854791

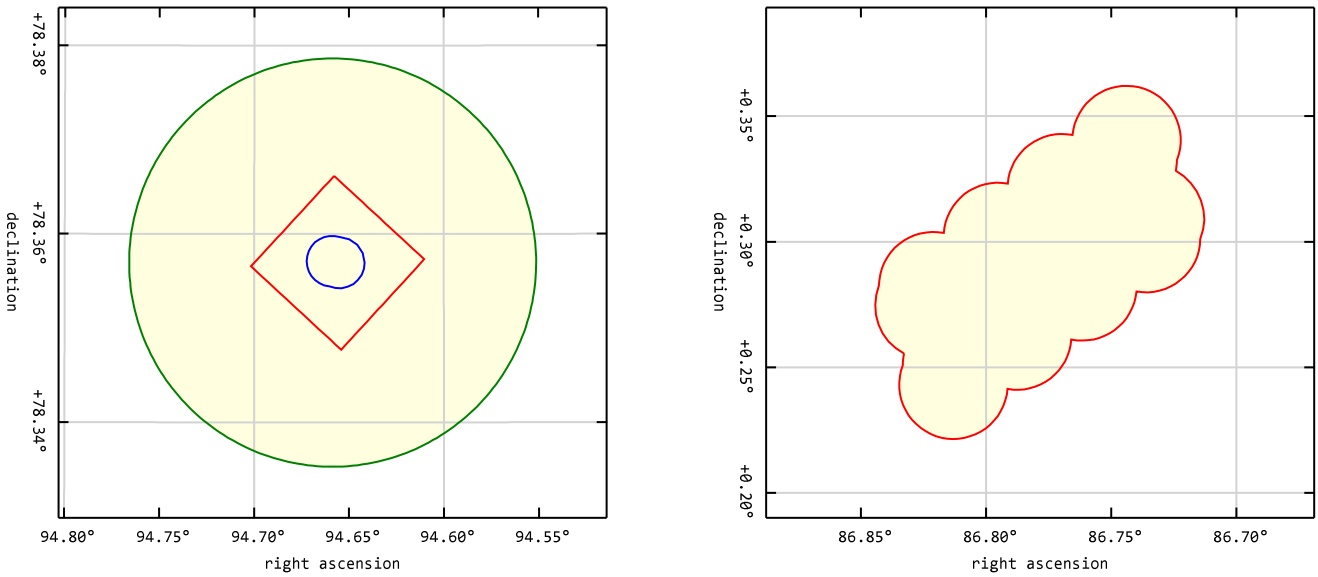

Fig. 2. Left panel: Illustration of the instantaneous sky coverage of the three spectrometers. The green line shows the footprint of the SPIRE integral field spectrometer which is represented by a circle, $2\overset{'}{.}6$ in diameter, equivalent to the vignetted coverage of the detector. The red line is the approximately $47'' \times 47''$ field of view of the PACS integral field spectrometer. (In reality, one of the rows of the 5×5 matrix of PACS is slightly offset.) The blue line marks the field of view of the HIFI instrument, which is represented as the union of two circles corresponding to the half-power beam widths of the two polarizations. HPBW of HIFI depends strongly on frequency between $11\overset{''}{.}1$ and $43\overset{''}{.}5$. Right panel: The footprint of a SPIRE spectroscopic raster map.

ing took place within the database server using user-defined functions to minimize data movement. Although this was not exactly a big data project, yet the raw pointing information amounts to almost 100 GB while the final data product is around 10 GB. The data processing pipe-line is illustrated in Fig. 1.

Herschel operated in numerous observational modes including single pointed observations, raster maps and scan maps. During single pointed mode observations, which mode was used primarily for spectroscopy, the telescope was aimed at a given position or it was alternating between the source and the calibration point. Footprints of pointed observations can simply be represented by the sky coverage of the detector of the integral field spectrometers (PACS and SPIRE) or the half-power beam width (HIFI). Fig. 2 illustrates the instantaneous sky coverage of the three spectrometers. Raster maps are basically $N \times M$ pointed observations organized in a grid. This observational mode was used mainly for spectroscopic measurements.

Almost all photometric observation by the PACS and SPIRE instruments were done using the scan map mode. Since the instantaneous sky coverage of the detectors is only $3\overset{'}{.}5 \times 1\overset{'}{.}75$ (PACS) and $4' \times 8'$ (SPIRE), extended areas were observed by directing the boresight of the telescope along straight, parallel trajectories (the so called scan legs) while reading out the imaging detectors continuously. Sometimes the areas were re-scanned in a pattern almost perpendicular to the first. Since the velocity of scanning was fixed to a few speeds, often multiple repetitions of the same area were done in order to increase the signal to noise ratio and depth of the observation.

Raw pointing data consisted of the coordinates, an instrument position angle and a time stamp, along with a few flags. At the end of each straight scan leg the telescope was slewed to the next leg. Since PACS raw pointing data contained the slew periods, telescope turn-arounds had to be identified and filtered out. Fig. 2 illustrates the method of scan map reconstruction. Once the ends of the scan legs had been identified, the footprints of individual legs were taken to be the convex hull of the detector corners, positioned at the beginning and at the end of the scan leg. The final footprints of scan maps were created by taking the union of the scan legs.

3. GEOMETRIC REPRESENTATION OF FOOTPRINTS

To represent the geometry of Herschel footprints we used the Spherical Library of Budavári et al [7]. Instead of pixelating the sphere, the software library describe spherical geometry in an exact analytic form. *Regions* (complex spherical shapes) are described as unions of *convexes*. A convex is the intersection of the surface of the unit sphere with a 3-dimensional

doi: 10.2788/854791

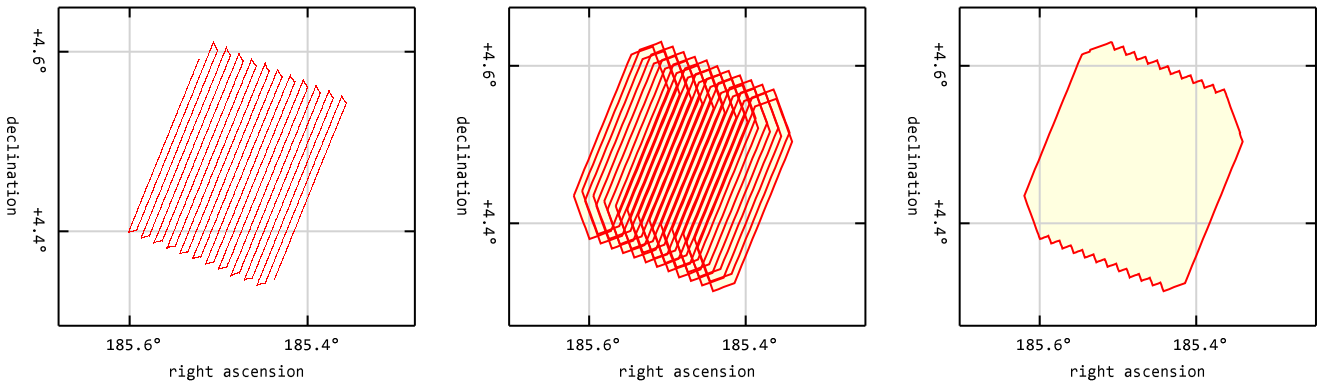

Fig. 3. Construction of a Herschel PACS scan map footprint. The left panel shows the original pointing information without turn-around filtering. The middle panel shows the overlapping, parallel scan legs. The right panel is the final footprint, the union of the scan legs.

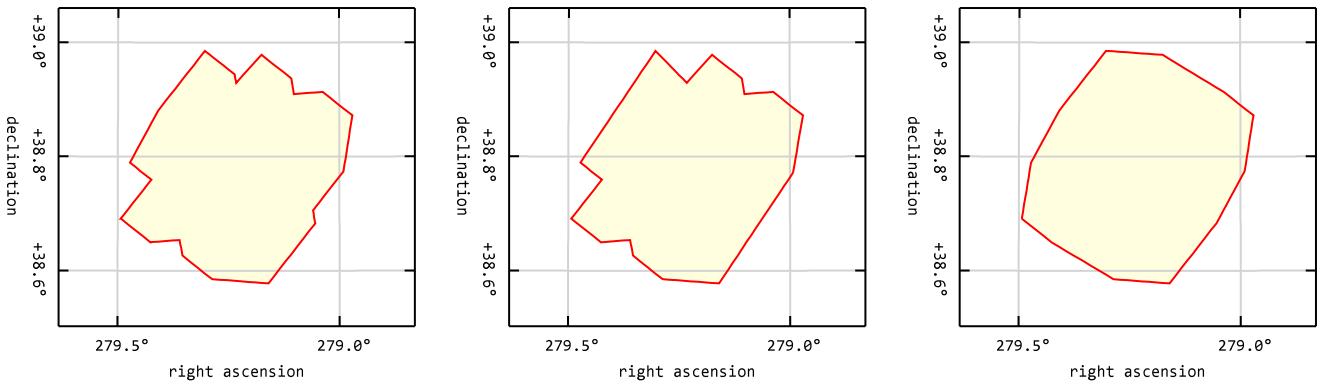

Fig. 4. Illustration of the convex hull reduction (middle panel) and the Douglas–Peucker algorithm (right panel) applied to the footprint of a SPIRE scan map. The Douglas–Peucker method can reduce the complexity of the outline of concave regions without removing the larger infolds, at the time it might cut off small tips and potentially reduce the area of the shape.

convex polyhedron, i.e. a set of *half spaces*. Hence, spherical convexes are bounded by arcs of great or small circles and are not necessarily contiguous. For example, the library can easily represent regions with non-trivial topology, such as rings. The main advantage of decomposition regions into convexes is that it makes point containment testing very simple and fast, an important feature when implementing search algorithms for large databases.

The aforementioned Spherical Library supports all necessary operations required by Herschel footprints, including boolean set operations and computation of the spherical convex hull of a list of points. We have extended the functionality of the library with two new features. The first is the computation of the great-circle-arc-bounded convex hull of entire footprints, the second is the Douglas–Peucker simplification of complex region [8]. Fig. 4 illustrates the results of the two methods.

In addition, we have implemented a fast and flexible visualization library to generate nice plots of footprints. The library supports a wide set of projections, various file formats including vector formats, automatic rotate and zoom functionality and streaming data retrieval from the database to plot large data sets. The figures of the paper were generated with this library.

4. INDEXING THE FOOTPRINTS

Herschel made tens of thousands of observations and the footprints are considerably complex. We used the Hierarchical Triangular Mesh indexing scheme to compute an adaptive triangular cover of the footprints and built a hierarchical database index to allow fast spatial search. HTM indexing is explained in details in [7]. Fig. 5 illustrates the HTM cover of the footprint of a SPIRE observation. When a complex shape is covered with HTM triangles (trixels), thanks to the hierarchical construction of the mesh, the inner, completely covered triangles can be merged into larger ones. Point containment testing with HTM is very simple. If a point falls into

doi: 10.2788/854791

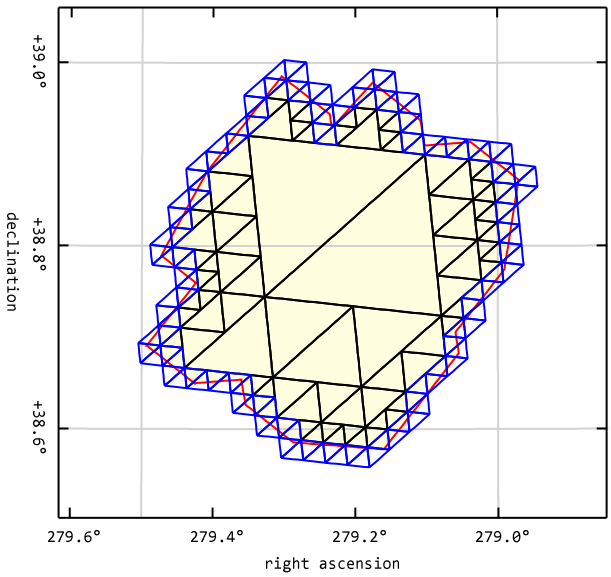

Fig. 5. Illustration of footprints of the three kind of pointed observations. HIFI: the diamater of the two circles varies between 11.1 and 43.5 arcsec, depending from the band used, PACS: 47 × 47 arcsec square, SPIRE: one circele with the diameter of 120 arcsec.

an inner triangle, one can always be sure it is within the shape. Shall the point fall into a boundary trixel, exact containment testing, a more expensive procedure, is necessary.

In practice, one generally wants to find all observations containing an object or covering (or intersecting with) a given region on the sky. Since the original implementation allowed for point containment search only, we implemented two major extensions to HTM indexing to support the Herschel Footprint Database.

5. REFERENCES

[1] Pilbratt et al., "Herschel space observatory-an esa facility for far-infrared and submillimetre astronomy," *Astronomy & Astrophysics*, vol. 518, pp. L1, 2010.

[2] P. Z. Kunszt, A. S. Szalay, I. Csabai, and A. R. Thakar, "The Indexing of the SDSS Science Archive," in *Astronomical Data Analysis Software and Systems IX*, N. Manset, C. Veillet, and D. Crabtree, Eds., 2000, vol. 216 of *Astronomical Society of the Pacific Conference Series*, p. 141.

[3] S. Lubow and T. Budavári, "Hubble source catalog," in *Astronomical Data Analysis Software and Systems XXII*, 2013, vol. 475, p. 203.

[4] de Graauw et al., "The Herschel-Heterodyne Instrument for the Far-Infrared (HIFI)," A&A, vol. 518, pp. L6, July 2010.

[5] Poglitsch et al., "The Photodetector Array Camera and Spectrometer (PACS) on the Herschel Space Observatory," A&A, vol. 518, pp. L2, July 2010.

[6] Griffin et al., "The Herschel-SPIRE instrument and its in-flight performance," A&A, vol. 518, pp. L3, July 2010.

[7] T. Budavári, A. S. Szalay, and G. Fekete, "Searchable sky coverage of astronomical observations: Footprints and exposures," *Publications of the Astronomical Society of the Pacific*, vol. 122, no. 897, pp. 1375–1388, 2010.

[8] D. H. Douglas and T. K. Peucker, "Algorithms for the reduction of the number of points required to represent a digitized line or its caricature," *Cartographica: The International Journal for Geographic Information and Geovisualization*, vol. 10, no. 2, pp. 112–122, 1973.

doi: 10.2788/854791

MULTIBILLION SCALE CATALOGUES IN THE GAIA ARCHIVE DATABASE AT ESAC

J. Gonzalez-Nunez **, J. Salgado*, R. Gutierrez-Sanchez*, J.C. Segovia Serrato*, J. Hernandez*, C. Arviset*, F. Aguado-Agelet**, B. Merin*, W. O'Mullane**

* ESAC Science Data Centre - European Space Agency ** ETSE Telecomunicacion - Universidade de Vigo

ABSTRACT

The ESA Gaia mission will survey the sky for at least 5 years measuring high accuracy astrometry, radial velocities and multi-colour photometry. The Data Analysis and Processing Consortium (DPAC) efforts will result in an astronomical catalogue with unprecedented accuracy and completeness of up to 1 billion (1E9) sources, about 1 percent of the Galactic stellar population. Efficient scientific exploitation of this data set will require the development of software services for the storage, access, retrieval and mining of this data. Specific developments applied to provide exposure to such a large dataset include, among others, relational databases able to cope with tables in the range of terabytes.

In the heart of the Gaia Archive, the Core Systems developed at ESAC (Madrid) will use, among other database technologies, PostgreSQL instances. Modules such as pgSphere and Q3C provide PostgreSQL capabilities for the storage, geometrical query and crossmatch of astronomical catalogues. Its utilisation in combination with Virtual Observatory (VO) protocols such as TAP has become a widespread practice in Astronomical Data Centres for serving their catalogues to the scientific community and general public. We will review the architecture, technology, computing models and configuration applied to this infrastructure, as well as details about the administration of this large instance.

Index Terms— Gaia, Relational databases, Scientific archives, Virtual Observatory

1. INTRODUCTION

The ESDC (ESAC Science Data Centre), located at ESAC, is the responsible for the design and implementation of the science astrophysical, planetary and heliospheric missions archives. This group is also responsible for the long term preservation of the data.

Gaia is part of the European Space Agency (ESA) Horizon 2000 long-term scientific program (1). Launched in late 2013, was inserted into a Lissajous orbit around the Lagrange point L2. Spinning for at least five years, it will provide a map of our Galaxy including comprehensive positional and radial velocity measurements that will allow the creation of a three-dimensional map of up covering about 1 percent of the Galactic stellar population.

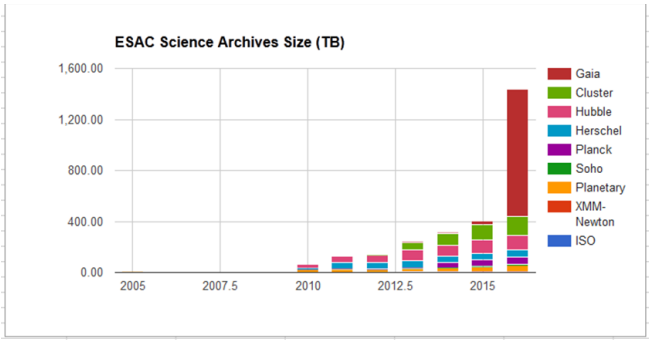

Fig. 1. Expected evolution of data size inside ESAC Science Archives

Data Analysis and Processing Consortium (DPAC) is the consortium created for processing and distributing the data Gaia . DPAC is organized into nine Coordination Units, each responsible for a well-defined set of tasks in the data processing. Gaia is the first ESA astronomy mission for which Archive developent will be part of a wider data analysis consortium, coordinated under Coordination Unit 9 (CU9). Its responsibilities, organization and structure are defined in (2).

Out of the data processing carried out by DPAC, CU9 will receive incrementally more precise releases of a catalogue comprising around one billion (1E9) sources, among other data products.

Along with the Gaia catalogue the Gaia Archive will deliver other DPAC produced catalogues such as the simulated Gaia Universe Snapshot Model (GUMS) or the Initial Gaia Source List (IGSL), with the original catalogues used for its production, including all catalogue-related DPAC products, such as the official crossmatches.

The Gaia Archive Catalogue Database will host all of the different releases of these catalogues, with special emphasis in delivering the greatest performance for the latest versions.

The final data delivery will not only include the catalogue of one billion sources but also the single epoch CCD transit data that was used in its computation, and, in addition to catalogues, the consortium will also deliver an estimated 12TB of science telemetry, 1.5E9 spectra and a full database content of 44TB of data still under reduction process. This will make an

doi: 10.2788/854791

Type of Data	Amount	
Science telemetry	17TB	
Astrometry transits	$22.5x10^9$	$225x10^9$ images
Photometry transits	$22.5x10^9$	$45x10^9$ images
Spectroscopy transits	$1.5x10^9$	$4.5x10^9$ spectra
Main Database	44TB	

Table 1. Gaia data statistics up to 5th June 2015

estimated total around 1 PB at the end of the mission (2022)

In addition to Gaia data, the Gaia archive will host other main catalogues in astronomy to maximise scientific exploitation, supporting science cases requiring for comparison or inclusion of data from these catalogues.

Being large for a space mission, Gaia numbers will be exceeded by future ESA missions like Euclid. Architectural patterns for astronomical archives will have to change to maximise the exploitation of large volumes of data where highly computationally intensive statistical analysis take relevance over human supervised analysis.

2. EXPOSING ASTRONOMICAL CATALOGUES

Storage and query of tables of data has long been the field of study of Relational Database Management System (RDBMS). This mature technology, based in relational algebra and ACID priciples (Atomicity, Consistency, Isolation, Durability) (3) has been applied with success to the storage of astronomical catalogues, particularly within astronomical surveys, producing query services such as SkyServer and CasJobs (4). The Gaia Archive provides similar mechanisms, bringing the user a query interface similar to database administration tools as shown in Figure 2.

Standardisation, through VO protocols such as TAP (5) and the usage of standard query languages as ADQL (6) have dramatically increased the interoperability of these kind of services. Implementations such as Taplib and ready to use toolkits (7) (8) reduce the entry cost barrier for data centres to expose their catalogues through their own services.

The Gaia Archive architecture is based in TAP+, this is, and extension of TAP with the addition of certain capabilities (9) (10):

- **Authentication:** Based in the ESA CAS server, allowing for interoperation with other ESA services and the ESA VOSpace area.
- **User DB areas:** local schema where user catalogues can be uploaded and results of queries can be stored persistently.
- **Authorisation:** Imposed by Gaia validation procedures, the system allows to manage access to certain items of data.
- **Sharing capabilities:** Tables can be shared to other users and that sharing mechanism will be notified by events to the different users.
- **Basic crossmatch support:** Possibility to create a close neighbour crossmatch between user tables and Gaia public schema tables.

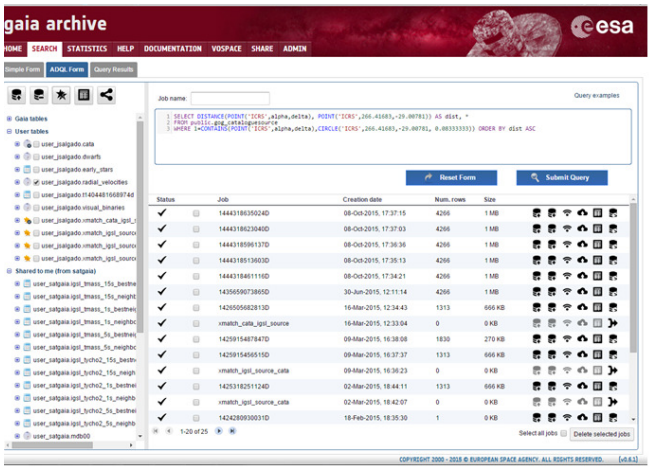

Fig. 2. Gaia Archive at ESAC UI: TAP query interface

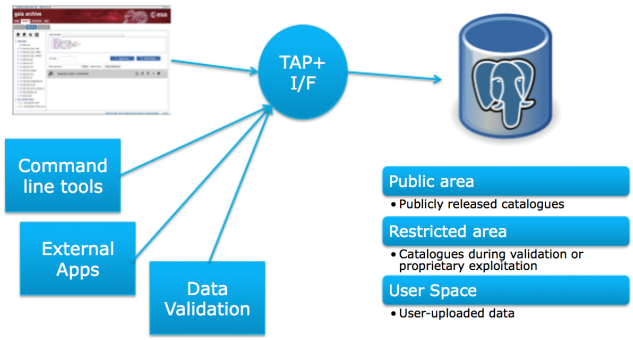

Fig. 3. Access to catalogues through TAP+

3. DATABASE TECHNOLOGIES FOR CATALOGUE HANDLING IN TAP SERVICES

3.1. Early technological studies

As part of Gaia CU9 and within the early Gaia Archive Preparation group (GAP) activities, several technological studies were conducted to address suitability of different technologies for handling large catalogues, being the most representative:

- **Java based Map-Reduce analysis on a Hadoop cluster (11):** Representative of the family of shared-nothing cluster systems based in MapReduce and distributed filesystems (GFS, HDFS).

doi: 10.2788/854791

Extremely scalable infrastructure (clusterized). Shared nothing architecture allows for aggregations of nodes virtually unlimited, also in cloud infrastructures. Allows for the execution of Machine Learning and statistical analysis for large-scale data processing, with derived technologies such as Apache Spark The absence of indexing schemas limits response times to be workable for interactive work. Being a NO-SQL platform, availability of SQL querying mechanisms is limited to implement ADQL and TAP access on top of it.

- **Greenplum Relational clustered database based in PostgreSQL:** Representative of the family of commercial relational databases with multiple nodes.

  Very scalable. Table partitioning, or clustering strategy sets limits in the types of queries that may be performed efficiently. Specific and complex strategies have to be put in place to perform efficient crossmatches based in Q3C join. Only proprietary technologies are available, technology dependence, licensing costs, very long term maintainability issues.

- **Standalone Relational PostgreSQL database with high performance hardware:** Open source and free. Easy maintenance. Full general purpose capability (no partitioning or clustering needed). Very good performance through intra-machine query parallelization. Vertical scalability, limited to present hardware, but with a fair margin over Gaia sized catalogues (1E9 sources).

3.2. The Gaia Archive Relational DB

The main database for handling catalogues in the Gaia Archive is based in PostgreSQL (12), and the astronomical geometry extensions PgSphere (13) and Q3C (14). The deployment architecture is described in Figure 4, and has the following features:

- Load Balancing: Scalability is ensured enabling streaming replication between the cluster master and slave node(s). As per the time of the first release, the Gaia Archive Database will be comprised of one master and one slave nodes. Load is balanced through pgpool-II, dispatching queries between DB servers.

- High Availability and Hot Standby: Combining the hot-standby capabilities of PostgreSQL with the watchdog and failover features of Pgpool-II, the DB system can ensure much higher uptime rates, preventing downtimes caused by HW failure of one of the nodes.

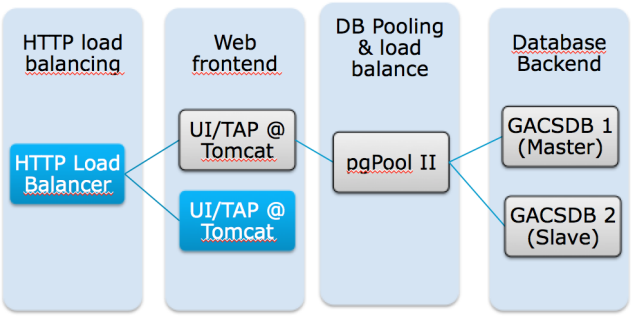

Fig. 4. Gaia Archive Database deployment architecture. Grey: comissioned as of 11th February 2016

Radius	Transactions/s	avg service time	output
1´	9430	2.5 ms	160 rows
5´	451	54 ms	4733 rows
10 deg	5.8	3.9 s	720000 rows

Table 2. GACSDB01 throughput querying cone searches to GUMS Milky Way simulation with 2.2E9 sources

3.3. Throughput and performance

Query throughput is maximised through the usage of high performance hardware, including high amounts of memory (up to 1TB on GACSDB01 and 1.5TB on GACSDB02) and the layering of storage volumes according to the usage patterns required to the data hosted (Figure 5). Performance has been addressed with stress testing on GUMS (15) data as shown in Table 2.

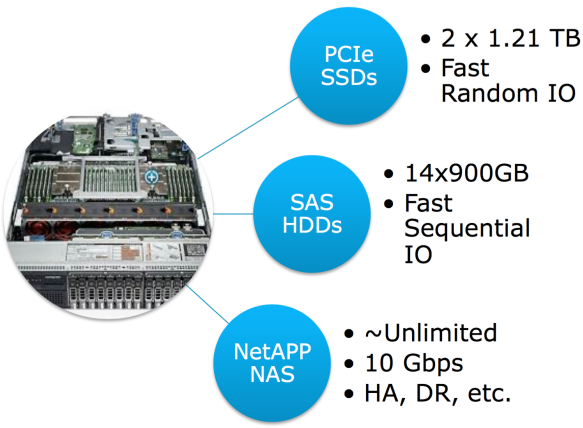

Fig. 5. Storage volumes layering for GACSDB01

4. ACKNOWLEDGEMENTS

This work has been performed at the ESAC Science Data Centre (ESDC), coordinated as part of the Gaia Coordination

Unit 9 of the Data Processing and Analysis Consortium for Gaia (DPAC).

We acknowledge CU2 for the production of independently simulated Gaia-like data for use in the system. The data simulations have been done in the supercomputer Mare Nostrum at Barcelona Supercomputing Center – Centro Nacional de Supercomputación (The Spanish National Supercomputing Center).

References

[1] M. A. C. Perryman, K. S. de Boer, G. Gilmore, E. Hg, M. G. Lattanzi, L. Lindegren, X. Luri, F. Mignard, O. Pace, and P. T. de Zeeuw, "Gaia: Composition, formation and evolution of the galaxy," *A&A*, vol. 369, no. 1, pp. 339–363, 2001.

[2] F Mignard, R Drimmel, F De Angeli, C Babusiaux, C.A.L. Bailer-Jones, U Bastian, M Beck, M Biermann, A.G.A. Brown, J de Bruijne, C Cacciari, A Cellino, J.F Claeskens, M Cropper, F Drimmel, P Dubath, D.W. Evans, L Eyer, C Fabricius, M Gai, C Huc, S Jordan, C Jordi, D Katz, S Klioner, M Kontizas, C Lagerkvist, U Lammers, M Lattanzi, F van Leeuwen, T Levoir, L Lindegren, X Luri, E. Masana, F Mignard, W O Mullane, X Passot, D Pourbaix, A Recio-Blanco, A Robin, S Rosen, C Soubiran, P Tanga, F Thevenin, J Torra, A Vallenari, A Volpicelli, and J Warell, "Gaia data processing and analysis consortium, response to esa announcement of opportunity. proposal for the gaia data processing," in *Gaia DPAC documentation*, 2007.

[3] Jim Gray, "The transaction concept: Virtues and limitations (invited paper)," in *Proceedings of the Seventh International Conference on Very Large Data Bases - Volume 7*. 1981, VLDB '81, pp. 144–154, VLDB Endowment.

[4] A. Szalay, J. Gray, A. Thakar, P. Z. Kunszt, T. Malik, J. Raddick, C. Stoughton, and J. vandenBerg, "The SDSS SkyServer, Public Access to the Sloan Digital Sky Server Data," *eprint arXiv:cs/0111015*, Nov. 2001.

[5] P. Dowler, G. Rixon, and D. Tody, "IVOA Recommendation: Table Access Protocol Version 1.0," *ArXiv e-prints*, Oct. 2011.

[6] I. Ortiz, J. Lusted, P. Dowler, A. Szalay, Y. Shirasaki, M. A. Nieto-Santisteban, M. Ohishi, W. O'Mullane, P. Osuna, the VOQL-TEG Working Group, and the VOQL Working Group, "IVOA Recommendation: IVOA Astronomical Data Query Language Version 2.00," *ArXiv e-prints*, Oct. 2011.

[7] L. Michel, C. Motch, H. N. Nguyen, and F. X. Pineau, "Building an archive with Saada," *Astronomy and Computing*, vol. 7, pp. 45–51, Nov. 2014.

[8] M. Demleitner, "The DaCHS Multi-protocol VO Server," in *Astronomical Data Analysis Software and Systems XXIII*, N. Manset and P. Forshay, Eds., May 2014, vol. 485 of *Astronomical Society of the Pacific Conference Series*, p. 309.

[9] J Salgado, J. Gonzalez, R. Gutierrez, JC. Segovia, P. Osuna, and C. Arviset, "Gaia Archive Implementation. Waiting for one billion stars," in *Astronomical Data Analysis Software and Systems XXIV*, 2014, Astronomical Society of the Pacific Conference Series.

[10] J Salgado, J González-Núez, R Gutiérrez Sánchez, JC Segovia, C Arviset, S Nieto, and B. Merín, "Access to massive catalogues in the Gaia archive: a new paradigm," in *Astronomical Data Analysis Software and Systems XXV*, 2015, Astronomical Society of the Pacific Conference Series.

[11] D. Tapiador, W. OMullane, A.G.A. Brown, X. Luri, E. Huedo, and P. Osuna, "A framework for building hypercubes using mapreduce," *Computer Physics Communications*, vol. 185, no. 5, pp. 1429 – 1438, 2014.

[12] Michael Stonebraker and Lawrence A. Rowe, "The design of postgres," *SIGMOD Rec.*, vol. 15, no. 2, pp. 340–355, June 1986.

[13] I. Chilingarian, O. Bartunov, J. Richter, and T. Sigaev, "PostgreSQL: the Suitable DBMS Solution for Astronomy and Astrophysics," in *Astronomical Data Analysis Software and Systems (ADASS) XIII*, F. Ochsenbein, M. G. Allen, and D. Egret, Eds., July 2004, vol. 314 of *Astronomical Society of the Pacific Conference Series*, p. 225.

[14] S. Koposov and O. Bartunov, "Q3C, Quad Tree Cube – The new Sky-indexing Concept for Huge Astronomical Catalogues and its Realization for Main Astronomical Queries (Cone Search and Xmatch) in Open Source Database PostgreSQL," in *Astronomical Data Analysis Software and Systems XV*, C. Gabriel, C. Arviset, D. Ponz, and S. Enrique, Eds., July 2006, vol. 351 of *Astronomical Society of the Pacific Conference Series*, p. 735.

[15] Y. Isasi, F. Figueras, X. Luri, and A.C. Robin, "Gums and gog: Simulating the universe for gaia," in *Highlights of Spanish Astrophysics V*, Jose M. Diego, Luis J. Goicoechea, J. Ignacio Gonzlez-Serrano, and Javier Gorgas, Eds., Astrophysics and Space Science Proceedings, pp. 415–415. Springer Berlin Heidelberg, 2010.

THE BENEFITS OF USING XML TECHNOLOGIES IN ASTRONOMICAL DATA RETRIEVAL AND INTERPRETATION

Guy Beech
School of Computing and Engineering, University of Huddersfield, HD1 3DH, UK

ABSTRACT

This paper describes a solution found during recent research that could provide improvements in the efficiency, reliability and cost of retrieving stored astronomical data. This solution uses XML Technologies in showing that when querying a variety of astronomical data sources a standardised data structure can be output into an XML query results Document. This paper shows the astronomical XMLSchema that has been partially developed in conjunction with simple custom supporting system software. It also discusses briefly possible future implications.

Index Terms - XML Schema, standard, data processing

1. INTRODUCTION

The current retrieval processes of existing astronomical records have a number of difficulties that are shown in the table below:

Multiplicity of formats
Transformation problems due to incompatible formats
Archaic problems with vintage data such as limited storage space, maximum string length and other restrictions of technology of their time
Need for digitisation
The increasing rate of astronomical data capture both by institutions and amateurs
The growing different types of astronomical data collectors [1]

Table 1: Retrieval Difficulties

If these could be solved it would result in more reliable and cost effective acquisition of useful and valuable data.

Although there are many institutions and individuals using good standards of astronomical data capture, there is a lack of any widely adopted comprehensive standards of astronomical data recording that are currently in use. Due to this variety of storage systems data retrieval is often expensive and incomplete. [2]

The purpose of the research carried out was to:

- Identify a precise way forward to achieving improved efficiency of data retrieval, through gaining understanding of current literature, developing a new retrieval process and working towards an astronomical schema to be adopted as standard. [3]

The literature review of this research confirmed that the great majority of all records are now held in digital format [4] in electronic databases, in a wide variety of data structures and that there has also been in recent years an extremely rapid increase in the rate at which new data is being saved, which is a trend that is predicted to carry on into the future. [5]

The literature review also indicated that there is no widely used, comprehensive standard of data recording used in astronomy. Schemas produced so far, including those of the Virtual Observatory, have been limited in scope in that they have been designed to cover designated sections of astronomy and not astronomy as a whole. Details of the evaluation of existing schemas are to be found in my thesis. [3] A single schema that is capable of providing a standard data structure across all of the branches of astronomy will take time to develop. Therefore care was taken during the research that in the initial stages of schema development the design of the structure enables later additions to it. The following Hypothesis was tested:

- That an extensible schema can be developed to provide a common data retrieval structure for astronomical data, it can be designed in such a way that it will extend to cover all areas of astronomical data and this schema can be shown to work within a software system which ensures that the queried data retrieved is subjected to schema validation.

This study was developed from a consideration of biochemistry information retrieval schemas in a paper by Marco Mesiti and colleagues which discusses XML solutions for the problems of the representation, integration and management of heterogeneous biological data, involving biological data types represented by a number of XML languages and schemas. These models discussed gave good guidance as they were developed to handle solutions to the problems of large amounts of disparate data over many disciplines, [6] not unlike the situation with Astronomical data.

doi: 10.2788/854791

2. THE LOGICAL MODEL

The extensible schema was developed for use as a controlling mechanism which only permits the use of astronomical data once it had been retrieved from databases and saved into an approved structure via schema validation. An XMLSchema was considered suitable for this for the reasons that Diagram 1 illustrates:

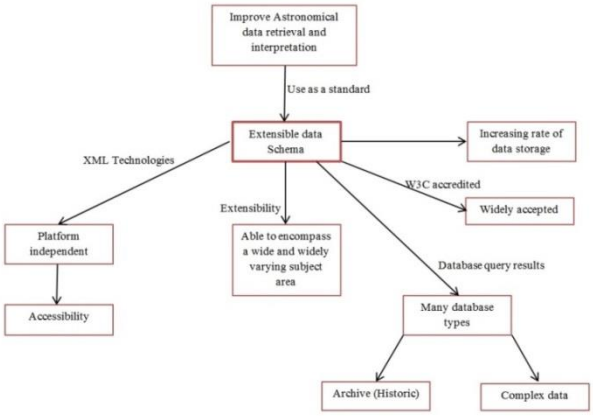

Diagram 1: Advantages of an XML schema

A custom software application developed as part of this research would be used to test the performance of the schema.

3. IMPLEMENTATION OF THE PHYSICAL MODEL

The structure of the astronomy XMLSchema was dictated by the known primary branches of Astronomy and sub branches that were needed for the astronomical data that was used during this research.

The testing that was carried out retrieved data of visible and infra-red type observations and so the XMLSchema was developed with the following structure in those sub-branches:

Astronomy (root)
Astrophysics
Astrobiology
Astrochemistry
Historical Astronomy
Planetary Science
Observational Astronomy
Visible
Infrared
Radio
Microwave
Gravity wave
Shortwave
Neutrino
Submillimetre

Table 2: Astronomical Schema high level nodes

All the branches of Astronomy other than 'Observational Astronomy' can be developed at some time subsequently as can all the types of 'Observational Astronomy' observations, other than 'Visible' and 'Infrared' types which were fully developed to the following detailed structure as shown in Figure 1 below:

Figure 1: Detail of Visible and Infra-red nodes

The inclusion of BLOB data (proposed as Base 64 encoding) storage nodes is particularly valuable as it allows for streams of data to be added to the xml that it is not possible to add as 'string' data.

The software developed to demonstrate the XMLSchema in a practical scenario was given the name Retrieval of Astronomical Data (ROAD). The structure of this software system is shown below in Diagram 2:

doi: 10.2788/854791

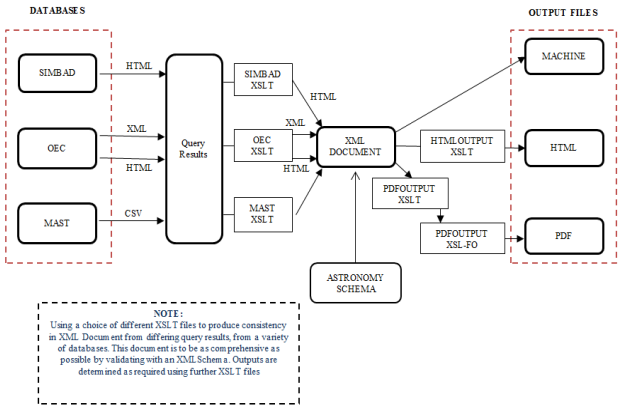

Diagram 2: ROAD system structure

The ROAD system currently has very basic, limited capabilities which are designed to just prove the concepts. A properly capable user friendly system has yet to be developed. It is able to query up to three separate databases, (those of OEC, MAST and SIMBAD) combining the returned data into a query results object which is then used to create an XMLDocument of standard structure containing data, as defined by the XSLT transformation files for each database. The user of the ROAD system can choose up to three output file types of PDF, HTML and MACHINE (which is basically XML for the purpose of machine to machine communication). The astronomy XMLSchema scans the XMLDocument and only permits the production of output files if the XMLDocument is of the approved structure.

The diagram above shows there are three different data types that were retrieved from the databases of HTML, CSV and XML which comprised four different formats in total as the HTML query responses were in different formats. The data was originally captured by both Optical and Infra-red telescopes. The table below shows those database data difficulties resolved by the ROAD system:

Multiplicity of formats	✓ use of XSLT
Transformation problems due to incompatible formats	✓ use of XSLT
Archaic problems with vintage data such as limited storage space, maximum string length and other restrictions of technology of their time	✓ schema nodes
Need for digitisation	Not resolved
The increasing rate of astronomical data capture both by institutions and amateurs	Not resolved
The growing different types of astronomical data collectors [1]	✓ schema nodes

Table 4: Retrieval Difficulties Resolved

4. RESULTS

Five test runs were made of the ROAD system of differing combinations of the databases being queried and different outputs requested. The results are shown in Table 5 below.

Run	Choice of Sources	Choice of Outputs	Actual Outputs
1	MAST SIMBAD OEC_XML OEC	HTML PDF Machine	resultsXmlDocument_15 03 2015 19-08-48.xml htmlOutput_15 03 2015 19-08-50.html machineOutput_15 03 2015 19-08-50.xml pdfOutput_15 03 2015 19-08-50.fo run1output.pdf (using Attic)
2	MAST SIMBAD OEC_XML	HTML	resultsXmlDocument_15 03 2015 19-13-28.xml htmlOutput_15 03 2015 19-13-30.html
3	SIMBAD OEC_XMLOEC	PDF	resultsXmlDocument_15 03 2015 19-17-05.xml pdfOutput_15 03 2015 19-17-06.fo run3output.pdf (using Attic)
4	MAST, OEC_XML	Machine	resultsXmlDocument_15 03 2015 19-24-50.xml machineOutput_15 03 2015 19-24-51.xml
5	MAST SIMBAD OEC_XML OEC.	HTML PDF Machine	resultsXmlDocument_15 03 2015 19-35-19.xml Schema error here

Table.5: ROAD Run Results

It can be seen from the table that runs 1 to 4 produced outputs of types chosen by the user within the available range of PDF, HTML and Machine (XML).

Note that in Run 5 the XMLDocument creation was deliberately made to be produced in an invalid format (by making changes to the XSLT files) and therefore no Output files were produced. Shown here is the log output for that test run which also gives information of the reason that the XMLDocument was found to be invalid:

```
==========================================
ROADS beginning data retrieval 15/03/2015 19:35:19
==========================================
Querying for data...
Querying for Mast
Querying for OEC
```


doi: 10.2788/854791

```
Querying for OEC_XML
Querying for Simbad
Query results returned
The element 'infrared' has invalid child element 'visualobservationdata'. List of possible elements expected: 'irobservationdata'.
resultsXmlDoc did not validate

XMLDocument created: C://roadLogs//resultsXmlDocument_15 03 2015 19-35-19.xml

*** Error *** No file outputs due to XMLDocument schema validation error
==========================================
ROADS finished data retrieval 15/03/2015 19:35:21
==========================================
```

Figure 2: Log output of Run 5

5. CONCLUSIONS and FUTURE DEVELOPMENTS

This extensible XMLSchema is capable of further development without compromising the structure created so far. This is due to the fact that the root node is that of 'astronomy' with all the different areas of astronomy as child nodes. More child nodes can be added or existing ones further developed in the future as they are needed.

All the test runs correctly produced one XMLDocument each containing combined data from the datasource queries that were chosen by the user. Output files were also produced from each run. The first four test runs correctly produced the output files requested by the user. The final fifth run did produce a schema error as expected and correctly, in accordance with expectations, did not create any user output files.

This research resolved several of the data retrieval difficulties by:

- Creating a partial astronomy XMLSchema
- Implementing the proof of concept ROAD software and then carrying out the test runs that retrieved data into XMLDocuments validated by the XMLSchema prior to any output of data to users.

It is this XMLSchema validation process that provides the control on the XMLDocument structure to ensure a standard structure. It is this standardization of structure that can reduce the loss of metadata (by ensuring that it is included) and enable easier retrieval and use. Additionally it is possible that building up a generally available library of XMLDocuments from completed database queries could be of benefit for accessing data more easily and completely. Also, building up an available library of XSLT files to be used for querying astronomical datasources and generating user outputs could be an increasingly valuable asset for the creation of XMLDocuments containing astronomical data. Also MACHINE type outputs could be a way of automating requests directly to telescopes and other types of detecting instruments to collect new data, if deficiencies in existing data were found.

Future Research

It is apparent that future research is required for the full development of the XMLSchema and a resource of XSLT and XMLDocument libraries. Such research would require skills of both a technical and sociological nature. On the technical side the XMLSchema needs to be extended across the full spectrum of astronomical disciplines and more XSLT files and data queries need to be developed. On the sociological side, this solution to the data retrieval problem needs to be attractive to the astronomical research and industry sectors. Achieving the W3C standard for the XML Schema would be a big step forward in achieving widespread acceptance.

6. REFERENCES

[1] VAO. (2015, May 1). *History of the Virtual Astronomical Observatory.* Retrieved from VAO: http://virtualobservatory.org/whatis/history.aspx

[2] NASA. (2014, October 21). *XDF: The Extensible Data Format Based on XML Concepts*. Retrieved from NASA: http://nssdc.gsfc.nasa.gov/nssdc_news/june01/xdf.html

[3] Beech, G. (2015, October). An Investigation of the Benefit of XML Technologies in Astronomical Data Interpretation. Huddersfield University, West Yorkshire, UK.

[4] Lopez, M. H. (2014, November 14). *The World's Technological Capacity to Store, Communicate, and Compute Information.* Retrieved from Science Magazine: http://www.sciencemag.org/content/332/6025/60

[5] Wall, M. (2014, November 17). *Astronomy Overload.* Retrieved from Space.com: http://www.space.com/9308-astronomy-overload-scientists-shifting-stargazing-data-mining.html

[6] Mesiti, M. (2011). XML-Based Approaches for the Integration of Heterogeneous Bio-Molecular Data. In H. H. Trimm, *Recent Advances in Biochemistry* (pp. 206 -). CRC Press.

TOWARDS SECURING REPRODUCIBILITY OF EARTH SCIENCE RESEARCH FINDINGS THROUGH AUTOMATIC POLICY ENFORCEMENT BASED ON MACHINE READABLE DATA MANAGEMENT PLANS

[1]*F. Engel,* [1]*H. Görzig,* [2]*P. Walsh,* [1]*M. Hemmje,*

[1]Fernuniversität in Hagen, Germany
[2] Cork Institute of Technology, Cork, Ireland

ABSTRACT

There exists an urgent demand to ensure full traceability and reproducibility of research findings to validate results for later reuse. The awareness of this fact is accompanied by increasingly stringent regulations from reimbursement agencies, which enforces compliance to specific policies for safety and reproducibility. Furthermore, research undertakings are expensive and the return on investment needs to be secured by research funding agencies and public bodies through proper management of the knowledge that is required for long term reuse. For this reason, reimbursement agencies are beginning to make Data Management Planning mandatory to ensure that later validation of research is feasible, to foster its proper reuse, prior to funding. Hence, this paper will discuss key aspects to enable future research validation to support reproducibility through proper enforcement of Data Management Planning.
Specifically, the approach discussed in this contribution is seeking to secure long term validation of research in compliance to widely accepted archival and digital preservation standards. In essence we propose a system that will enable researchers to harvest research results far into the future. We will discuss solutions and open challenges regarding this problem in the Earth Science domain on the basis of our work in the European FP7 projects SHAMAN, SCIDIP-ES, and APARSEN.

Index Terms— Reproducibility, Preservation, Data Management Plans

1. INTRODUCTION

SCIDIP-ES (cf. [1]) is a European Commission (EC) co funded research project whose aim was to support the preservation of research data in the Earth Science (ES) domain. Comprehensive access to and proper interpretation of preserved research results is essential for its reuse e.g. as basis for new research undertakings. However, beside pure bit stream preservation of research results, there is an increasing demand to secure the knowledge required for later result validation by ensuring research traceability and reproducibility. To accomplish this task the EC and further reimbursement agencies enforces the disposal of Data Management Plans (DMP) to ensure compliance to research data quality, sharing and security aspects, as an integral part for funding research projects in the recent EC horizon 2020 framework program (cf. [2]). Furthermore, publishing groups such as *Nature* (cf. [3]) and others have already start to mandate checklists to ensure the existence of a minimal set of information about all the material that is required to secure quality standards and support the validation of published findings. To address the issue, we believe that comprehensive support for DMP enforcement is critical for successful research reuse. In the following we will discuss the recent state of the art and related work of relevant technologies and elaborate on our solution towards supporting automated enforcement of DMPs.

2. STATE OF THE ART AND RELATED WORK

In the scientific literature there exists a multitude of definitions to describe the nature of research reproducibility. Just to mention a few: Dalle provides in (cf. [4]) a quite generic definition: "*Applied to in-silico experiments (i.e. computer simulations), we consider reproducibility as the ability to produce similar data by running a similar simulation*". Furthermore, he mentioned the importance of traceability that he defines as: "*Traceability is the ability to trace the dependencies of results and publications to the software elements that were used to produce them*". However, Freire et al. provides in (cf. [5]) a more formal definition: "*A computational experiment that has been developed at time t on hardware/operating system s on data d is reproducible if it can be executed at time t' on system s' on data d' that is similar to (or potentially the same as) d*".

To foster the reproducibility of funded research, reimbursement agencies have begun to make the provision of a *Data Management Plan* (DMP) mandatory, prior research funding. In general, a DMP is a full text document that elaborates how the research data is handled during- and after the project runtime. Since no common template for DMPs exist, reimbursement agencies provide specific frameworks that lists general policies for data management that are need to be addressed at a minimum within a DMP. In this paper we will address the validation of research through the ability of

doi: 10.2788/854791

preserved research reproduction. Hence, we will now introduce some models and technologies in the scope of digital preservation.

Our work in the SHAMAN research project has already outlined that from the perspective of long-term preservation, information objects undergo an archive-centric information life cycle with the phases *Creation*, *Assembling*, *Archiving*, *Adoption*, and *Reuse* (cf. [6]). Within the *Creation* phase new information comes into existence. The *Assembly* phase denotes the appraisal of objects relevant for archival and all processing and enrichment for compiling the complete information set to be sent into the future, meeting the presumed needs of the Designated Community. Assembly requires in depth knowledge about the Designated Community in order to determine objects relevant for long term preservation together with information about the object required for identification and reuse some time later in the future. The *Archival* phase addresses the life time of the object inside the archive. The *Adoption* phase encompasses all processes by which archival packages are unpacked, examined, adapted, transformed, integrated and displayed to be used and understood by the consumer. The *Reuse* phase however, means the exploitation of information by the consumer. In particular, reuse may be for purposes other than those for which the information object was originally created.

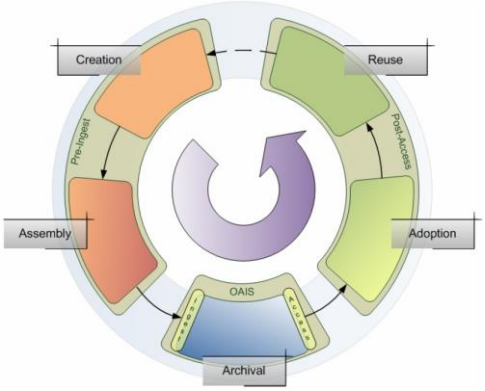

Figure 1: Archive centric Information lifecycle (cf. [6])

The information life cycle proves the complex relations between an information object and further information that come into existence before and after its archival. In the SHAMAN and the SCIDIP-ES project we therefore have envisioned and partly implemented a *Packaging Service* for packaging information for long term preservation of information that came into existence during the single phases of the archive centric lifecycle, in conformance to the accepted OAIS Information Model (cf. [7]). OAIS is an ISO standard (ISO 14721), commonly applied as framework of terms and concepts for an archival information system and is tightly coupled to the ISO 16363 standard that ensures audit and certification of trustworthy digital repositories. Further OAIS compliant services and toolkits that are related to the Packaging Service have been developed during the runtime of the SCIDIP-ES project like the *HAPPI Toolkit* and the *GAP-Identification* Service. The HAPPI-Toolkit manages the evidence histories of a preserved object. The evidence history contains various evidence records that are effectively added to an evidence history after each change on an information object, representing the object's provenance. In contrast the GAP-Identification Service could be applied to calculate the intelligibility gap for a certain piece of information regarding a specified designated community.

Since OAIS is a framework for any kind of digital encoded information object the traceability and reproduction of research findings are not explicitly addressed. Therefore, DMP and reproducibility support has not been explicitly handled within the SCIDIP-ES project. However, reproducibility support has been intensively discussed and outlined during our work in the educational part of the European APARSEN project (cf. [8]). In the ES and related communities the demand for traceable and reproducible research is addressed among others through the Earth-cube initiative (cf. [9]) or the Earth Science Informatics Technical Committee (ESI TC, cf. [10]).

However, if a project with DMP is funded from a reimbursement agency, the policies specified in the DMP are required to be enforced. A readily available solution to implement such policies on data management level is the implementation and enforcement of rules that specify under what conditions what action should be executed. A first investigation towards automation of policies derived from DMPs, on this basis has been undertaken by Chen et al. on basis of the *integrated Rule Oriented Data System* (iRODS, cf. [11]). Within this publication Chen et al. are providing a set of so called actionable rules that are formalized and enforced through the application of the iRODS rule language. Though, our experience in recent research projects indicates that often common content management systems, in combination with version control systems and further technologies, are applied in this task.

Aside of the iRODS rule language various further technologies to formalize and enforce rules has been developed in the special field of knowledge management, like in the context of the Semantic Web. The Semantic Web is an advancement of the World Wide Web that yields towards a machine readable reprocessing of resources, for easy exchange of information over system barriers (cf. [12]). This approach should enable effective access to resources and automation of processes along with integration and reuse of information. Ontologies play a key role in technologies that are part of the Semantic Web, as such they are applied to build a shared vocabulary as a basis to enable access to distributed datasets over system barriers. Such technologies are among others W3C specification of description languages as OWL to enables expression of ontologies. Beside representation of ontologies the specification and enforcement of rules is also

doi: 10.2788/854791

a part of the Semantic Web. In general, rules in the context of the Semantic Web denote the represented knowledge that could not be expressed via existing description languages (cf. [14]). Among others the W3C publishes the *Semantic Web Rule Language* (SWRL, cf. [15]) and the *Rule Interchange Format* (RIF, cf. [16]). Furthermore, to make use of such encoded knowledge a so called *Reasoner* is applied. A Reasoner thereby deduces logical consequences on basis of a provided knowledge base. A huge amount of different Reasoner implementations exist, including the Jena Reasoner (cf. [17]). Most of these implementations come with an Application Program Interface that supports programmatic access to the underlying ontology and deduction of statements.

3. AUTOMATIC ENFORCEMENT OF DMP

A DMPs specifies the policies that determine how research data is handled during the runtime of a project and after the project is ended to ensure its proper reuse. Keeping the reproducibility of research findings into the long term is thereby strongly influenced through a proper enforcement of such policies. Starting from our participation in recent EC funded project our observation is that various actors participate in the process of DMP specification and influence the research data management in different dimensions (cf. [18]). The Formal Dimension involves the project administration, the Managerial Dimension the project management and the Operative Dimension the project implementation and execution. More specifically, the ***Formal Dimension*** of DMPs is spanned by the funding agencies' Grant Agreements (GA), corresponding laws and policies. The GAs are usually providing the contractual framework for DMP. Therefore, it specifies what the DMP has to accomplish and to comply with. Corresponding laws and regulations therefore provide the legal, regulatory and subsequent policy building framework. To comply with the requirements and challenges created by the analysis of the formal layer, a RDM work plan is developed in the ***Managerial Dimension*** of DMP. The RDM work plan describes the RDM scenario that has to be created to comply with the DMP requirements and challenges and their corresponding representation schema set up by the analysis of the formal dimension. This RDM work plan includes strategic and organizational aspects, concrete activities, and deliverables. In the RDM work plan sequences of activities and their dependencies are formulated. Therefore, the implementation of the DMP is based on this RDM work plan. The data producers such as software developers and researchers in the project form the ***Operative Dimension*** of the DMP. Tasks and activities listed in the work plan are executed by them and thereby produce and use the data to be archived and preserved for effective later reuse. In essence, all of the three dimensions are influence the strategy of research data management for its comprehensive later reuse, in various stages of the archive centric information life cycle. As indicated through the work conducted by Chen et al. the available solution to implement and enforce policies derived from DMPs through rules is appropriate. Since iRODS is a complex system, well suited for projects that needs grid computing support, but not applicable for every application, we will follow a more generic approach that is applicable, on top of existing data management systems. From the experience gained during our research in several EC co-funded research projects, we thereby expect a software system based on Semantic Web technologies is appropriate in this task. This will support easy exchange and processing of information on basis of standardized software in a data management independent approach.

Hence, we suggest a web based *Service System of Rules for Research Data Management* (SSRRDM, cf. Figure 2), that could be partly implemented through existing technologies stemming from the several EC co-funded research projects, like the HAPPI Toolkit for proper assembly of provenance trails, the GAP-IS Service to keep preserved information understandable in the long term or the Packaging Service to assemble all required information to ingest, manage and access preserved findings for its later reproduction. In essence we envision on base of the SSRRDM support for data managers and funders with respect to the following three challenges: ***1.) Assembly of rules:*** A DMP is a full text document that specifies policies for research data management. Before policies, derived from a given DMP on the operational dimension could be enforced and validated automatically, it is required to assemble a set of corresponding rules that fulfils derived policies. The SSRRDM will provide a registry for the cataloging and access to configurable rules based on Semantic Web technologies for derived DMP policies on top of existing reasoning systems. ***2.) Rule enforcement and Management:*** Beside registration, it is necessary to configure rules given the specific requirements, to update these in case that the policies of a DMP are changing and to enforced them through an underlying rule system on top of a data management system. The SSRRDM will provide a generic interface to configure a set of rules and enforce rules through application of appropriate reasoned APIs. ***3.) Archival and validation:*** To prepare the archival of data for later reuse, all applied rules and available constituents of research, like its provenance, will be recorded through the HAPPI Toolkit and serialized and packed through the Packaging Service. Since research data is becoming increasingly complex and is often interlinked with (external) resources, archiving them requires a structured and machine serialization mechanism as provided by Semantic Web technologies that it is already foreseen within the Packaging Service implementation. Further important knowledge for comprehensive later reuse

doi: 10.2788/854791

could be gathered through an interface with the GAP-IS service.

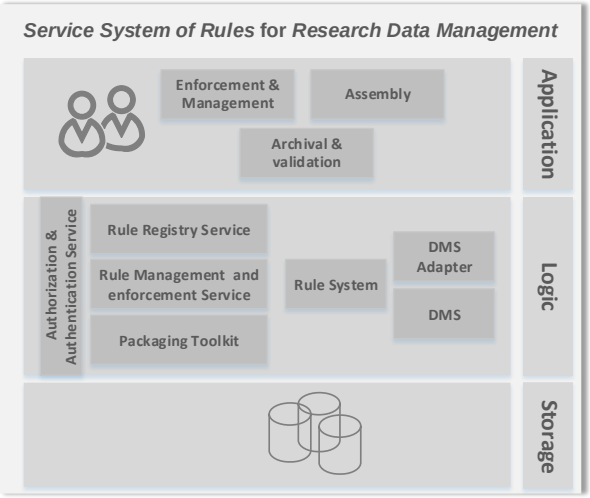

Figure 2: SSRRDM architecture

Our next steps are a deep analysis of public accessible DMPs and corresponding frameworks as basis for the rule registry instantiation, beside a first harmonization and integration of existing SCIDIP-ES Services and Toolkits.

4. CONCLUSION

Reproducibility of research findings is a growing demand to secure result validation from reimbursement agencies. Data Management Planning is the approach of reimbursement agencies to ensure the existence of critical information to enable reproducibility of research findings for its later validation. We argue that technologies in the context of the Semantic Web are applicable to automate enforcement of policies derived from DMP for proper research data management. The implementation of a system that enables creation and automation of enforcement of DMPs could partly be built on services and toolkits from the SCIDIP-ES project building on methods from the SHAMAN project and insights gained in the APARSEN project.

5. REFERENCES

[1] S. Crompton, D. Giaretta, B. Matthews, F. Engel, H. Brocks, A. Shoan, F. Marelli (2014). SCIDIP-ES: A Sustainable Data Preservation Infrastructure to Support OAIS Conformant Archives (2014). In: APA/C-DAC International Conference on Digital Preservation and Development of Trusted Digital Repositories

[2] EC, Guidelines on Data Management in Horizon 2020 (2013). URL:http://ec.europa.eu/research/participants/data/ref/h2020/grants_manual/hi/oa_pilot/h2020-hi-oa-data-mgt_en.pdf. (last accessed 2015.10.21)

[3] Editorial, Joining the reproducibility initiative (2014), Nature Publishing Group. http://dx.doi.org/10.1038/nnano.2014.287

[4] Dalle, O. (2012). On reproducibility and traceability of simulations. Simulation Conference (WSC) (pp. 1 - 12). IEEE

[5] Freire, J., Bonnet, P., & Shasha, D. (2012). Computational reproducibility: state-of-the-art, challenges, and database research opportunities. SIGMOD International Conference on Management of Data (pp. 593-596). ACM.

[6] Brocks, H., Kranstedt, A., Jäschke, G., and Hemmje, M. (2010). Modeling context for digital preservation. In Smart Information and Knowledge Management. Springer Berlin Heidelberg.

[7] CCSDS: Reference model for an open archival in-formation system (OAIS). Pink Book 1, Consultative Committee for Space Data Systems (2012).

[8] Proll, S., & Rauber, A. (2013). Scalable data citation in dynamic, large databases: Model and reference implementation. In Big Data, 2013 IEEE International Conference on (pp. 307-312). IEEE.

[9] Earthcube. URL http://earthcube.org/ (last accessed 2015.10.21)

[10] Ramachandran, R., & Khalsa, S. J. S. (2015). Moving from Data to Knowledge: Challenges and Opportunities. Geoscience and Remote Sensing Magazine, IEEE, 3(2), 51-54.

[11] Chen, S. Y., Conway, M., Crabtree, J., Lee, C., Misra, S., Moore, R. W., . . . Xu, H. (2015). Policy Workbook -iRODS 4.0. Chapel Hill, North Carolina, USA: iRODS Consortium

[12] (2016, 01 12). Semantic Web. Retrieved 01 28, 2016, from Semantic Web: https://www.w3.org/standards/semanticweb/

[14] Hebeler, J., Fisher, M., Blace, R., & Perez-Lopez, A. (2012). Semantic web programming. John Wiley & Sons.

[15] W3C. (2004, 05 21). SWRL: A Semantic Web Rule Language Combining OWL and RuleML. Retrieved 01 28, 2016, from SWRL: A Semantic Web Rule Language Combining OWL and RuleML: https://www.w3.org/Submission/SWRL/

[16] W3C. (2013, 02 05). RIF Primer (Second Edition). Retrieved 01 28, 2016, from RIF Primer (Second Edition): https://www.w3.org/TR/rif-primer/

[17] Jena. (2015). Reasoners and rule engines: Jena inference support. Retrieved 01 28, 2016, from Reasoners and rule engines: Jena inference support: https://jena.apache.org/documentation/inference/

[18] Görzig H., Engel F. Hemmje M. Lessons Learned and Open Challenges Regarding DMP and RDM Support. iPRES 2015. Accepted, to be published

doi: 10.2788/854791

WED BASED NASA BIG DATA VISUALIZATION AND ANALYSIS

Emily Law[1], George Chang[1], Richard Kim[1], Shan Malhotra[1]

(1) Jet Propulsion Laboratory, California Institute of Technology (Emily.S.Law@jpl.nasa.gov, George.Chang@jpl.nasa.gov, Richard.M.Kim@jpl.nasa.gov, Shan.Malhotra@jpl.nasa.gov)

ABSTRACT

NASA's earth and planetary spacecraft return large amounts of remote sensing data, such as imagery and raw science measurements, in support of remarkable research. Not only does the data lead to new scientific discoveries about our planet and the solar system, it provides a wealth of information to educate, inspire, and engage the public at large. To leverage this rich data for mission planning, scientific research and public outreach, it is essential to make it accessible and understandable, analyzable, all while appealing to their interests. This poster will cover web-based capabilities that showcase NASA's large volume of remote sensing data. We will illustrate it in easy-to-use and interactive mediums for diverse use via advanced visualization and analysis tools. We will also present future work supporting big data from space.

1. INTRODUCTION

Meeting the challenges of space exploration has resulted in new knowledge that has kept NASA and JPL world leaders in big data science and technology. We, NASA's Lunar Mapping and Modeling Project (LMMP) team at JPL, under the management of Solar System Exploration Virtual Research Institute (SSERVI), have developed a system that includes a set of web-based portals and a suite of interactive visualization and analysis tools to enable mission planners, lunar scientists, and engineers to access a large volume of mapped data products from past and current missions. This presentation will provide an overview of this system highlighting its web portals that are designed for space and Earth exploration. We will also demonstrate their uses and capabilities, highlighting big data visualization and analysis innovations.

2. LUNAR MAPPING AND MODELING PORTAL (LMMP)

LMMP (http://lmmp.nasa.gov) provides a suite of interactive tools that incorporate observations from past and current lunar missions, creating a comprehensive lunar research Web portal. The online Web portal allows anyone with access to a computer to search through and view a vast number of lunar images and other digital products. The portal provides easy-to-use tools for browsing, data layering and feature search, including detailed information on the source of each assembled data product and links to NASA's Planetary Data System. While mission planning is LMMP's primary emphasis, LMMP also addresses the lunar science community, the lunar commercial community, education and outreach, and anyone else interested in accessing or utilizing lunar data. Its visualization and analysis tools allow users to perform analysis such as lighting and local hazard assessments including slope, surface roughness and crater/boulder distribution. LMMP features a generalized suite of tools facilitating a wide range of activities including the planning, design, development, test and operations associated with lunar sortie missions. Sharing of multi-layered visualizations is made easy with the ability to create and send LMMP bookmarks. LMMP is also a powerful tool for education and outreach, as is exemplified by its mobile clients (Moon Tours for iOS and Android), serving of data to NASA's Eyes on the Solar System, and serving of data to a growing community of digital planetariums.

3. VESTA PORTAL

Vesta Trek (http://vestatrek.jpl.nasa.gov), a web-based application applying LMMP technology to visualizations of the asteroid Vesta. Data gathered from multiple instruments aboard Dawn have been compiled into Vesta Trek's user-friendly set of tools, enabling users to study the asteroid's features. The application includes:

- Interactive maps, including the ability to overlay a growing range of data sets including topography, mineralogy, abundance of elements and geology, as well as analysis tools for measuring the diameters, heights and depths of surface features and more.
- 3-D printer-exportable topography so users can print physical models of Vesta's surface.
- Standard keyboard gaming controls to maneuver a first-person visualization of "flying" across the surface of the asteroid.

doi: 10.2788/854791

4. MARS TREK

Using large amount data from more than 40 years of Mars scientific instruments and missions, Mars Trek (http://marstrek.jpl.nasa.gov) allows astronomers, citizen scientists and students to visually study the Red Planet's features. The application allows users to easily visualize data collected from multiple instruments on multiple spacecraft orbiting Mars. Mars Trek is currently being used by the team working on selecting possible landing sites for NASA's upcoming Mars 2020 rover. Mars Trek will also be used as part of NASA's newly announced process to examine and select candidate sites where humans will first explore on Mars. The recent release of the movie, "The Martian", provides an excellent opportunity for public outreach, showcasing NASA's plans for the exploration of Mars. One of the ways in which NASA and the movie studio are collaborating is to have features specific to the story of the Martian featured in the recent update of Mars Trek. This new feature of Mars Trek allows users to explore the path across Mars described in the story while learning from NASA commentary.

5. SMAP VIEWER

Soil Moisture Active Passive (SMAP) Viewer (http://smap.jpl.nasa.gov/map/) is web-based portal application applying LMMP technology to visualizations and analysis of data captured by the SMAP mission. SMAP is an orbiting observatory that measures the amount of water in the top 5 cm (2 inches) of soil everywhere on Earth's surface. The topsoil layer is the one in which the food we eat grows and where other vegetation lives. The SMAP viewer presents the data in advance visualization allowing science teams to compare and contrast between the different data layers in support of calibration and validation of their algorithms, and to characterize the quality and level of uncertainty. In addition, analysis tools provided by the viewer support statistical analysis that supports agricultural models.

6. TECHNOLOGIES

The suite of web portals provided by the LMMP project is built on a Service Oriented Architecture (SOA)[1] that is scalable and extensible for all planetary bodies. The portals are supported by a common backend infrastructure and use a common frontend visualization framework. The infrastructure provides core services for data ingestion, data management, image tiling using hadoop [2], arcGIS [3], and workflow using cloud computing for various computation and data analysis services, search via SOLR [4] and download. The frontend framework takes advantage of HTML5 [5] frontend user interface. It employs a flat open space design and implementation that maximizes usability with modular tools and widgets. It includes in browser over sampling that stretches the image dynamically when zoom in past its resolution level. It is responsive to different sizes and form factors. It is embeddable into other browsers.

7. SUMMARY AND CONCLUSION

NASA's Lunar Mapping and Modeling Project (LMMP) has grown considerably from its initial goal of providing a mission planning tool for Lunar exploration. Recognizing the big data trend and challenges, the team had the foresight to architect and design the system (including a combination of a backend infrastructure and a common user interface framework) to be scalable and extensible. The system is now successfully providing capabilities for multiple activities including mission planning, scientific research, decision making, as well as public outreach for other planetary bodies via its web portals. In addition to continual enhancements to the system, the team continues to expand its capabilities to new destinations and new research. We also encourage and invite the user community to provide suggestions and feedbacks as the development team continues to expand the capabilities of the system, its related products, and the range of data and tools that we have provided.

8. ACKNOWLEDGEMENT

The authors and the LMMP team would like to thank the Advanced Explorations Systems Program of NASA's Human Exploration Operations Directorate, the Planetary Science Division of NASA's Science Mission Directorate, and the Solar System Exploration Virtual Research Institute for their support and guidance in the continuing development of LMMP.

.

9. REFERENCES

[1] "What Is SOA?". opengroup.
[2] "Welcome to Apache Hadoop!". hadoop.apache.org
[3] https://www.arcgis.com/features/
[4] http://lucene.apache.org/solr/
[5] Leslie Sikos. "HTML5 Became a Standard, HTML 5.1 and HTML 5.2 on the Way"

IMPROVING EARTH SCIENCE DATA DISCOVERABILITY AND USE THROUGH METADATA RELATIONSHIP GRAPHS, VIRTUAL COLLECTIONS, AND SEARCH RELEVANCY

Andrew Mitchell, Dawn Lowe, Kevin Murphy, Daniel Pilone

NASA

ABSTRACT

NASA's Earth Observing System Data and Information System (EOSDIS) is a multi-petabyte-scale archive of Earth science data that supports global climate change research by providing end-to-end services from instrument data collection to science data processing to full access to earth science data. This paper discusses efforts to improve the discoverability and usability of earth science data by exploiting EOSDIS' extensive metadata holdings in NASA's Earth Science Metadata system, the Common Metadata Repository (CMR), which provides sub-second complex search capabilities across 275,000,000 pieces of science metadata over an archive growing at a sustained rate of 8.5TBs per day. With the initial deployment of the CMR, NASA EOSDIS made a significant step forward in metadata quality. CMR enhancements are now focused on improving the discoverability and accessibility of high value Earth Science data by exploiting the richness of EOSDIS metadata, dramatically improving the discoverability of tailored data products for end users and more comprehensive responses to search requests.

***Index Terms*—** Common Metadata Repository, metadata catalog, unified metadata model, relevancy

1. BACKGROUND

NASA's Earth Observing System Data and Information System (EOSDIS) has continually evolved to improve the discoverability, accessibility, and usability of high-impact NASA data spanning a multi-petabyte-scale archive of Earth science data products. NASA has supported multiple catalogs enabling discovery of this data for many years including the Global Change Master Directory (GCMD) and the EOSDIS Clearing House (ECHO). EOSDIS needed a sub-second search capability to address discovery needs for its Big Data holdings while continuing an average daily archive growth of 8.5TBs. Over the last several years NASA has conducted a number of stakeholder surveys to help satisfy the expectations of a user community that consumes a daily distribution of 22TBs of data from EOSDIS archives.

The Common Metadata Repository (CMR) builds on the work done by ECHO and the GCMD to provide a unified, authoritative repository for Earth Science metadata. The CMR enables Earth Science applications to provide end users with nearly immediate access and interactivity across massive stores of Earth Science data by providing *high performance, standards compliant*, *temporal*, *spatial*, and *faceted* search of the associated metadata.

The CMR supports an environment that has decades of existing data and metadata; guidance for creation and maintenance of quality metadata is critical to the success of the system. To address this, NASA developed the Unified Metadata Model (UMM) and corresponding documentation[1].

The UMM, at the very top level, defines the idea of Metadata Profiles; potentially complex, cohesive ideas in EOSDIS that may be related to other profiles through parent-child or simple association relationships. The current version of the UMM identifies Science Collections, Science Granules, and Meta-Metadata profiles. Visualization, Parameter, and Service metadata profiles are slated for a subsequent UMM version. The UMM is intended to be a living model – evolving on quarterly updates.

The UMM does not, however define a file format. Instead, the UMM defines the EOSDIS metadata model, identifies required and recommended fields, then provides mappings to widely supported and standardized metadata formats including GCMD's DIF, NASA ECHO's ECHO10, and ISO19115.

Leveraging the UMM for its internal metadata model, the CMR exposes a set of public facing RESTful APIs intended to be consumed by client applications[2]. These APIs provide search, metadata retrieval, metadata validation, and ingest capabilities to both NASA and third party applications throughout the Earth Science community. An end user, unless they are writing their own search tool, typically doesn't interact directly with the CMR, instead they leverage one of the many visual or automated discovery tools built on top of the CMR APIs. This once-removed from the end user aspect of the CMR becomes a

doi: 10.2788/854791

complicating factor in the discussion of user-intent modeling below.

2. CHALLENGES OF DATA DISCOVERY

While the CMR and UMM made major advancements in terms of metadata quality and providing a single point of high performance, authoritative metadata for the EOSDIS holdings, the sheer number of data collections available to client applications and therefore end users made finding and returning relevant datasets a critical factor in streamlining data discovery. For example, finding a data collection with the words "sea ice" from the metadata holdings is simple; finding the latest version of a collection that best fits the end-user's needs of performing a time series analysis of sea ice over a given period of time is more complicated. An informal investigation of finding ozone data distributed by the GES-DISC DAAC using three top EOSDIS search tools by searching for the text "ozone" illustrated the discovery challenges with current data holdings as shown in Table 1.

	Results	Precision	Recall	False Positive	False Neg.
GCMD	209	61%	98%	81	2
ECHO	179	65%	89%	63	14
Mirador	65	82%	41%	12	77

Table 1: Effectiveness of sample search tools looking for Ozone data with a text search for "ozone".

Further exacerbating the situation is that some data can be processed "OnDemand" to produce derivative products that were not described by EOSDIS metadata. Fundamentally, these products are no different than Level 3 products being produced from processing Level 2 products which were produced by processing Level 1 products, etc. except that these products, given their targeted user base and limited distribution, were not cataloged and described within EOSDIS. Discovery of these products without contacting the source Distributed Active Archive Center (DAAC) was extremely difficult.

In order to improve the discoverability of EOSDIS products and quality of CMR search results NASA is pursuing two different enhancements to the CMR: adding support for virtual products and improving the search implementation to enhance result relevancy.

3. CMR VIRTUAL PRODUCTS

In addition to standard data products, NASA's EOSDIS has the ability to create dynamic derived products when requested by a user. Called OnDemand Products, these products share a common source but are the result of additional processing or user requested extraction (subsetting) of specific pieces of information.

For example, the ASTER L1A product is used to create multiple OnDemand products including Level 2 Surface Emissivity, Surface Reflectance, and Surface Temperature as shown in Figure 1.

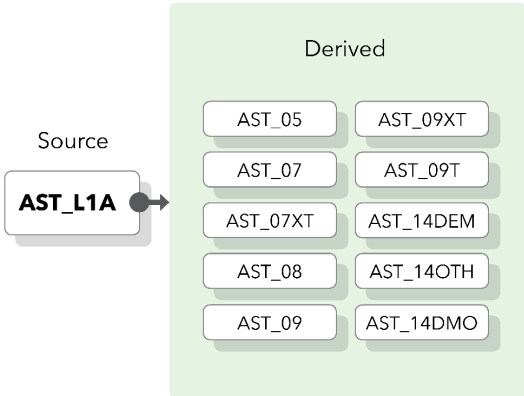

Figure 1: A single source product can be processed into multiple derived, possibly custom products

In other cases, many different physical variables are packaged into a standard product but not all are applicable for a given user. Aqua AIRS Level 3 Data is a great example of a broad range of measured parameters packaged into a single product as shown in Figure 2. Subsetting tools like OPeNDAP[3] can be used to extract the desired parameters, but require effort on the user's part.

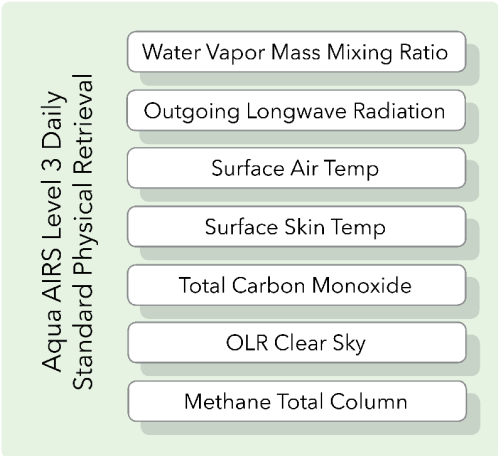

Figure 2: Products can contain hundreds of measured parameters, users typically want a relatively small subset

doi: 10.2788/854791

While derived and custom data products are tremendously valuable to end users, discovery of these products is difficult and inconsistent explicitly because of their dynamic nature. OnDemand products typically have limited or no discovery level metadata that tools can leverage to help users find and access the data.

The CMR's microservice architecture and asynchronous ingest pipeline provide hooks to capture incoming source data products. Rather than create derived product metadata by hand, a configurable metadata generation service was developed and attached to the CMR. This metadata generator can watch the CMR ingest flow and dynamically create and maintain metadata for derived products as shown in Figure 3.

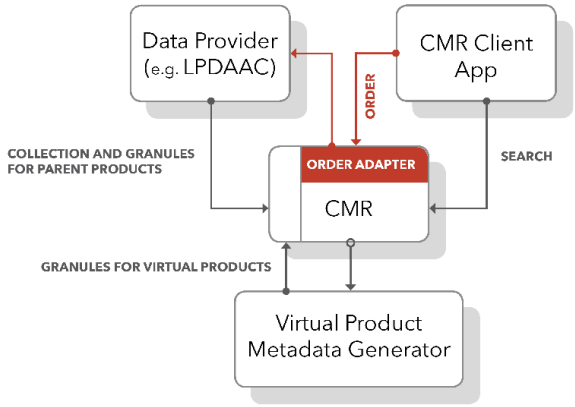

Figure 3: The CMR Virtual Product Metadata Generator monitors incoming source metadata and maintains derived virtual product information.

While the actual product may not exist on disk, metadata fully describing the collections and granules exists in the CMR and as far as client applications are concerned, represents real products available to users. The CMR can transparently unwrap virtual product requests for compatibility with existing subsetting services while users can request derived products that meet their unique needs, such as an AIRS derived methane product.

With OnDemand products properly supported, configured derived products were now theoretically discoverable in CMR client applications, however the challenge of returning more relevant search results for users was made even more important.

4. IMPROVING CMR SEARCH RESULT QUALITY

The CMR provides high performance spatial, temporal and keyword search capabilities but the initial implementations of each algorithm incorporated basic assumptions of importance when ranking matches. With the increased number of possible matches NASA started an effort to provide more accurate search results.

NASA DAACs have a wealth of knowledge about the data hosted by the DAAC and their user community. Through their user services they have expertise in providing guidance on best fit data for user needs. This expertise can be used to evaluate current relevancy with respect to the "ideal" rankings.

In addition to direct ranking information, an initial set of metrics can be established to measure the effectiveness of the results against a set of representative searches using traditional relevancy measures of precision and recall. These metrics can be augmented with user behavior information such as searches-to-download conversions, workflow monitoring and measurement, and actual A/B testing with relevancy variations.

While these metrics are being developed, an effort has been undertaken to improving the relevancy of the search results returned by CMR keyword searches. Keyword searches are free text searches across multiple fields within the metadata. Unlike other types of searches in the CMR, results matching keyword searches are assigned relevancy scores based on the fields that match the given query. Each field is assigned a "boost" value and the boosts of all the fields that match a given query are multiplied together (boosts are always >= 1.0) to generate a relevancy score. The more fields that match, the higher the relevancy.

The initial boosts used by the CMR were chosen based on perceived importance of the fields, so the "optimal" values for these boosts is not known. The CMR APIs were updated to provide a mechanism for client applications to provide desired boost values for specific metadata fields while executing queries.

The addition of the boost APIs provides the foundation for several significant relevancy improvements for CMR client applications:

- Applications aware of their user base can appropriately weight boost scores to provide results matching their users' expectations.
- Programmatic A/B testing of boost score variations can be done with no CMR changes and correlated to user behavior.

doi: 10.2788/854791

- Basic machine learning approaches can be used to develop boosting scores that produce results closely matching DAAC provided "ideal" rankings.

These APIs have been deployed in the CMR test environment and are available for public use. It is expected that there is no single optimal weighting for **all** scenarios and the availability of the boost APIs allows for user or use case specific weightings specified by client applications at runtime.

Moving beyond basic keyword relevancy improvements, additional relevancy improvements include temporal and spatial algorithm enhancements. Temporal relevancy ranking can be improved in several ways, but our initial investigations explore two approaches:

- Adjusting temporal relevancy weight based on the temporal overlap of the match with the temporal search space. The greater the overlap, the higher the relevancy score.
- Adjusting temporal relevancy weight based on correlation with an external, temporally significant event. For example, aerosol measurement searches corresponding with date ranges of volcanic eruptions could be weighted more heavily the more closely they align with the actual eruption event if "volcano" is a search term.

Similar approaches can be applied to to spatial relevancy improvements by computing percentage of spatial overlap and / or association with external, spatially bound events.

5. OUTSTANDING CHALLENGES

With the initial deployment of the CMR, NASA EOSDIS made a significant step forward in metadata quality. CMR enhancements are now focusing on improving the discoverability and accessibility of high value Earth Science data by exploiting the richness of EOSDIS metadata.

The CMR virtual product capability dramatically improves the discoverability of OnDemand and tailored data products for end users. Future enhancements can leverage the virtual product capability to create bundled products of related data, complex transformations, or even enable on demand generation of currently processed data in real-time, avoiding the need for long term storage of reproducible products.

By further leveraging the UMM's ability to model relationships between profiles along with improved relevancy ranking, the CMR can be enhanced to provide more comprehensive responses to search requests. An additional UMM profile to support unstructured content (UMM-D) is being developed. By associating UMM-D records with existing metadata collection records (UMM-C), parameter information (UMM-P), and visualization information (UMM-V), the CMR will be capable of providing comprehensive results to searches across EOSDIS holdings. This diversity of information enables client applications, such as NASA's EarthData[4], to display search result summary sheets as shown in Figure 4.

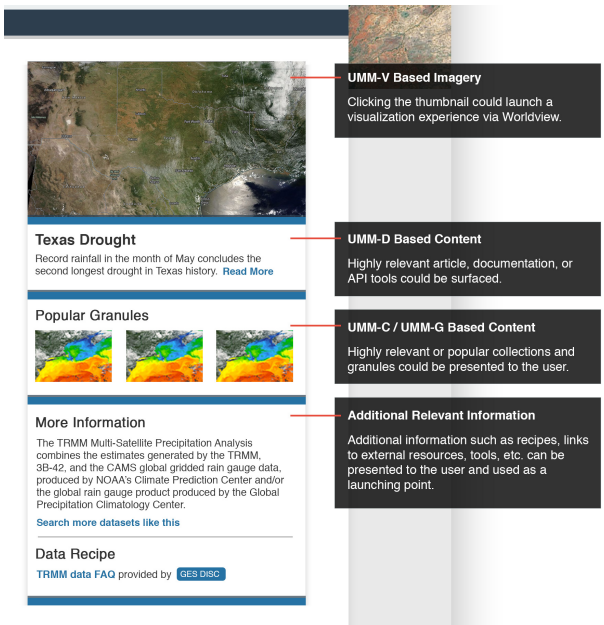

Figure 4: Mock up of a search result summary sheet that could be produced leveraging rich CMR search results.

Summary sheets would contain highly relevant, diverse information pulled from multiple sources across EOSDIS with drawn from the high quality metadata made available through the CMR.

The CMR and all of the associated efforts on relevancy and complex response handling is available to the international community at https://cmr.earthdata.nasa.gov and the system is being open sourced for enhanced reuse.

[1] NASA, https://wiki.earthdata.nasa.gov/display/CMR/UMM+Information
[2] NASA, http://cmr.earthdata.nasa.gov/
[3] OPeNDAP, http://www.opendap.org/
[4] NASA, http://earthdata.nasa.gov/

INNOVATIVE WAYS OF VISUALIZING META DATA IN 4D USING OPEN SOURCE LIBRARIES

Balhar Jakub, Veselka Jonáš, Vlach Pavel, Voumard Yann

Gisat s.r.o., Solenix inc.

ABSTRACT

There are more and more data being measured by different Earth Observation satellites around the world. Ever-increasing amounts of these data present new challenges and opportunities for their visualization.

In this paper we propose how to visualize the amount, structure and distribution of such data in a transparent way, which takes into account the time-dimension as well. Our approach allows to get both the global overview and detailed regional information about EO missions distribution.

We introduce our java application for dealing with heterogeneous ESA data sources as well as a mobile-friendly and easy-to-use web mapping application for 4D visualization of this data.

***Index Terms*—** Heterogeneous data sources, Open Data, Big Data Exploitation, 4D Data visualization

1. INTRODUCTION

There are more and more data being measured by different Earth Observation satellites around the world. Ever-increasing amounts of these data are made publicly available free of charge, presenting new challenges and opportunities for their visualization. The data as well as the metadata needs to be presented to the public in an engaging and transparent way and the final visualization needs to be appealing, interactive and easy to use. It is important to present the data in a way which allows mission managers, stakeholders and the general public to have a punctual overview of which different earth observation missions are providing useful information to scientific communities for the different areas of our planet and how are they doing it.

In this paper we discuss how to handle the data from various sources available by ESA. There are large volumes of the data available. Our aim is to store them in a way, which allows a simple aggregation needed for the visualization in the different levels of detail.

Second issue is the visualization of the data and metadata about ESA missions in a way, which will be responsive, fast and easy to use and therefore deliver the information to a broad scale of users efficiently. In this case, we are focusing on displaying the earth coverage built from data acquired by different missions and sensors in the time series. It is important for us to display how the amount of products evolves over the time.

Our goal in general is to provide open source application presenting the metadata about the data retrieved by different Earth Observation satellites, which takes into account all the aspects mentioned above.

The paper can be divided into two parts. In the first chapter we mention the ESA data sources for missions and discuss the possibilities for visualizing the missions' metadata. In the second chapter we present the API and the mobile-friendly application we developed for this purpose.

2. METHODOLOGY

In this part we introduce the main data sources, which we are using in our application. In the second part we propose the way to visualize the missions' metadata in a meaningful way, with the main focus on their amount and distribution.

2.1. Access to available data sources

Our application is interacting with three ESA data and metadata sources. For access to the Copernicus data the Sentinel Scientific Data Hub (SciHub) is used. Federated Earth Observation (FedEO) catalogue is used for access to other ESA and Third-Party missions. GPod is used for the acquisition of aggregated data for Copernicus as well as for other ESA and Third-Party missions.

The SciHub exposes two dedicated APIs for browsing and public access of the EO data stored in the rolling archive, Open Search and OData. OpenSearch is a set of technologies that allow publishing of search results in a standard and accessible format. OpenSearch is RESTful technology and it's complementary to the OData. The OData interface is a data access protocol built on core protocols like HTTP and commonly accepted methodologies like REST that can be handled by a large set of client tools as simple as common web browsers, download-managers or computer programs.

FedEO is built around the concept of dataset series (collections). It helps to consider each dataset series (collection) as a set of metadata records. One such record can either describe another dataset series or an individual dataset (granule). Each dataset series is identified by an

doi: 10.2788/854791

identifier, which allows to address search requests to a particular dataset series and to receive metadata records as part of the search response [1]. The FedEO query interface is based on OpenSearch. In order to provide a well-defined search path from a collection of interest to granules associated with that collection, it is recommended to use the so called 2 step searching, which consists of a collection level search and the subsequent granule level search.

GPod contains information about aggregated acquisitions in either spatial or temporal format. It allows us to get pre-processed information for missions like Meris, Sentinel-1A and some others. It is internally stored in a directory like structure, where different directories represents information for different missions, sensors, acquisition modes and years. It provides the data in the .html, .svg, .png and the .dat formats. In our application we take the data from the .dat formats. For spatial coverage it offers the information in a format, where every line contain amount of the products for given bounding box.

2.2. Data visualization study

In the scope of this paper we discuss visualization of the metadata for ESA missions. It means amount of products taken in the time series, type of the products distinguished by the instrument and the mission. In order to deliver the information to the broad scale of users, it was necessary to find the way to visualize the data efficiently. With the respect towards the type of data, it is useful to display these information using map as the basis for display whenever possible, because map is efficient in enabling human users to understand complex situations [2].

Map can be understood as a tool to order information by their spatial context and therefore it is perfect interface between a human user and all that big data. Such map enables human user to answer location-related questions. In this case it means getting the punctual overview of the amount and distribution of ESA missions products around the world.

The proper choice of the cartographic method for visualization plays the key role for the right communication of spatial information. The choice depends on many aspects, most importantly on the type of spatial data, the volume of delivered information and the targeted group of users [3]. We describe the used methods below.

2.2.1. Visualizing the data in the 2D over the globe

We decided to use a special type of a density map, the choropleth grid, for displaying the amount of the product (e.g. total for all missions, total for one mission or total for one type of instrument). This method is based on division of the area to the same (or similar) geometric objects. In our case, the earth is divided into 1° x 1° rectangles in the most detailed resolution. This method allows us to compare different time observations due to the fact that the reference areas contain their shape and size. Furthermore, usage of this method allows us to easily display aggregated data in different levels of detail. This approach is convenient for the user, because he or she can easily move from a global perspective to regional and local details and so get both the general overview and the detailed information.

Another aspect we took into account in our visualization is the viewport. For instance when the user is looking on the Europe, he doesn't care about amount of products over the America. He is interested in seeing the density map relevant for Europe. We are therefore adapting our density map to currently visible area. This approach has two main advantages. First, only pieces of the data that the user is actually interested in are delivered, so the application is much faster and that's why the better user experience is provided. Second, a colour scale is always adjusted according to maximum and minimum value in current viewport and user is able to distinguish the local differences. Because the data which are displayed are quantitative, the colour scale is based on the intensity colour palette, where simply the more intense the colour is, the more products were acquired for given space.

2.2.2. Visualizing the data in the 3D over the globe

To display the contributions made by different missions and different instruments we used the Web World Wind extruded polygons functionality. The extruded polygons are in fact 3D summarizing diagrams.

Each diagram represents the total amount of the missions' products or instruments products acquired in the given area. Different colours are used for distinguishing the particular missions/instruments. Therefore using this method enables us to see the contribution of each mission or instrument.

At the same time there is only one additional information displayed on the top of the map. This means that user either displays the share of the missions on the total amount of product or the share of the instruments on the total amount of products and therefore the map is not visually overloaded. However, the User Interface provides an option to select between missions and instruments.

If the 3D representation isn't for the user for some reason suitable, it is possible to switch from 3D to 2D representation. The information about the distribution of the amount of products among missions and instruments is then displayed as a pie charts. The relative changes in the amount of products taken is displayed through the size of the circle. The sizes of the pies are relevant to the currently displayed area. Therefore the pies size is changing together with the size of the visible area.

doi: 10.2788/854791

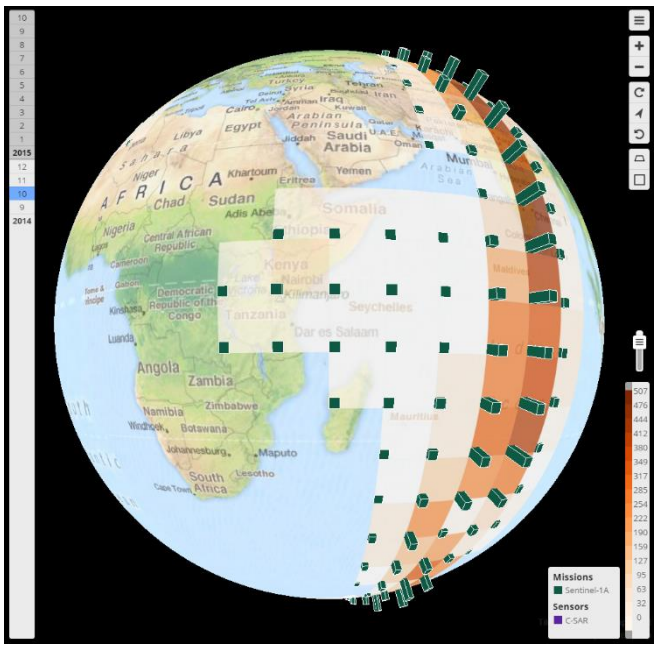

Image 1. Visualization of aggregated Sentinel-1A data acquired in October 2014.

Basically, we combine both of the methods above in our application. The 3D (or 2D) diagrams are typically displayed on the top of the choropleth map (Image 1). The user gets the information about total amount of the retrieved products in the area as well as about their structure.

3. RESULTS

To support our goals, we developed an application consisting of three parts. First is a java application for pre-processing data from different data sources and outputting them into different formats. It is possible to implement support for other data sources than the ones explained above. The same goes for the output formats.

The tool for pre-processing data prepares the files in the valid structure usable as an API if deployed anywhere on the server. At current moment we focus on the API with the .json files. In development version it is deployed at http://api.eoapps.eu/production/. The API contains missions.json file, which expresses available missions and sensors. The information about products acquired over month in certain area is available at the path {missionCode}/{instrumentCode}/{year}/{month}.json

This API also contains pre-aggregated data for different levels of detail representing the area of n * n degrees, where n represents the zoom level. These data are available on the path: /{zoomLevel}/{missionCode}/{instrumentCode}/{year}/{month}.json

In order to display the information over the globe, we decided to use WebWorldWind framework (https://webworldwind.org/), an open-source virtual globe solution. The WebWorldWind is a library and API rather than a stand-alone application. This enables it to be included in any web page or web application as a component. It allows us to visualize the information over the 3D globe instead of just visualizing it over the 2D map. It is very well optimized to work on the mobile as well as in desktop browser. It also supports handling of large amounts of data pretty well. This is crucial for us to display the polygons over half of the visible world. It also allows us to show 3D extruded polygons on top of the globe.

Our application was built in a responsive way allowing us to achieve very good results also on the mobile devices. There are two important parts of this effort. First part is optimizing the UI so that all important parts are visible. We tried to design a lightweight UI, where the key role plays the map and the other components are optimized for touch-screen devices. Second part is optimizing performance so that it runs adequately even on the mobile devices, which performance is still worse than that of desktops. We achieved that besides the other ways by processing the data just for currently visible area of the globe.

Our current implementation allows us to visualize the amount of products in the temporal perspective. We can show the user information about products acquired in single month or in multiple months. Our approach also allows the user to choose year for which he displays the data. In this way we present user the information in correct time based perspective. We also want the user to see the evolution of the amount of the data over time either by seeing the month by month information or by adding amount for months with play functionality.

The main asset of our solution is adapting the choropleth to the currently visible area in currently visible zoom level. By this we means that when user zooms in, zooms out or pan, the choropleth map will change and adapt to currently visible area. Therefore the application is faster because of much smaller amount of data loaded. Moreover, the user gets the information about distribution of the data in particular area, even if the area is not covered heavily.

The second most important asset is the fact that we take into the account the whole bounding box of acquired products unlike other solutions which take into the account just the centre of bounding box.

4. CONCLUSION

In this paper we propose how to visualise the ESA missions metadata (with the main focus on their amount and distribution) in a meaningful way in order to allow mission

managers, stakeholders and the general public to have a punctual overview. Our approach is based on data visualization in a way which allows to get both the general overview and the detailed information about amount of ESA missions' products and their distribution.

For this purpose we developed the java application for pre-processing data from different data sources and output them into different formats and the responsive, fast and easy-to-use web mapping application for data visualization which is built on top of open source technologies like WebWorldWind.

5. REFERENCES

[1] Coene, Y., *FedEO Client Partner Guide*. 1st ed. [eBook], European Space Agency, 2014.

[2] Gartner, G., "The relevance of Cartography," *Proceedings of GeoCart'2014 National Cartographic Conference*, New Zealand Cartographic Society Inc., Auckland, 2014.

[3] Vozenilek, V., *Metody tematické kartografie*, Univerzita Palackého, Olomouc, 2011.

doi: 10.2788/854791

ON THE EFFECT OF AIR POLLUTION ON CLINICAL RECORDS OF DIABETES-RELATED PATIENTS

Andrea Marinoni and Paolo Gamba

Dipartimento di Ingegneria Industriale e dell'Informazione, Università degli Studi di Pavia
Via Ferrata 5, I-27100, Pavia, Italy. E-mail: {andrea.marinoni, paolo.gamba}@unipv.it

ABSTRACT

The investigation and analysis of interactions between human phenomena and natural events has recently become a very interesting research field, as it can play a key role for several applications, from sustainable development to community policy design and short-, medium- and long-range resource allocation planning. In this paper, we provide a study of the interplay between air pollution (as estimated by remotely sensed data processing) and clinical records, so that inferences and correlations among black particulate concentration, micro- and macro-vascular disease onsets and hospitalization tracks can be efficiently drawn. We focused on the second order administrative area of the city of Pavia, Italy, from 2009 to 2014. Experimental results show how effective connections between the estimated air quality and the clinical records behavior can be accurately drawn and derived.

Index Terms— Big data, air pollution, exposomics, remote sensing.

1. INTRODUCTION

As recent developments in data collection have been achieved and implemented in social, financial, environmental and healthcare research, understanding and quantifying human-environment interaction (HEI) has become a big challenge in the spotlight of the scientific community [1]. The branch of medical research that is focused on this goal is called exposomics [2], [3]. In order to provide a complete characterization of the exposure behavior and interaction with the human health, exposomics requires a complete spatial-temporal description of the whole set of agents that are involved in the study. Thus, geo-referentiation of the patients is crucial to provide a reliable dataset over which clinical, medical and epidemiological records retrieving is possible [4], [5].

On the other hand, this constraint implies the environmental data to be collected with a very fine resolution over space and time. This requirement might be cumbersome to achieve if only ground probes are employed, as they can just give information on the environmental situation over a very restricted area[6]. Hence, Earth observation (EO) analysis can play a key-role within the exposomics framework. Indeed, the HEI assessment architecture can take advantage of the ability of remotely sensed images to properly provide reliable and effective information on wide spatial regions. Further, specific EO analysis techniques can deliver trustworthy estimations and evaluations on physical-chemical composition of the given scene, as air quality and water pollution ratios as well [1]. Actually, it is possible to infer particulate concentration from proper processing of thermal infrared records in remotely sensed data [6].

The city of Pavia, Italy, represents a valid test site for inferring correlations between clinical data and remotely sensed environmental characterization, in order to provide information to the healthcare community on interactions between environmental agents and disease evolution. Specifically, taking advantage of the clinical records collected by the Pavia local healthcare agency and Fondazione Salvatore Maugeri, Pavia, Italy, it is possible to develop the aforementioned analysis over consistent datasets, s.t. the information on undesired symptoms show-up and exposure outline over a given time period can be derived.

In this paper, we deliver the first attempt of drawing correlations and interactions between health records and air quality maps obtained from remotely sensed EOs. Specifically, air quality behavior over the year has been evaluated along with the information on glycated hemoglobin records that have been collected over the second order administrative area of Pavia, according to the clinical data collected within the Pavia local healthcare agency database. The outcomes of the proposed analysis showed how actually a first correlation between air quality and diabetic disease evolution can be drawn. The next Section introduces the description of the considered datasets and the experimental results that have been achieved. Section 3 reports the final remarks and next steps.

2. DATASET AND EXPERIMENTAL RESULTS

The experiments have been carried over a dataset that results from gathering together multispectral records and clinical data collected over the second order administrative area of Pavia, Italy, from 2009 to 2014. The second order administrative area of Pavia is spread over an area of 2968.64 km^2 in the north-western region of Italy. It counts 189 municipal-

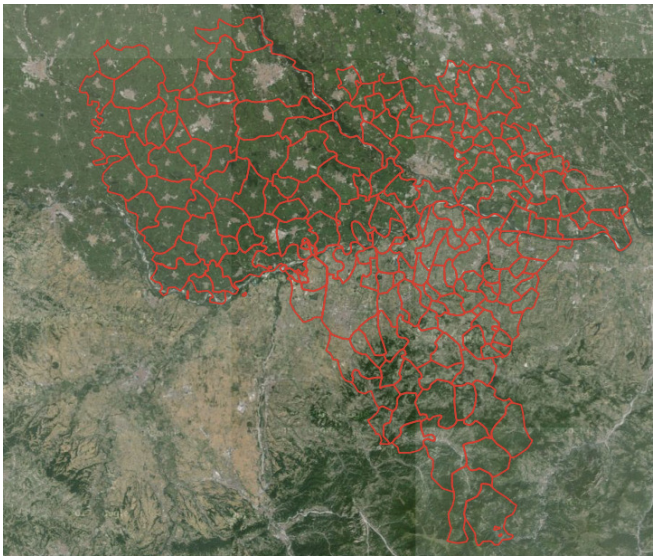

Fig. 1. Boundaries of all the municipalities in the Province of Pavia, Italy.

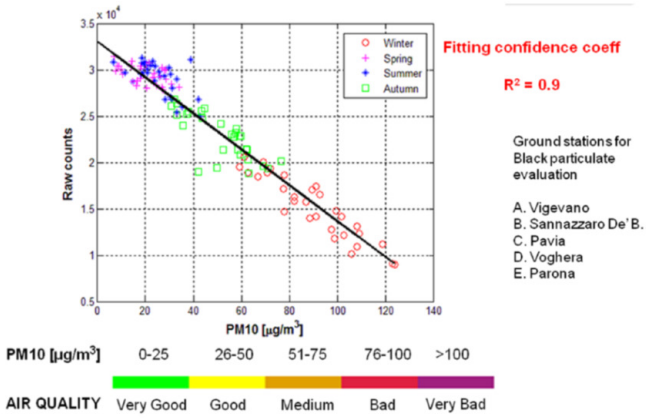

Fig. 2. Raw counts of the thermal infrared data of LandSat L8 images acquired over path 194 row 29 orbit on the 4 seasons in the 2009-2014 interval as a function of the black particulate PM10 concentration as recorded by the Pavia second order administrative area ground stations (located in Pavia, Voghera, Vigevano, Parona, and Sannazzaro de B.). Polynomial fitting between black particulate concentration and raw counts of the thermal infrared imagery is also reported: its confidence coefficient R^2 reached the value of 0.9.

ities which are grouped in 9 districts. We collected 35 Landsat images with cloud cover less than 5% in the above mentioned temporal interval over this test location, where each Landsat image consists of 2800 $\times$ 2800 pixels and has a 30m spatial resolution. Since it has been proven that the thermal infrared band of multispectral imagery can be inversely correlated to the quantity of black particulate concentrated over a given area [6], we can achieve relevant information on the human-environment interaction phenomena that are occurring in the Pavia area by looking for correlations between Landsat-based air pollution analyses and clinical data collected on the ground.

Specifically, in this experiment we considered a set of glycated hemoglobin and body mass indices (BMIs) as they have been collected out of a group of 1084 patients affected by diabetes and living in the area during the same time interval [7]. Those clinical records are time-referenced and georeferenced according to each municipality. In order to provide a statistically relevant dataset, glycated hemoglobin and BMI values were aggregated for each year. Further, each patient has been monitored through the 2009-2014 interval s.t. the records of every hospitalization, disease onset and check-up have been properly saved and considered.

The correlation between the presence of black particulate and the recorded temperature can be thoroughly characterized by radiance processes that affects the energy transmission throughout the atmosphere. Specifically, a pollution layer delivers a decrease of the atmospheric transmission factor. This effect impacts on the thermal infrared acquisition, since the solar heating is decreased as well. Consequently, the emitted radiance is lower, s.t. the signal recorded by the sensor is lower. Simultaneously, the pollution layer absorbs the emitted radiance, i.e., causes a strong impoverishment of the energy that is radiated upward. Hence, the aforementioned physical processes contribute to outline the correlation between the increase of pollution and the decrease of apparent temperature.

Several models have assessed the magnitude of these processes and their effects [8]. In order to achieve the air quality maps, the data acquired by LandSat L8 mission by means of Operational Land Imager (OLI) and Thermal Infrared Sensor (TIRS) have been considered. Collecting the data over the thermal infrared band and plotting them w.r.t. the black particulate concentration as reported by ground stations, it is apparent as the physical processes that affect the radiance transmission and absorption drive the correlation between air pollution and temperature records. Hence, taking into account the overall pattern of the raw counts thermal signal as a function of black particulate concentration, a polynomial fitting model has been implemented to estimate the air quality of the scene [6]. Specifically, for each temporal series we considered the thermal infrared signals acquired by Landsat 8 over the areas covered by the ground probes that have been taken into account. Then, we have been able to draw a polynomial relationship between black particulate concentrations and sensed reflectances. Hence, for each pixel in every temporal series, we apply the polynomial function resulting from the aforementioned fitting process to estimate the air quality, which has been quantized on five levels over the black particulate concentration estimate.

Specifically, for each temporal series we considered the

doi: 10.2788/854791

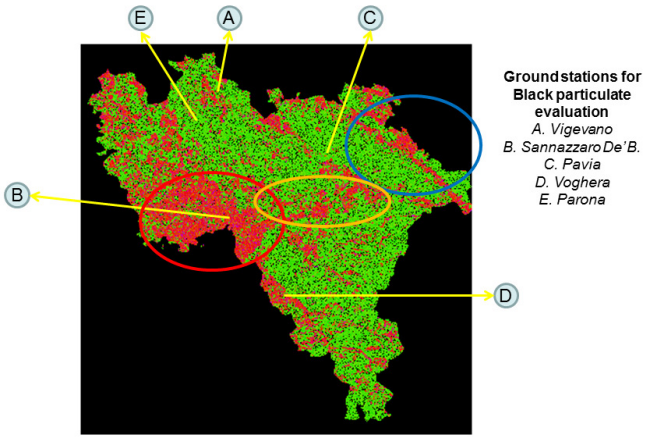

Fig. 3. Distribution of aggregate estimates over the 2009-2014 years of air pollution over the Province of Pavia, Italy. The same air quality color map in Fig. 2 applies here.

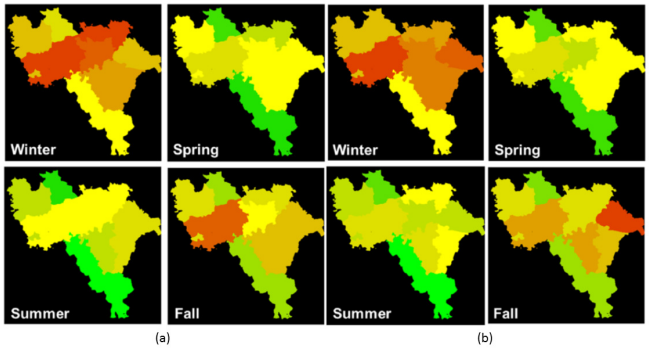

Fig. 4. Seasonal air quality maps for 2011(a) and 2012(b) estimated over the nine districts which compose the Province of Pavia, Italy. The same air quality color map in Fig. 2 applies here.

thermal infrared signals acquired by Landsat 8 over the areas covered by the ground probes that have been taken into account. Then, we have been able to draw a polynomial relationship between black particulate concentrations and sensed reflectances. Hence, for each pixel in every temporal series, we apply the polynomial function resulting from the aforementioned fitting process to estimate the air quality, which has been quantized on five levels over the black particulate concentration estimate.

Fig. 2 reports the thermal infrared raw counts of the images collected on the four seasons in the 2009-2014 interval as a function of the black particulate PM10 concentration as recorded by the aforementioned ground stations that are located in Pavia and 4 towns within its second order administrative area (ie, Voghera, Vigevano, Parona, Sannazzaro de B.). Such area is grouped for health care purposes into three districts, Pavese (including Pavia), Lomellina (comprising Vigevano, Parona and Sannazzaro de B.), and Oltrepo (including Voghera). Apparently, Fig. 2 shows how the black particulate concentration impacts on the recorded reflectance signals according to the aforementioned thermal infrared effects provided by air pollution. Thus, it is possible to infer air quality maps by properly processing the remotely sensed images according to the scheme that has been previously introduced. Specifically, we aim at characterizing the air quality of the given scene according to the quantization that is proposed at the bottom of Fig. 2. Finally, Fig. 2 displays also the polynomial fitting between black particulate concentration and raw counts provided by remotely sensed thermal infrared imagery as a solid black line. It is worth to note that the aforementioned fitting can be considered as very accurate and reliable, as its confidence coefficient R^2 reached the value of .9.

In order to appreciate the quality of the estimates we obtained according to the aforesiad procedure, Fig. 3 reports the distribution of the aggregate estimated air pollution over the Province of Pavia, Italy, through the 2009-2014 years. Specifically, it is possible to note how the estimated air pollution concentrate throughout the whole considered time span over particular areas. Indeed, it is possible to higlight how the air quality in the red, orange and blue boxes in Fig. 3 is dramatically worse w.r.t. to other regions over the Province of Pavia. On the other hand, the aforesaid boxes identify specific areas where particular anthropogenic extents and activities are located. In fact, the red box points the area around an oil refinery. Moreover, the blue box identifies the region across a chemical factory and a state highway. Finally, the three larger blops of high air pollution estimate within the orange box are located over two cheese factories and a concrete factory. Therefore, it is possible to understand how the aforementioned correlation between air pollution and thermal infrared contribution in multispectral images can be thoroughly motivated on ground.

In order to provide statistically reliable evaluations for the correlation inference analysis that had to be carried out over the glycated hemoglobin dataset, we managed to aggregate the air pollution estimates that have been achieved starting from the thermal infrared contribution investigation at a seasonal scale. Further, as the clinic data samples have been collected at a district level scale over the territory, we had to aggregate the measures over the spatial dimension as well. The resulting products are maps as those in Fig. 4, where the seasonal air quality estimates for the year 2011 (a) and 2012 (b) are reported. Then, it is possible to appreciate how the air quality behavior as estimated according to the proposed procedure can actually highlight the different scenarios occurring over the considered area. Further, according to the map in Fig. 3, it is possible to note how the contribution of the polluting anthropogenic extents affects larger areas.

Experimental results show how a strong correlation be-

tween air pollution as estimated by processing remotely sensed data and clinical records behavior through the considered time interval can be observed and extracted. Specifically, it is possible to infer how air quality actually affects the health of patients who suffer of diabetes-related diseases. Indeed, the pattern of hospitalizations of those patients can be accurately described according to the estimated distribution of air pollutants.

Further, glycated hemoglobin behavior is strongly coupled to the air quality estimate distribution, s.t. the onsets for microvascular diseases can be understood and explained by the aforementioned interplay (see Fig. 5). Actually, the correlation factor between air pollution estimates and glycated hemoglobin is about 91%. Specifically, four districts showed a very strong deviation from the mean values that have been computed along the clinical records. Indeed, the p-values resulting from mixed model analysis in order to extract the fluctuations of the data values through the clinical dataset showed how the behavior of high relevance samples (i.e., records which follow particular paths w.r.t. the mean distribution of the glycated hemoglobin values) is supported by the air quality maps that have been estimated.

Therefore, the correlation that has been achieved makes the proposed framework actually suitable to understand and quantify the effective impact of air pollution on human health. Hence, further interactions can be characterized and detailed by including large datasets of heterogeneous sources within proper exposomics architecture.

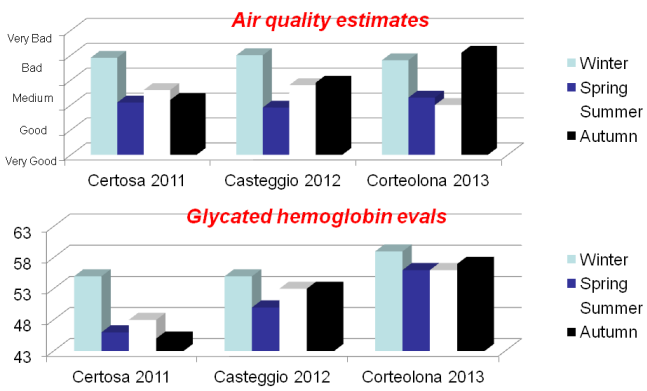

Fig. 5. Air pollution as estimated by remotely sensed data and glycated hemoglobin behaviors over three towns within the second order administrative area of Pavia, Italy.

3. CONCLUSIONS

A novel framework for integration of heterogeneous datasets collecting clinical, administrative, environmental and remote sensing records for assessing the impact of air quality on human health is introduced. The area of the Province of the city of Pavia, Italy, is used as test site as time- and geo-referenced data have been acquired through the 2009-2014 years time span. Relevant correlation between the air pollution maps as estimated by properly processing multispectral image stacks and the glycated hemoglobin behavior of a cohort of diabetes-related patients has been discovered. This result provides an interesting point-of-view for thoroughly investigating the effect of air quality on human health. Future works can be dedicated to explore the actual human-environment interplay in different scenarios.

4. REFERENCES

[1] A. Marinoni and P. Gamba, "Big data for human-environment interaction assessment: challenges and opportunities," in *Proc. of 2014 ESA Big Data from Space Conference (BiDS)*, Frascati, Italy, Nov. 2014.

[2] S. Martin, "Exposome informatics: considerations for the design of future biomedical research information systems," *J Am Med Inform Assoc.*, vol. 21, no. 3, pp. 386–390, 2014.

[3] S. Allender, C. Foster, L. Hutchinson, and C. Arambepola, "Quantification of Urbanization in Relation to Chronic Diseases in Developing Countries: A Systematic Review," *Journal of Urban Health: Bulletin of the New York Academy of Medicine*, vol. 85, no. 6, pp. 938–951, 2008.

[4] M. Faber, N. Niemes, and G. Stephan, *Entropy, environment and resources*, Berlin, D: Springer Verlag, 2nd edition, 1995.

[5] D. Pimentel, C. Harvey, P. Resosudarmo, K. Sinclair, D. Kurz, M. McNair, S. Crist, L. Shpritz, L. Fitton, R. Saffouri, and R. Blair, "Environmental and economic costs of soil erosion and conservation benefits," *Science*, vol. 267, pp. 1117–1123, February 1995.

[6] L. Wald and J.-M. Baleynaud, "Observing air quality over the city of nantes by means of landsat thermal infrared data," *Int. J. of Remote Sensing*, vol. 20, no. 5, pp. 947–959, 1999.

[7] A. Dagliati, A. Marinoni, C. Cerra, P. Decata, L. Chiovato, P. Gamba, and R. Bellazzi, "Integration of Administrative, Clinical, and Environmental Data to Support the Management of Type 2 Diabetes Mellitus: From Satellites to Clinical Care," *SAGE Journal of Diabetes Science and Technology*, vol. 10, no. 1, pp. 19–26, 2016.

[8] L. Wald, L. Basly, and J.-M. Baleynaud, "Satellite data for the air pollution mapping," in *18th EARSeL Symposium on Operational Remote Sensing for Sustainable Development*, Enschede, The Netherlands, 1998.

doi: 10.2788/854791

USING ELK STACK TO INDEX, ANALYZE AND VISUALIZE GAIA DPCT INFRASTRUCTURE LOG DATA

Andrea Fonti[1], Ruben De March[1], Rosario Messineo[1], Angelo Fabio Mulone[1], Fabio Filippi[1], Fabio Sella[1], Carmelo Manetta[1]

[1]ALTEC S.p.A. Corso Marche 79, 10146 Torino

ABSTRACT

We describe a software stack enabling indexing, analysis and visualization of log data collected from the Gaia DPCT infrastructure. Exploiting the described system the DPCT staff is able to gather meaningful insights about operational platform status, presented as real time updating dashboards, while being able to perform tasks of greater complexity like anomaly detection or post failure analysis leveraging a huge database of diagnostic data.

The presented system leverages big data technologies making such a large scale diagnostic analysis feasible. Different data sources are exploited in order to perform the anomaly detection: SNMP metrics, operational database, application log data, network and SAN data.

Algorithms for anomaly detection successfully deployed to the described system will be presented as well.

Index Terms — ELK, Elasticsearch, Logstash, Kibana, logs, time series, anomaly detection, monitoring, Gaia, DPCT, TECSEL2

1. INTRODUCTION

Management of the Gaia DPCT infrastructure (a large scale science processing system) poses a number of different challenges in log data and metrics analysis due to the scale of the data collected from system monitoring, application logs and science DB lying in the TB per month range. Analyzing such a large data set in order to prevent failures or to diagnose the source of an already fixed failure presents multiple challenges in:

1. **ETL**: in order to integrate data coming from different monitoring sources and with wildly varying structure;
2. **Indexing**: in order to reduce the time took to explore, aggregate or retrieve data;
3. **Visualization**: enabling the discovery of statistical trends or efficient information presentation via dashboards;

These three points would be already difficult to deal while storing small scale data but as the system grows toward a data processing centre for the Gaia science mission the conventional systems are not able to respond as quickly as needed, exploration of the data-set should be always viable in near real-time, as well as a subset of the visualization requirements. To solve the scale-related challenges a stack employing big data technologies was deployed and integrated into the DPCT infrastructure.

2. INDEXED DATA

Data that will be used in the analysis process should be transformed and indexed; three major data sources are subject to this process.

2.1. Log data

Log data is shipped from the processing cluster to the ETL processing host as structured documents; log data is not flattened as lines and thus can be transformed field by field (i.e. the Mapped Diagnostic Context is shipped as a series of key-value pairs), fields are then parsed to generate other structured data or renamed to have a uniform naming scheme across different data sources (i.e. from stack trace field the exception type is extracted to a new field).

2.2. SNMP metrics

Hosts participating into the processing cluster are queried at regular intervals via SNMP protocol to gather metrics. Stored metrics include: load average, CPU load, memory consumption, mount point metrics, network throughput metrics.

2.3. Distributed application progress data

By querying the database used by the cluster to publish processing jobs a snapshot of the ongoing processes can be extracted. These data can be correlated with metrics and log data.

3. ELK STACK

ELK is a software stack comprising Elasticsearch, Logstash and Kibana. It is aimed at indexing, ETL and visualization of data; the deployed service topology is described in Fig. 1.

doi: 10.2788/854791

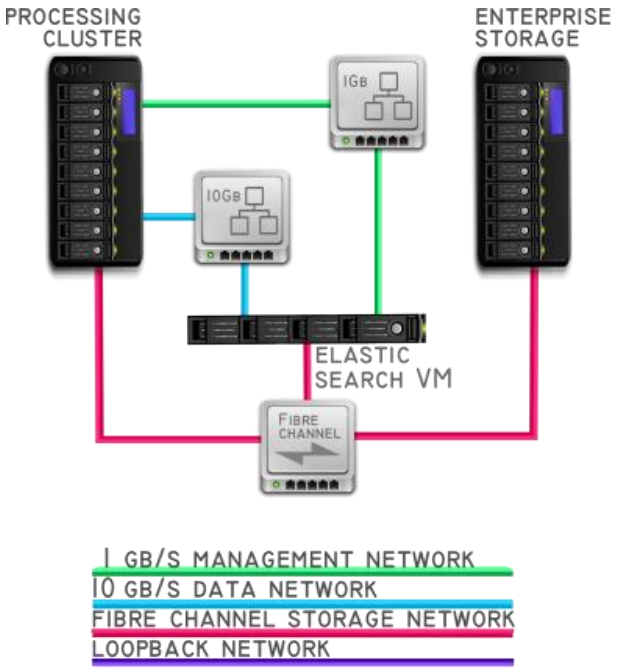

Fig. 1. Deployed service topology

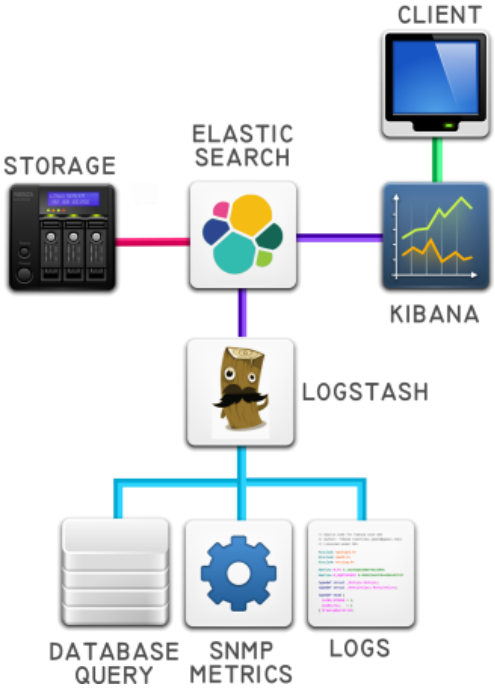

Fig. 2. Indexed data flow

3.1. Elasticsearch

Elasticsearch is a search server based on Lucene [1]. It provides a distributed, multitenant-capable full-text search engine with a RESTful web interface and schema-free JSON documents. Elasticsearch is developed in Java and is released as open source under the terms of the Apache License. It provides scalable search, has near real-time search, and supports multitenancy. Elasticsearch is distributed, which means that indices can be divided into shards and each shard can have zero or more replicas. Each node hosts one or more shards, and acts as a coordinator to delegate operations to the correct shard(s). Rebalancing and routing are done automatically.

3.2. Logstash

Logstash is a tool to collect, process, and forward events and log messages thus solving the ETL problem [2].

Collection is accomplished via configurable input plugins including raw socket/packet communication, file tailing, and several message bus clients.

Once an input plugin has collected data it can be processed by any number of filters which modify and annotate the event data adapting its schema. Finally Logstash routes events to output plugins which can forward the events to a variety of external programs including Elasticsearch, local files and several message bus implementations.

3.3. Kibana

Kibana is an open source analytics and visualization platform designed to work with Elasticsearch [3]. Kibana is used to search, view, and interact with data stored in Elasticsearch indices. Advanced data analysis and visualization are available via a variety of charts, tables, and maps.

Kibana makes it easy to understand large volumes of data via a simple browser-based interface, creating and sharing dynamic dashboards that display changes to Elasticsearch queries in real time.

doi: 10.2788/854791

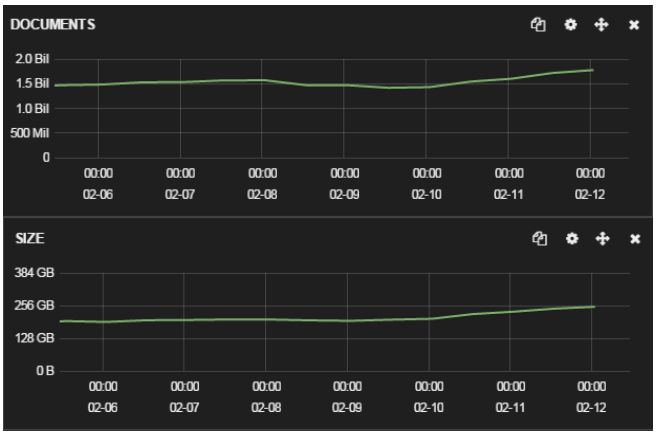

Fig. 3. Number of indexed log events and datastore size over a 7 days period

4. ANALYSIS METHODS

The mining of data available in log files has been widely investigated and many applications of interest have been identified [4].

Our log analyses are mainly focused on finding relations and co-occurrence between different terms (i.e., words or group of words) but also on detecting unusual behaviors. This work is already made by Kibana in a certain way, but we can also use its dashboard to drive our further investigations. For example, we can use Kibana to focus our attention on a particular time range or on some particular issue, thus filtering among all the possible terms.

A simple method to detect term co-occurrence is computing mutual information measures as done by B.J. Jansen when examining search logs with term level analysis [5]. After this computation, we obtain a large amount of values for every unordered pair of terms considered. This could be enough to check which are the most recurring paired terms, but what if such terms come in a group of three, four, n? We could generalize the mutual information measurement, but maybe it is better to use some kind of topological approach.

4.1. The topological approach

After the preliminary step of taking mutual information measures between the wanted pairs of terms, we obtain a list where every element can be seen as two distinct vertices along with the weigth of the edge between them. Therefore a weighted graph can be constructed using this data.

After having put in descending order the edges by their weights, we can think of adding in this order the edges to the graph, one by one. This operation creates a sequence of structures, equivalent to *simplicial complexes*, where the previous one is always included in the following one.

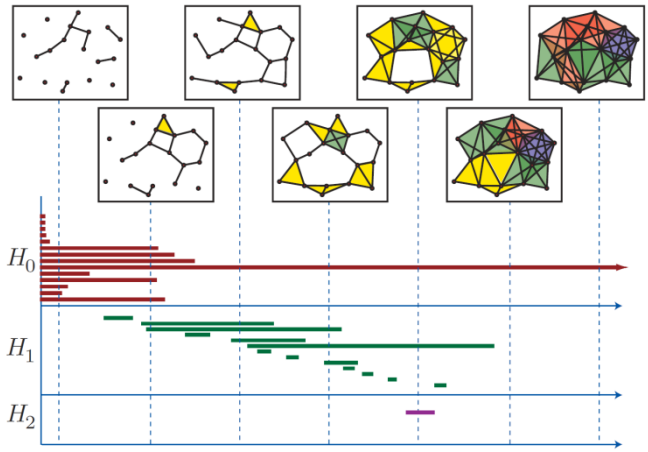

Fig. 4. Barcode diagram of a simplical complex constructed on a weighted graph. H_k is the *k-th* homology group and its rank equals the number of intervals intersecting the dashed line at a specific time. You can see that the rank of H_0 is the number of connected components, the rank of H_1 is the number of 1-dimension holes, and so on. This plot is taken from [8].

Given such a chain of complexes, one can compute the *homology groups* in order to reveal the structures of the dataset of chosen terms. Moreover, these topological structures can persist for a certain number of "edge additions", and the more persistent ones are the ones the dataset's skeleton consists of. The definitive selection, however, is decided by the data scientist, in an analogue way of the "cutting" in a hierachical clustering tree. This is what *persistent homology* stands for. The reader can find more details and the mathematical explanation in [5], [6], [7], [8].

As described in [8], this "skeleton" can be detected through specific diagrams, called *barcodes*. If we think as the edges addition as time passing by, seeing in the diagram a long time range where the barcodes are stationary means that the topological structures in that period are persistent, thus important for the structure of the dataset. An example of this can be seen in Fig. 4.

Barcodes have been computed with the help of dedicated Python libraries ([9]).

4.2 The time series approach

Another approach is generating time series datasets with our logs. We track the relative frequency (i.e., the number of occurrences of a term normalized by the total occurrence) for each term during a certain interval of time (e.g., 5

doi: 10.2788/854791

minutes). This leads to 5-minutes-resolution time series which can be compared between each other to find out high correlation or even high *partial* correlation, like it has been done on another DPCT project, TECSEL2 [10].

Furthermore, TECSEL2 tries to infer cause-and-effect delay between different time series. Log analysis could benefit from this feature too, if, for example, some (logged) events happened earlier introduce some issue that causes some effects later on. This could mean importing our logs time series data inside TECSEL2 architecture, to exploit the features and performance of the Cassandra database connected to the Spark framework.

Analyses input data can also be pre-processed in order to evaluate only terms in a certain frequency range or the components selected through dimensionality reduction techniques like Principal Component Analysis (PCA) or Independent Component Analysis (ICA). Another possible selection is the one that takes advantage from the topological analysis explained above to group some terms with others and then treat every group as a distinct series.

Real-time processing is also performed, to point out if frequencies of some terms are different from chosen or nominal ones.

5. CONCLUSIONS

Implementation of the described system fullfilled the goal of responding to infrastructural and applicative failure as quickly as possible to guarantee availability of Gaia DPCT science operations. The system is able to cope with the cardinality of the ingested datasets.

All the input data and output can then be a precious source of information for further analyses. In this paper we presented a solution which uses computational topology and another one oriented on time series analysis. These analyses are both hosted at DPCT on dedicated machines that use big data technologies, that constitute the state of art for the treatment of huge amount of data like these.

6. REFERENCES

[1] C. Gormley and Z. Tong, Elasticsearch: The Definitive Guide, O'Reilly Media, 2015.

[2] J. Turnbull, The Logstash Book, James Turnbull, 2015.

[3] Elasticsearch BV, "Kibana User Guide," Elasticsearch BV, 2 November 2015. [Online]. Available: https://www.elastic.co/guide/en/kibana/4.1/index.html. [Accessed 2 November 2015].

[4] A. Chuvakin, "Log Data Mining," 2003-2005. [Online]. Available: http://www.slideshare.net/anton_chuvakin/log-data-mining. [Accessed 2 11 2015].

[5] A. Zomorodian, "Computational Topology," in *Algorithms and Theory of Computation Handbook*, 2 ed., 2009.

[6] R. Ghrist, "three examples of applied & computational homology," [Online]. Available: https://www.math.upenn.edu/~ghrist/preprints/nieuwarchief.pdfg2=CH-gStuTj6. [Accessed 9 February 2016].

[7] R. J. Adler, O. Bobrowski, B. S. Matthew, E. Subag and S. Weinberger, "Persistent Homology for Random Fields and Complexes," *Borrowing Strength: Theory Powering Applications, A. Festschrift for Lawrence D. Brown, IMS Collections 6*, pp. 124-143, 2010.

[8] R. Ghrist, "Barcodes: the persistent topology of data," *BULLETIN (New Series) OF THE AMERICAN MATHEMATICAL SOCIETY*, vol. 45, no. 1, pp. 61-75, January 2008.

[9] G. Petri, "Holes: A Python module for detecting, analysing and visualising persistent homological features of complex networks," [Online]. Available: https://github.com/lordgrilo/Holes. [Accessed 2 November 2015].

[10] R. De March, D. Busonero, R. Messineo, A. Bemporad., F. Vaccarino, A.F. Mulone, A. Fonti, and M.G. Lattanzi, "TEmporal Characterization of the remote SEnsors response to radiation damage in L2," *Proc. Big Data from Space (BiDS16), IEEE*, 2016.

[11] B. Jensen, Handbook of Research on Web Log Analysis, IGI Global, 2008.

doi: 10.2788/854791

NANSAT - A SCIENTIST FRIENDLY PYTHON TOOLBOX FOR PROCESSING 2D SATELLITE EO AND MODEL RASTER DATA

Aleksander Vines, Anton Korosov, Morten W. Hansen, Asuka Yamakawa, Tor Olaussen, Torill Hamre

Nansen Environmental and Remote Sensing Center
Thormøhlens gate 47,
NO-5006 Bergen,
Norway

ABSTRACT

NANSAT [1], the **Nan**sen Center **Sat**ellite data processing toolbox, is an open source framework for processing 2D satellite EO and model raster data.

The main goal of NANSAT is to facilitate easy development and testing of scientific algorithms, easy analysis of geospatial data, and efficient operational processing.

The main functions of NANSAT are to read and export geospatial data, to assign discovery and use metadata to the datasets and to implement basic operations, such as reprojection and subsampling, and a few pixel functions, e.g. for calculating incident angles. In addition, it is designed to be easy to extend, allowing users to add new scientific algorithms or support for new data formats.

NANSAT supports a large number of different satellites/sensors out of the box, among others, data from Envisat, Radarsat and Sentinel-1, as well as model datasets stored in NetCDF [2] format.

Index Terms— Earth Observation, Data Analysis, Processing, Visualization, Geospatial Data Processing, Scientist Friendly Framework, Python, Open Source, GDAL

1. INTRODUCTION

One of the driving forces behind the NANSAT development was to provide a framework which enables non-programmers, such as students and scientists, to process geospatial data. Basic functions, e.g., reprojection, subsampling and extracting transects, are available out of the box and new functionality should be easy to add.

Efforts were made to create such a framework in Matlab. The result was a complex system, which in itself makes it harder to implement new functionality and the system more prone to bugs [3]. In addition, Matlab is not free, but requires an expensive license. Python is free and open-source. Python is an object-oriented language which allows a clear modular structure and makes it easy to maintain and modify existing code. Moreover, scientific Python libraries, e.g., Numpy [4] and SciPy [5], are available for free. Consequently, Python was adopted and the development of NANSAT started in 2011.

To enable the use of several data sources for co-location, comparison and validation, NANSAT takes benefit of the already existing open source Geospatial Abstraction Data Library (GDAL) C/C++ library [6]. This is used to read geospatial data from a number of different sources and to perform basic operations on the data. A core feature of the system is that data is read from a file only when needed, and it is possible to extract subsets of the full dataset. This is more efficient than reading the full dataset into memory, as is the case with, e.g., Matlab, and in particular if it is stored on a remote server and accessed via, e.g., OPeNDAP [7].

NANSAT provides high flexibility in a command-line interface and a Python API to achieve its goal. NANSAT can be run interactively from within the Python interpreter, or included in your own Python code. When a workflow has been established, the Python script can be called from a shell script or a cron job.

2. SYSTEM DESIGN

Figure 1 shows an overview of the system concept. NANSAT is built around GDAL, extends it and employs standard controlled metadata vocabularies, e.g. Climate and Forecast [8] (CF) conventions and Global Change Master Directory (GCMD) keywords [9].

To complement the retrieved raster data, a package called mappers adds meaning to the data as metadata. Presently, NANSAT supports over 40 types of dataformat, different satellite formats and numerical models and any gridded datasets following the CF standard for NetCDF files.

The *Nansat* class is the core of NANSAT and represents a dataset.A *Nansat* object contains metadata on the geographical reference of the data and the parameters/bands with geophysical variables. It also contains a reference to a virtual GDAL vrt XML file. The vrt file is a metadata file with instructions for GDAL on how to read and interpret a datafile. Some basic operations in NANSAT, like reprojecting or crop-

doi: 10.2788/854791

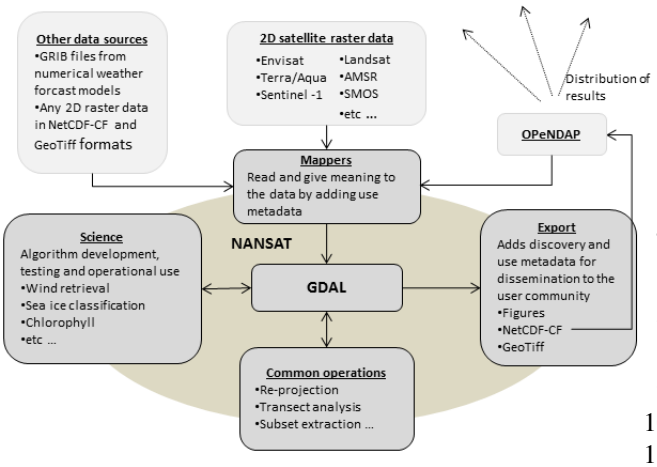

Fig. 1. The NANSAT concept.

ping data, does not modify the underlying data itself, but modifies this vrt file. When a vrt file is modified, a reference to the previous version is retained. This allows to easily undo operations on the data. All high-level operations are performed on the *Nansat* object, e,g, reprojecting, cropping and exporting.

A basic idea behind NANSAT is that it should be easy to add new mappers. Users can easily develop and add their own mappers if data formats are not supported by NANSAT.

When we open a file with NANSAT (within Python):

```
n = Nansat(filename)
```

these steps follow:

- The constructor of NANSAT calls gdal.Open(uri) to open the file/url with GDAL, and returns a GDAL dataset with a list of available raster bands
- It loops through available mappers and parses the dataset to the mapper
- Each mapper checks if the input dataset is appropriate for the mapper, i.e. if the format, the metadata and the set of bands in the dataset correspond to the mapper. If the dataset is not valid, the mapper silently fails and the next mapper is tested
- If the mapper fits the dataset: the mapper creates a GDAL VRT file which has Raster Bands corresponding to the variables and adds respective metadata to each band (standard_name, units, etc).
- The mapper opens the VRT file with gdal.Open() and returns this GDAL dataset back to Nansat.

3. USAGE

A set of functions to perform common operations on the data, like georeferencing and reprojection, averaging, transect extraction and analysis, as well as visualization and generation of high quality maps, is available. In addition to this, NANSAT contains a set of common export functions.

```
1  n = Nansat("MER_FRS_1PNPDK_example.N1")
2  n.write_map('Meris-map.png')
3  n.write_figure('Meris-band-1.png',
4                 bands=['L_413'], clim='hist')
5  dLatlong = Domain(
6      4326,
7      "-te 17 62 23 65 -ts 2000 1000")
8  n.reproject(dLatlong)
9  n.write_map('Meris-reproj-map.png')
10 n.write_figure('Meris-reproj-band-1.png',
11                bands=['L_413'], clim='hist')
```

In the above example, a MERIS L1 image is opened with NANSAT(line 1) and then reprojected and cropped(line 8) into a specific domain(lines 5-7) with longitude and latitude coordinates. A bounding box of the domain, both before(line 2) and after(line 9) the reprojection, is exported to an image with a crude map as shown in Figure 2. In Figure 3, the two output images for the first band (lines 3-4 & 10-11) are displayed. This band contains the top of atmosphere (TOA) spectral radiance at 413 nm wavelength.

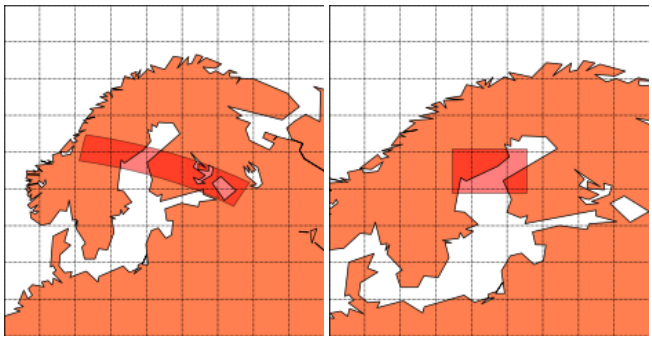

Fig. 2. The domain on a map before and after reprojection

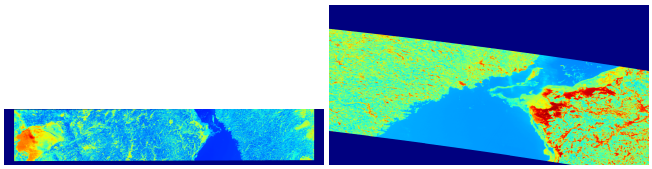

Fig. 3. MERIS L1 413nm band before and after reprojection

A more advanced output example is displayed in Figure 4. Here, the sea surface salinity(color), derived from MODIS data, is combined with surface currents(streamlines), from TOPAZ reanalysis data. The gray areas represent land and the white areas represent clouds.

doi: 10.2788/854791

```
from matplotlib.colors import LogNorm
from nansat import *
# Access local SMOS data downloaded
# from http://cp34-bec.cmima.csic.es/
n_smos = Nansat('BEC_OI_20100811.nc')
# Access local TOPAZ data
n_topaz = Nansat('TP4DAILY_2010_222.nc')
# Define region of interest near Amazon esturay
dstDomain = Domain(
  NSR().wkt, '-te -60 -6 -35 15 -tr 0.1 0.1')
lon, lat = dstDomain.get_geolocation_grids()
# Co-locate datasets on the same grid
n_smos.reproject(dstDomain)
n_topaz.reproject(dstDomain)
# Fetch equal-size arrays of data
# Fetch sea surface salinity from SMOS-BEC
sss = n_smos['SSS']
# Fetch eastward and northward components
# of surface current from TOPAZ
u = n_topaz['utot01']
v = n_topaz['vtot01']
# Create map comparing SSS from
# MODIS and surface currents from TOPAZ
# create map canvas
nmap = Nansatmap(dstDomain, resolution='h')
# show MODIS SSS distribution
nmap.imshow(sss, vmin=32.5, vmax=35.5)
# plot streamlines of surface current
nmap.streamplot(lon, lat, u, v, 3,
                linewidth=np.hypot(u, v),
                color='k')
# add colorbar and save the plot
nmap.add_colorbar(shrink=0.7)
nmap.drawgrid()
nmap.save('test.png', dpi=300)
```

4. INSTALLATION

NANSAT depends on multiple external, free and open source libraries, e.g., GDAL, Numpy, Scipy, Matplotlib, py-thesaurus-interface and netcdf4.

GDAL could be cumbersome to install if you want to include all drivers for full functionality of NANSAT(hdf4, hdf5, netcdf, dap, geos, jasper). Therefore, we have pre-built fully functional binaries of GDAL for Linux(32 and 64 bits) which are available at the nersc repository on the Anaconda Cloud [10].

In order to provide a fully compatible environment, we also provide a Vagrant [11]/Ansible [12] setup [13], which allows the user to set up and provision a virtual machine configured with all the dependencies.

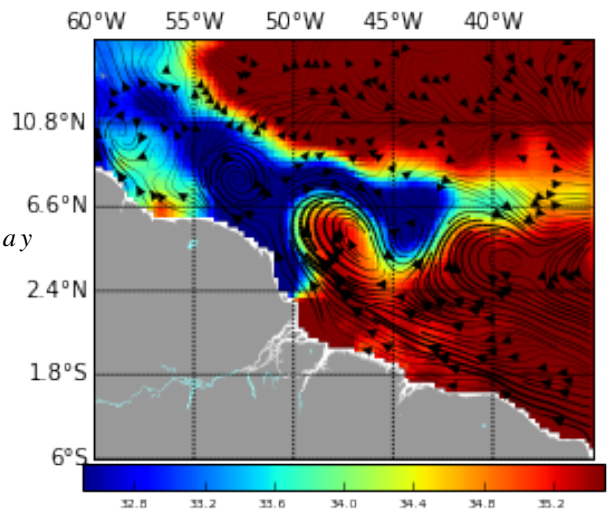

Fig. 4. Sea surface salinity and surface currents

5. USE CASES

NANSAT is widely used by several researchers and students in Norway, Russia, China, India, South Africa, France, USA, and Germany and in projects , like NORMAP [14] and MyOcean/Copernicus Marine Environment Monitoring Service [15], for instance to generate sea ice, ocean color, wind and ADT products.

An example workflow using the Support vector machines (SVM) ice classification algorithm [16] is shown in Figure 5. Nansat is used throughout this chain for basic tasks, e.g. reading data, calculating incident angles and removing noise using pixelfunctions, and exporting images.

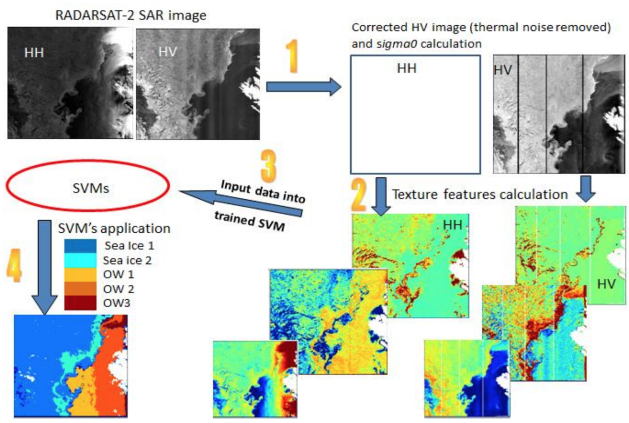

Fig. 5. Basic flow diagram on SVM ice classification

NANSAT is also used in a process chain in the Metvind project to provide an operational weather/wind forecast and verification service with synthetic aperture radars and in situ measurements. This project is in collaboration with the Nor-

doi: 10.2788/854791

wegian Meteorological Institute and the service will be available by the end of 2016.

Nansen-Cloud [17] has been an internal project at the Nansen Environmental and Remote Sensing Center where we have developed a Scientific Platform as a Service. NANSAT is one of the key elements in this platform, and is used to read the datasets in the pilot version.

6. SUMMARY AND FUTURE WORK

NANSAT has been in use for several years, mostly for processing of satellite images. NANSAT has been our preferred interface towards satellite data, and we are constantly working on improving it, extending NANSAT to support new dataformats and new functionality, e.g. streaming data through the OPeNDAP protocol.

The future lies in Big Data and we have noticed the need for combining satellite data with model and in-situ data. NANSAT is being incorporated in a scientific platform in Nansen-Cloud. Through this platform, we will support a large variety of data, where more specialised interfaces to in-situ and model data will be combined with the satellite data. It is designed to easily search and collocate datasets for scientists to be able to navigate the large volumes of data that are available today.

Future plans include further work on NANSAT and Nansen-Cloud to better be able to handle the large amounts of data that scientists will need. On the NANSAT development side we want to:

- Improve the test coverage to further ensure the quality of the product. Currently the unit-tests cover 56% of the code. Integration tests covers a lot more (not measured). We aim at 90% coverage on the unit-testing.
- Speed up and optimise code for reprojected data.
- Speed up and optimise code for reading generic streamed data through OPeNDAP.

7. ACKNOWLEDGEMENTS

NANSAT was developed as part of the NORMAP [14] project, funded by the Research Council of Norway grant no.195397.

We would like to thank the open source communities for all the packages that are used in NANSAT.

8. REFERENCES

[1] Anton Korosov, Morten W. Hansen, Asuka Yamakawa, Knut-Frode Dagestad, Aleksander Vines, Maik Riechert, and Alexander Myasoedov, "nansat: Nansat-0.6.8," http://dx.doi.org/10.5281/zenodo.45703, Feb. 2016.

[2] R. Rew and G. Davis, "NetCDF: an interface for scientific data access," *Computer Graphics and Applications, IEEE*, vol. 10, 2002.

[3] Robert C. Martin, *Clean Code: A Handbook of Agile Software Craftsmanship*, Prentice Hall, 2008.

[4] S. Van Der Walt, C. Colbert, and G. Varoquaux, "The NumPy Array: A Structure for Efficient Numerical Computation," *Computing in Science and Engineering*, vol. 13, no. 2, pp. 22–30, 2011.

[5] E. Jones, T. Oliphant, and P. Peterson, "Scipy: Open source scientific tools for Python," http://www.scipy.org, 2001.

[6] GDAL Development Team, "GDAL - Geospatial Data Abstraction Library, Version 2.0.0," http://www.gdal.org/, 2015.

[7] Peter Cornillon, James Gallagher, and Tom Sgorous, "OPeNDAP: Access data in a distributed, heterogeneous environment," *Data Science Journal*, 2003.

[8] "CF Conventions and Metadata," http://cfconventions.org/, 2015.

[9] Global Change Master Directory (GCMD), "Gcmd keywords, version 8.1," http://gcmd.nasa.gov/learn/keywords.html, 2015.

[10] "Anaconda Cloud," https://anaconda.org, 2016.

[11] "Vagrant by HashiCorp," https://www.vagrantup.com/, 2016.

[12] "Ansible is Simple IT Automation," http://www.ansible.com/, 2016.

[13] "Nansen-Cloud provisioning with vagrant," https://github.com/nansencenter/nersc-vagrant/, 2016.

[14] The NORwegian Satellite Earth Observation Database for MArine and Polar Research, "NORMAP," http://normap.nersc.no/, 2015.

[15] "Copernicus Marine Environment Monitoring Service," http://marine.copernicus.eu/, 2015.

[16] A. Korosov, N. Zakhvatkina, and S. Muckenhuber, "Ice/water Classification of Sentinel-1 Images," European Geophysical Union General Assembly, 17 Apr 2015, Vienna, Austria.

[17] Morten Wergeland Hansen, Anton Korosov, Aleksander Vines, Tor Olaussen, and Torill Hamre, "The Nansen-Cloud: a Scientific Platform as a Service for the Nansen Group," in *Proceedings of BiDS'16*, 2016.

doi: 10.2788/854791

BROADVIEW RADAR ALTIMETRY TOOLBOX

R. Escolà[1], A. Garcia-Mondéjar[1], G. Moyano[1], M. Roca[1]
M. Terra-Homem[2], A. Friaças[2], F. Martinho[2], E. Schrama[3], M. Naeije[3],
A. Ambrózio[4], M. Restano[5], J. Benveniste[6]

1 isardSAT Ltd., 2 DEIMOS Engenharia, 3 TU Delft,
4 DEIMOS/ESRIN, 5 SERCO/ESRIN, 6 ESA/ESRIN

ABSTRACT

The universal altimetry toolbox, BRAT (Broadview Radar Altimetry Toolbox), already capable of dealing with past and current altimetry missions' data, incorporates now the capability of reading the upcoming Sentinel-3 L1 and L2 products. ESA promotes and supplies this new feature, supporting users of the future Sentinel-3 SAR Altimetry Mission.

***Index Terms*—** BRAT, toolbox, altimetry, radar, Sentinel-3

1. INTRODUCTION

The BRAT project started in 2005 from the joint efforts of ESA (European Space Agency) and CNES (Centre National d'Études Spatiales). The project aimed at creating a collection of tools and tutorial documents designed to facilitate the processing of radar altimetry data.

2. THE TOOLBOX

The toolbox enables users to interact with the most common altimetry data formats. Moreover, BRAT can be used in conjunction with MATLAB/IDL (via reading routines) or in C/C++/Fortran/Python (via an API). This allows for the desired data to be obtained and manipulated in the environment with which the user is more familiarized, bypassing the data-formatting hassle. It also allows for fast visualisation, data conversion into other formats such as NetCDF, ASCII text files, KML (Google Earth), and plotting into raster images (JPEG, PNG, etc.).

A wide range of computations can be done in BRAT, from the application of built-in formulas available within the tool's library to user defined ones, which can be saved for future purposes.

The BRAT Graphical User Interface (GUI) is the front-end for the powerful command line tools that are also part of the BRAT suite.

2.1 The BRAT Graphical User Interface

A new GUI is currently being built, based on the feedback from the community, which was obtained through workshops. The current iteration is at its beta phase, of which some results can be seen in Figures 1, 2 and 3.
The main drive behind the new UI developments is an increase in usability: making BRAT easier to be used, more intuitive, and ever closer to the community needs.

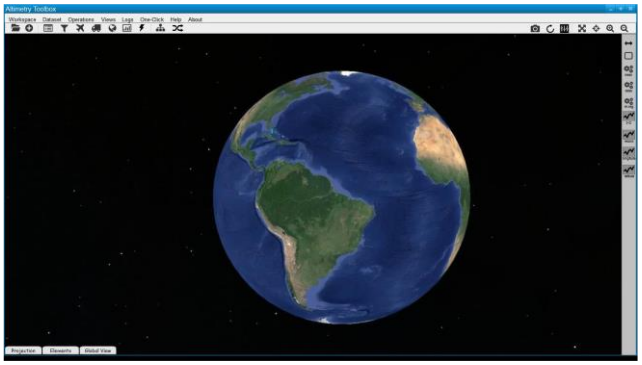

Figure 1. Main screen of the BRAT.

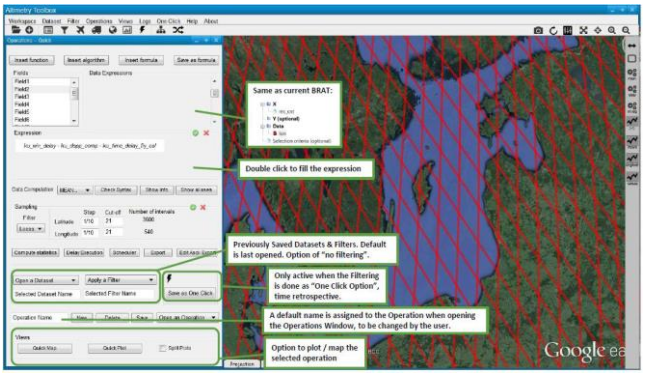

Figure 2. 'Dataset' screen of the BRAT.

doi: 10.2788/854791

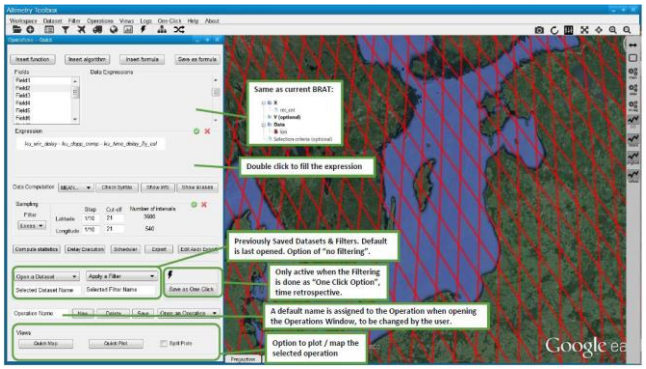

Figure 3. 'Operations' screen of the BRAT.

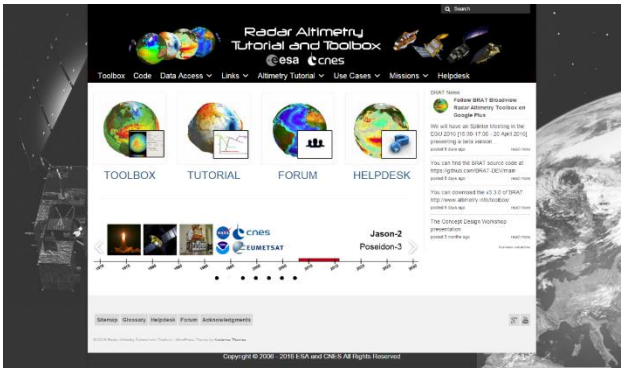

Figure 4. Home page of the BRAT website.

3. THE TUTORIALS

The **Radar Altimetry Tutorial**, now updated after the first release issued at the beginning of the project, contains an extensive introduction to radar altimetry describing its applications in different fields such as Oceanography, Cryosphere, Geodesy and Hydrology. Moreover, it has been updated with a new section covering **SAR altimetry**. This has been specifically conceived for the current project underlining the great potential of SAR altimetry, especially for coastal and inland applications.

Other written and video tutorials, showing how to use the toolbox and presenting some use-cases for both conventional and SAR altimetry can also be found on the website. All this content can be accessed through **http://www.altimetry.info/**.

4. THE USER COMMUNITY

One of the main goals of the BRAT consortium is to create a user community around the project. Apart from periodically organizing webinars, workshops and training courses, the project has promoted a Forum on the website (Figure 4 shows a snapshot of the main page) where users can discuss and share their knowledge, also suggesting modifications or additions to the code, which can be freely downloaded from the website.

Moreover, the BRAT consortium has created a helpdesk, which is meant to be an interactive channel of communication for both users of the toolbox and readers of the tutorials.

5. ACKNOWLEDGEMENTS

5.1 Background

The Broadview Radar Altimetry Toolbox, developed under an ESA contract within the SEOM (Scientific Exploitation of Operational Missions) program element, continues the work performed between 2006-2011 by CLS and S&T under ESA and CNES contracts, and focuses on the exploitation of the new missions Jason-3 and Sentinel-3.

5.2 Authors and Editors

The original authors of the Radar Altimetry Tutorial are:

V. Rosmorduc (CLS). O. Lauret (CLS),
C. Maheu (Akka), M.-P. Milagro (SERCO/ESRIN),
N. Picot (CNES) and J. Benveniste (ESA).

The editors of the Radar Altimetry Tutorial are J. Benveniste (ESA) and N. Picot (CNES).

doi: 10.2788/854791

INTEGRATION OF TOOLS FOR LARGE-SCALE EXPOSURE AND VULNERABILITY ASSESSMENT INTO THE REMOTE GEOSPATIAL PROCESSING ENVIRONMENT OF THE EUROPEAN SPACE AGENCY: AN EXPERIENCE ON "SENSUM TOOLS"

Daniele De Vecchi[1,2], *Fabio Dell'Acqua*[1,2], *Roberto Cuccu*[3,4], *Carla Orrù*[4], *Giancarlo Rivolta*[3,4]

[1] University of Pavia, Dep. of Industrial and Information Eng., via Ferrata 5, I-27100 Pavia, Italy
[2] EUCENTRE foundation, Aerospace section, via Ferrata 1, I-27100 Pavia, Italy
[3] ESA Research and Service Support, via Galileo Galilei, I-00044 Frascati RM, Italy
[4] Progressive Systems Srl, Parco Scientifico di Tor Vergata, I-00133 Roma, Italy

ABSTRACT

In the context of risk monitoring, some indirect measures called proxies, can feed vulnerability and exposure models over large areas. The extension of a built-up area, for example, and its change over time, are useful pieces of information. In the context of several FP7 projects, the research group at the University of Pavia developed a set of Python algorithms for local machines. In collaboration with the ESA Research and Service Support (RSS) team, two new services have been created based on the Earth Observation tools: a pre-configured customised virtual machine (RSS CloudToolbox) and a G-POD (Grid Processing On Demand) service. This manuscript includes a general description of the new parallel design and a comparison between the proposed services.

Index Terms— ESA G-POD, Research Support Service, Big Heritage Data, Risk Mapping, Integration, Built-up area.

1. EARTH OBSERVATION AND RISK ASSESSMENT

In the context of risk monitoring on large areas, complete assessment of fragility and/or economic value for each single building in an urban area would be infeasible. Yet some indirect measures, called proxies, can be defined [1] which help reckoning at least in a statistical way, the vulnerability and exposure over large areas without diving into unmanageable detailed assessments. Although not many, some of these proxies can also be extracted from satellite Earth Observation data.

EO-based exposure and vulnerability proxies can be defined at different spatial levels. Extension and outline of a built-up area, for example, is a valuable piece of information when assessing exposure to various natural risks on a vast area. Zooming into the urban areas, other proxies can be defined like buildings regularity and density, height, shape etc. Large-scale extraction of proxies is of interest to both public and private stakeholders seeking assistance in improving their assessment of the risk level distribution across a region, for various purposes. Thus, proxy extraction from satellite data has a potential market to generate, once some issues have been solved, one being the huge amount of satellite data to be processed and the heavy computational burden to be sustained.

Thanks to the EU for funding this work trough FP7 projects SENSUM, RASOR and MARSITE.

2. SOFTWARE TOOLS AND ESA SERVICES

Different open-source C and Python libraries are already available for image processing and geospatial information handling (e.g. OrfeoToolbox, OpenCV and GDAL) providing basic processing tools, producing results that reach a point located one step behind the exposure target. Therefore, it is of significant importance to provide end-users with a set of tools based on simple work flows capable of returning information at a higher level, in other words to "include the last mile".

Our research group at the University of Pavia has already developed a set of Python algorithms, a result of integrating low-level image processing and geospatial tools with high-level exposure-oriented workflows. A single processing chain leads to e.g. urban expansion monitoring based on multitemporal series; a sample result on Izmir, Turkey is shown in figure 1. These tools were developed and tested within the SENSUM project [2] framework and are released in two different forms: source code [3] and QGIS plugin [4]. The former solution is targeted to developers, offering a chance to instantaneously modify the workflows; the latter version is oriented to GIS users with less programming skills and used to interact with a graphic interface. Further development and maintenance is ongoing within the RASOR project framework [5] onto which geospatial platform the algorithms are being integrated, besides tested on various test sites including Istanbul, Turkey, in the framework of the MARSITE project [6].

With the lack of a unified software suite for vulnerability indicators extraction, the proposed solution can provide inputs for existing models like the Global Earthquake Model [7]. The inclusion of the proposed set of algorithms within

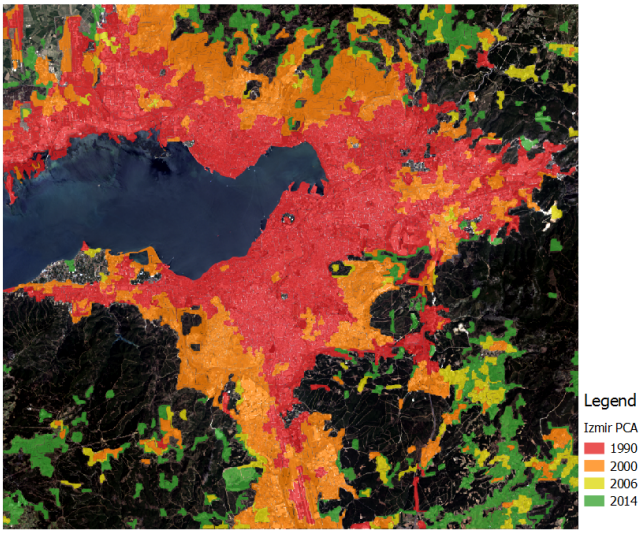

Fig. 1. Urban expansion map obtained through a SENSUM processing chain including image selection, registration of multitemporal stacks and hybrid urban area extraction.

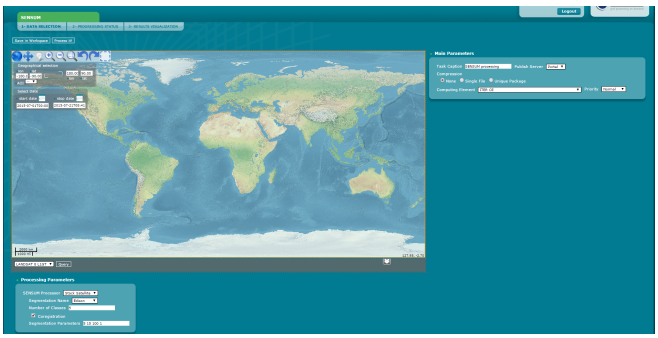

Fig. 2. SENSUM tools interface on the ESA G-POD system.

the RASOR platforms can guarantee support and enlarge the community of end-users.

All the above software was originally designed to run on a local machine, taking advantage of multi-core machines whenever possible. Due to the dependecies the code relies on, difficulties were registered by new users wishing to try the algorithm. A fruitful collaboration with the ESA RSS team [8] [9] led to software adaptation and testing by exploiting the RSS CloudToolbox service for the subsequent integration as G-POD service [10].

2.1. RSS CloudToolbox service

The RSS CloudToolbox service consists of an on-demand provisioning of customised virtual machines equipped with pre-installed software according to the user requirements. The virtual machine is created upon request with flexible hardware resources depending on the user needs. This resource is hosted on a cloud infrastructure and it is accessible via a secure shell connection. This solution was adopted in order to avoid endless installation and compatibility issues while developing and adapting the algorithm in preparation of the G-POD integration. The SENSUM Earth Observation Tools have been successfully tested into this service and therefore they're ready for the integration into the G-POD environment.

2.2. G-POD service

Following the objectives of the Big Data from Space challenge, a new parallel-driven structure has been designed to run on the ESA G-POD system [11] provided by the agency to support Earth Observation research, development and data exploitation [12]. The algorithm, known as Stack Satellite [13] [14] has been re-arranged and configured in order to retrieve data from the ESA repository and directly process it. The parallel architecture and the time saved in configuring the machine and downloading the data are among the main advantages. The processing workflow "backbone" is unchanged, meaning that the algorithms designed to extract built-up areas are formerly equivalent to the ones included in the QGIS plugin. Major adjustments, like re-writing the most demanding operations in Cython, have been carried out to further reduce the amount of time needed. The data flow has been redesigned, instead.

The new structure, shown in figure 3, takes advantage of multiple worker nodes depending on the amount of data to process. The *data handler* module (DH) converts the set of input imagery to a stack of bands over the area of interest. If requested, DEM mask and segmentation are computed using the most recent year, aiming at a stable set for all the other inputs. The *processing* module (PR) includes the different built-up extraction algorithms, including 3 pixel-based and 2 hybrid-based modules.

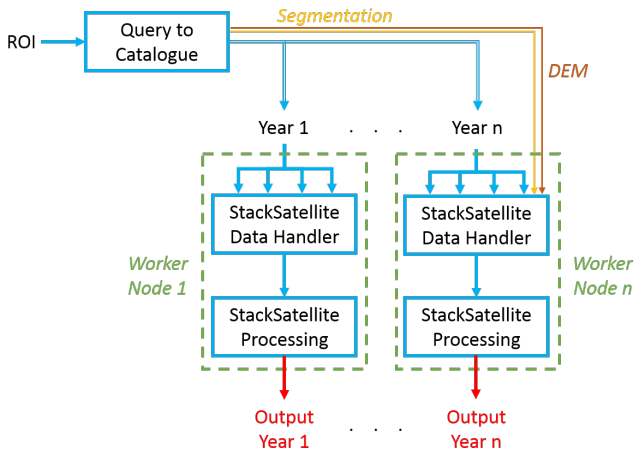

Fig. 3. New parallelized structure.

The user experience is quite straightforward:

doi: 10.2788/854791

1. Define the region of interest (a rectangle), the desired range of dates and the data set; a query to the catalogue is submitted in order to retrieve a list of matching items.
2. Choose from the list according to what is the objective of the extraction (e.g. the year gap) and define the methods to apply.

Once the inputs are defined, processing tasks are distributed within the system according to the available resources. The user is notified as soon as the process is completed with results available from the "Results visualization" tab.

3. RESULTS OF THE IMPLEMENTATION

CloudToolbox-related processing times are listed in table 1 and compared with the expected G-POD processing times running on 3 worker nodes concurrently. Tests have been carried out in an area covering Milan and part of Northern Italy, at the intersection of 3 different Landsat scenes. The final size is around 5435x3331 pixels. The CloudToolbox machine has been configured to match the G-POD worker nodes specifications, equipped with 8 CPUs and 32 GB of RAM.

4. CONCLUSIONS

This paper describes the efforts spent to integrate the SENSUM Earth Observation Tools as a G-POD service by taking advantage of the resources and flexibility provided by the RSS CloudToolbox. In the context of the Big Data from Space challenge, the Grid Processing On Demand can represent a key service with its parallel architecture while the CT service is a way to "taste" the algorithms capabilities. Further improvements are already planned for the integration, like the adaptation of the algorithms to Sentinel-2 imagery.

5. REFERENCES

[1] Harb, M.M.; De Vecchi, D.; Dell'Acqua, F., "Physical Vulnerability Proxies from Remotes Sensing: Reviewing, Implementing and Disseminating Selected Techniques" in IEEE Geoscience and Remote Sensing Magazine, vol.3, no.1, pp.20-33, March 2015 doi: 10.1109/MGRS.2015.2398672

[2] Web site of the EU FP7 project SENSUM [Online]. Available: http://www.sensum-project.eu/ Accessed 2015-10-14

[3] SENSUM source code [Online]. Available: https://github.com/SENSUM-project/sensum_rs Accessed 2015-10-14

[4] SENSUM QGIS plugin [Online]. Available: https://github.com/SENSUM-project/sensum_rs_qgis Accessed 2015-10-14

[5] Web site of the EU FP7 project SENSUM [Online]. Available: http://www.rasor-project.eu/ Accessed 2015-10-14

[6] Web site of the EU FP7 project MARSITE [Online]. Available: http://marsite.eu/ Accessed 2015-10-14

[7] The Global Earthquake Model (GEM) foundation web site [Online]. Available: http://www.globalquakemodel.org/ Accessed 2015-10-15

[8] ESA Research and Service Support (RSS) home page [Online]. Available: wiki.services.eoportal.org/tiki-custom_home.php Accessed 2015-10-15

[9] Marchetti, P.G.; Rivolta, G.; D'Elia, S.; Farres, J.; Mason, G.; Gobron, N., "A Model for the Scientific Exploitation of Earth Observation Missions: the ESA Research and Service Support", in IEEE Geoscience and Remote Sensing Society Newsletter, vol.162, pp.10-18, March 2012

[10] Cuccu, R.; Sabatino, G.; Delgado, J.M.; Rivolta, G., "Enabling SAR data exploitation by processing on-demand", in Geoscience and Remote Sensing Symposium (IGARSS), 2015 IEEE International, vol., no., pp.1476-1479, 26-31 July 2015 doi: 10.1109/IGARSS.2015.7326058

[11] ESA Grid Processing On Demand (G-POD) service page [Online]. Available: http://gpod.eo.esa.int/ Accessed 2015-10-14

[12] Farres, J.; Mathot, E.; Pinto, S., "G-POD: A collaborative environment for Earth Observation at the European Space Agency", in Proceedings of the ESA Living Planet Symposium, Bergen, Norway. Special publication SP-686 on CD-ROM, ESA Publications Division, European Space Agency, Noordwijk, The Netherlands, 2010

[13] Harb, M.; De Vecchi, D.; Dell'Acqua, F., "Automatic hybrid-based built-up area extraction from Landsat 5, 7, and 8 data sets", in Urban Remote Sensing Event (JURSE), 2015 Joint , vol., no., pp.1-4, March 30 2015-April 1 2015 doi: 10.1109/JURSE.2015.7120475"

[14] De Vecchi, D.; Harb, M.; Dell'Acqua, F., "A PCA-based hybrid approach for built-up area extraction from Landsat 5, 7 and 8 datasets", in Geoscience and Remote Sensing Symposium (IGARSS), 2015 IEEE International , vol., no., pp.1152-1154, 26-31 July 2015 doi: 10.1109/IGARSS.2015.7325975

Task Description	**ClTbx** (sec)	**G-POD** (sec)	**Input data**
Download data	~1800	600	LC81930292015196LGN00, LC81940282015267LGN00, LC81940292015267LGN00, LE71940282003114ASN00, LE71940292003114ASN00, LE71930292003123ASN00, LT51930292011233MOR00, LT51940282011272MOR00, LT51940292011272MOR00
DH (single year)	268	268	LE71940292003114ASN00, LE71930292003123ASN00, LE71940282003114ASN00
DH (single year, +DEM +segmentation)	726	726	LC81930292015196LGN00, LC81940282015267LGN00, LC81940292015267LGN00
DH (3 years)	804	278	LC81930292015196LGN00, LC81940282015267LGN00, LC81940292015267LGN00, LE71940282003114ASN00, LE71940292003114ASN00, LE71930292003123ASN00, LT51930292011233MOR00, LT51940282011272MOR00, LT51940292011272MOR00
DH (3 years, +DEM +seg)	2780	746	LC81930292015196LGN00, LC81940282015267LGN00, LC81940292015267LGN00, LE71940282003114ASN00, LE71940292003114ASN00, LE71930292003123ASN00, LT51930292011233MOR00, LT51940282011272MOR00, LT51940292011272MOR00
PR (single year, Pixel-based)	16	16	2015 stack
PR (single year, PCA-hybrid)	8276	8276	2015 stack
PR (single year, Dissimilarity-hybrid)	11169	11169	2015 stack
PR (3 years, Pixel-based)	48	21	2015 stack, 2011 stack, 2003 stack
PR (3 years, PCA-hybrid)	24828	8296	2015 stack, 2011 stack, 2003 stack
PR (3 years, Dissimilarity-hybrid)	33507	11190	2015 stack, 2011 stack, 2003 stack

Table 1. Real versus estimated processing time for the CloudToolbox and ESA G-POD services respectively.

doi: 10.2788/854791

HIGH-ACCURACY PHOTO-Z MEASUREMENTS FOR GALAXIES BASED ON SDSS-III PHOTOMETRY

*A.V.Meshcheryakov**

Space Research Institute,
Profsoyuznaya ul. 84/32,
Moscow, 117997 Russia

V.V.Glazkova, S.V.Gerasimov

Moscow State University,
Faculty CMC, GSP-1, Leninskie Gory,
Moscow, 119991, Russia

ABSTRACT

A new method for measuring the photometric redshifts (photo-z) of galaxies based on SDSS-III data is presented. We use the Random Forest empirical model trained on all of the potentially informative object features collected in the photometric catalog (apart from the standard flux measurements in five filters u, g, r, i, z). The method allows us to perform a highly accurate photo-z predictions based on SDSS photometric data, especially for low-redshift ($z < 0.35$) galaxies. We present the application of the method to the photo-z estimation for more than 70 million galaxies from the SDSS-III Final Data Release. The proposed method is finely scalable in application to photo-z's estimation in much bigger photometric datasets.

Index Terms— sky surveys, SDSS, photo-z, galaxies, machine learning, big data

1. INTRODUCTION

Two main approaches are currently used to measure the photometric redshifts of galaxies based on broadband photometry: photometric data fitting by well-known spectral templates (see, e.g., Bolzonella et al. 2000; Bruzual and Charlot 2003) and empirical machine learning methods with model training on a sample of objects with already measured spectroscopic z (see, e.g., Connolly et al. 1995; Carliles et al. 2010).

As was shown in a number of papers that have appeared in recent years (see e.g. D'Abrusco et al. 2007; Hoyle et al. 2014; Meshcheryakov et al. 2015), using various additional properties of objects contained in the photometric SDSS catalog (such as the extent and morphological type, the additional colors and magnitudes of objects, etc.) for each galaxy, apart from the standard u, g, r, i, and z flux measurements, allows the accuracy of photo-z estimates for galaxies by machine learning methods to be improved significantly (see e.g. Hoyle et al. 2014).

*The work was supported by the Russian Foundation for Basic Research (projects no.14-22-03111 ofi_m and 15-29-07085 ofi_m). We thank Microsoft for the cloud resources provided to our team within the Microsoft Azure for Research Program.

In this paper we show how to build photo-z empirical model, which uses in training the whole set of available object features stored in the photometric catalog of SDSS-III survey (which became a sort of benchmark for testing various photo-z methods). We present our method in the section 2. Then, the reference sample used to train the photo-z empirical model, is described in section 3. In the next section we estimate a photo-z uncertainty of our method, by using test galaxy sample with known z. In the section 4 we present SDSS catalog with measured photometric redshifts and shortly discuss the expected scalability of photo-z measurements, in the context of future massive photometric data sets.

2. METHOD

We investigated the efficiency of applying the currently available machine learning algorithms using the ideas of bagging (Breiman 1996), boosting (Kearns 1988), and the random subspace method (Bryll 2003), namely the Random Forest of Decision Trees (RFDT) (Breiman 2001), Extremely Random Forest of Decision Trees (ERFDT) (Geurts et al. 2006), and Gradient Boosting of Decision Trees (GBT) (see Friedman 1999, 2001) algorithms. The listed methods for most applied problems show results comparable to such highly accurate algorithms as the support vector machines (SVM) (Cortes and Vapnik 1995) and neural networks (McCulloch and Pitts 1943), while having a number of advantages: stability to overtraining and the absence of requirements on the distributions of input data (see Breiman 2001; Friedman 1999). The algorithms showed similar results in accuracy, but the best result on the prediction of photo-z for galaxies was obtained using the Random Forest (RFDT) algorithm.

The key concept of our photo-z method is the use of the entire space of available photometric attributes in training of Random Forest empirical model. We used all available photometric features of objects found in the PhotoObj table of SDSS photometric catalog, — only "technical" records in the catalog, which surely do not contain any observational information about sky objects, were excluded from the consideration. We used 581 photometric features (for each galaxy

doi: 10.2788/854791

extracted from the PhotoObj table of SDSS catalog) in the training of the empirical photo-z model.

3. DATA

As a reference sample (to train the RFDT model) we used the catalog of galaxies with measured redshifts from the SDSS-III spectroscopic survey (Data Release 12, Alam et al. 2015).

To download all the necessary information about reference galaxies the SDSS CasJobs database was used (see http://skyserver.sdss3.org/CasJobs/). We adopted the following conditions for the selection of reference objects from the database. Firstly, we selected only spectroscopically confirmed galaxies with reliably redshift measurements:

```
SELECT .. INTO ..
FROM PhotoPrimary as p, SpecObj as s
WHERE s.class = "GALAXY"
AND s.zWarning = 0
AND p.SpecObjID = s.SpecObjID
```

Then, we applied the standard filtering by the quality of the photometry (see http://www.sdss.org/dr12/algorithms/):

```
AND ( p.flags & 0x20) = 0
AND ( p.flags & 0x80000) = 0
AND ( p.flags & 0x400000000000) = 0
AND ( p.flags & 0x800000000000) = 0
AND ( p.flags & 0x10000000000) = 0
AND ( p.flags & 0x100000000000) = 0
AND ( p.flags & 0x10000000) != 0
AND ( p.flags & 0x40000) = 0
AND ( p.flags & 0x80) = 0
```

In order to produce more accurate photo-z measurements for galaxies, we decided to select from the catalog only objects detected in all g,r,i,z filters (with signal-to-noise ratio $S/N > 3$):

```
AND p.psfmagerr_g<0.3
AND p.psfmagerr_r<0.3
AND p.psfmagerr_i<0.3
AND p.psfmagerr_z<0.3
```

Finally, 1,060,304 reference galaxies were selected from the SDSS spectroscopic catalog. In this paper we focus on the application of our photo-z empirical model (trained on the reference sample) to the SDSS-III DR12 photometric catalog of galaxies.

4. PHOTO-Z ACCURACY

In order to estimate the accuracy of our method of photo-z measurements, we randomly split the reference catalog into two equal parts. 50% of objects from the reference catalog were used for training and cross-validation of the photo-z empirical model. The rest of galaxies were used as a test sample to estimate photo-z accuracy. To make direct comparison of our results with results from other photo-z methods in the literature, the standard statistical indicators (see, e.g., Brescia et al. 2014) were used: μ (the bias, mean value of residuals), σ (standard deviation), NMAD (normalized median absolute deviation), and RMS (root-mean-square of residuals), and the percentage of catastrophic outliers $|\Delta z| > 0.15$. The above quantities were calculated on normalized residuals Δz, defined in the following way: $\Delta z = (z_{ph} - z_{sp})/(1 + z_{sp})$, where z_{ph}, z_{sp} — photo-z estimate and spectroscopic redshift of galaxy, respectively.

Figure 1 shows the scatter diagram (measured photo-z of a galaxy versus its spectroscopic redshift) constructed for randomly chosen 30,000 galaxies from the test.

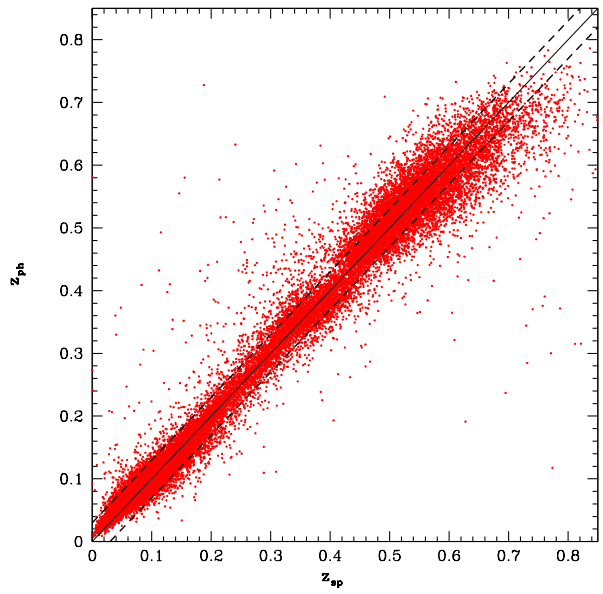

Fig. 1. Scatter diagram (measured photo-z versus spectroscopic redshift) constructed for randomly chosen 30,000 galaxies in the test sample. The solid line corresponds to $z_{ph} = z_{sp}$ case, and dashed lines indicate the relative $\pm 3\%$ deviation of photo-z measurements.

The achieved accuracy of photo-z measurements for galaxies on the SDSS-III test sample for different redshift in-

doi: 10.2788/854791

z range	$\|bias\|$ 10^{-2}	σ 10^{-2}	NMAD 10^{-2}	RMS 10^{-2}	$\|\Delta z\|$ > 0.15
Our results					
Overall	0.09	2.5	1.5	2.5	0.19%
$0, 0.115$	0.54	1.9	1.2	2.0	0.19%
$0.115, 0.177$	0.26	2.0	1.1	2.0	0.20%
$0.177, 0.345$	0.50	3.0	1.6	3.0	0.69%
> 0.345	0.24	2.8	2.1	2.9	0.02%
0.05..0.6	0.23	2.3	1.5	2.3	0.19%
Brescia'14					
Overall	0.06	2.3	1.7	2.3	0.12%
$0, 0.115$	0.79	2.0	1.7	2.2	0.03%
$0.115, 0.177$	0.78	2.0	1.7	2.2	0.02%
$0.177, 0.345$	0.56	2.8	1.9	2.9	0.36%
> 0.345	0.57	2.4	1.8	2.5	0.50%
0.05..0.6	0.16	2.2	1.6	2.2	0.11%

Table 1. Achieved accuracy of photo-z measurements for galaxies on the SDSS-III test sample. The redshift ranges for which the statistical indicators were calculated are given in the left column. Number of test objects within redshift ranges: 506479, 159192, 86289, 76859, 228343, 471592 (from top to bottom rows). In the bottom panel of the table we show results obtained by Brescia et al. (2014) in the same redshift ranges.

tervals, together with results obtained by Brescia et al. (2014) for the same z range, are shown in the Table 1. Note that Brescia et al. (2014) trained their MLPQNA empirical model on the 5 basic photometric features from SDSS catalog. As can be seen from the Table 1, the photo-z uncertainty in our method in the overall redshift range is very close to results of Brescia et al. (2014). On the other hand, our method gives substantially better photo-z results at low redshifts $z < 0.35$. This is not surprising, as the predictive power of additional photometric features from the SDSS catalog increases in the low-redshift regime.

5. THE PHOTO-Z CATALOG

We applied our photo-z empirical model, trained on the whole reference sample, to predict photometric redshifts of galaxies from the photometric catalog of SDSS-III survey. As we are interested in precise photo-z measurements, we decided to include in the target sample only galaxies detected ($S/N > 3$) in all main SDSS filters (the same condition was adopted to select the reference sample). A target sample was selected from the Casjobs SDSS database using the following SQL code:

```
SELECT p.* INTO ..
FROM Galaxy as p
WHERE p.psfmagerr_g<0.3
AND p.psfmagerr_r<0.3
AND p.psfmagerr_i<0.3
AND p.psfmagerr_z<0.3
```

The final sample, selected by the above query, contained 70,167,712 galaxies from the SDSS-III photometric catalog.

All photo-z measurements for the SDSS target sample were performed in the Microsoft Azure cloud computing platform. The virtual machine D14_v2 (16 cores, 112 GB RAM) with Ubuntu was used for calculations. The resulting photo-z photometric catalog is publicly available on the web page: http://astromining.org/projects/photoz/photoz_SDSS_BiDS16/.

It is worth noting, that Random Forest algorithm used in this work is finely scalable both for train (because of independence of regression trees construction) and for predict stage. We have tested these expectations, by measuring the duration of training and photo-z catalog creation for various number of computing cores. The duration of the train stage (using the whole reference sample — 1 million galaxies) almost linearly depends on the number of cores: 16 cores - 68 minutes, 12 cores - 88 minutes, 8 cores - 115 minutes, 4 cores - 189 minutes. For the predict stage we estimated the time needed for photo-z catalog (70 millions galaxies) creation: 16 cores - 12 minutes, 12 cores - 14 minutes, 8 cores - 18 minutes, 4 cores - 29 minutes). The parallel predictions were implemented by the "worker threads" model. Predictions on pieces of data were made by a thread pool. So we expect linear scalability on predict stage also for further increase of data volume.

We conclude, that the photo-z empirical method presented in this paper, can be easily applied to substantially ($\times$10-100) larger volumes of photometric catalogs, which will be produced by the next generation of ground-based and space-based photometric surveys (e.g. PanSTARRS, Euclid, LSST).

6. REFERENCES

S. Alam, F. Albareti, C. Prieto, F. Anders, S. F. Anderson, B. Andrews, et al., Astrophys. J. Suppl. Ser., 219, 12 (2015).
M. Bolzonella, J.-M. Miralles, and R. Pello, Astron. Astrophys. 363, 476 (2000).
L. Breiman, Machine Learning 24, 123 (1996).
L. Breiman, Machine Learning 45, 5 (2001).
M. Brescia, S. Cavuoti, G. Longo, and V. de Stefano, Astron. Astrophys. 568, 1126 (2014).
G. Bruzual and S. Charlot,Mon. Not. R. Astron. Soc. 344, 1000 (2003).
S. Carliles, T. Budavari, S. Heinis, C. Priebe, and A. Szalay, Astrophys. J. 712, 511 (2010).
A. Connolly, I. Csabai, A. Szalay, D. Koo, R. Kron, and J. Munn,Astron. J. 110, 2655 (1995). 21. C. Cortes and V. Vapnik, Machine Learning 20, 273 (1995).
R. D'Abrusco, A. Staiano, G. Longo, M. Brescia, M. Paolillo, and E. de Filippis, Astrophys. J. 663, 752 (2007).
J. Friedman, Comput. Stat. Data Anal. 38, 367 (1999).

J. Friedman, Ann. Stat. 29, 1189 (2001).
P. Geurts, D. Ernst, and L. Wehenkel, Machine Learning 63, 3 (2006).
B. Hoyle, M. Rau, R. Zitlau, S. Seitz, and J. Weller, arXiv:1410.4696, Mon. Not. R. Astron. Soc. 439, 48 (2014).
M. Kearns, Unpublished manuscript, Machine Learning class project (1988).
W. McCulloch and W. Pitts, Bull. Math. Biophys. 5, 115 (1943).
A. V. Meshcheryakov, V. V. Glazkova, S. V. Gerasimov, R. A. Burenin, G. A. Khorunzhev, Astronomy Letters 41, 307 (2015)

doi: 10.2788/854791

A FLEXIBLE AND DYNAMIC APPROACH TO COPE WITH THE NEW CHALLENGE OF BIG DATA IN EARTH OBSERVATION

Błażej Fitrzyk[1,2], Giovanni Sabatino[1,2], Roberto Cuccu[1,2], Júlio Carreira[1,3], Giancarlo Rivolta[1,2]

[1] ESA Research and Service Support, via Galileo Galilei, 1, 00044 Frascati (Italy)
[2] Progressive Systems Srl, Parco Scientifico di Tor Vergata, 00133 Roma (Italy)
[3] Critical Software, Taveiro , Coimbra (Portugal)

ABSTRACT

In Earth Observation (EO), the Big Data era started with the launch of Sentinel-1A in April 2014. Sentinel missions deliver data volumes a factor of 10 larger than previous EO missions. New services supporting EO research shall remove the obstacles deriving from manipulation and processing of such large data volumes. At the same time, in order to keep the cost low, an efficient service shall provide dynamic and scalable solutions capable to work automatically with limited actions by the operators.
In this paper, a flexible and dynamic solution for EO data processing on-demand is presented.

***Index Terms*— Big data exploitation, earth observation, research support service, dynamic processing resources, cloud, grid computing, flexible, scalable**

1. INTRODUCTION

The new era of Big Data poses new challenges in many fields of human activity, from social life to finance, from banking to science and ICT. Not only data storage and data handling are becoming more and more demanding, but also data processing can be difficult to carry out when data volumes increase this rapidly. The new requirements are giving rise to new issues and impose to review the current computing models which were thought only few years ago but for lower data volumes.
The Grid Processing on-Demand environment (G-POD) [1] made available and operated by ESA Research and Service Support (RSS) [2] service, is a platform based on the Grid technology for running processing jobs on demand. Users' algorithms are integrated in the system and have access to a large amount of Earth Observation satellite data (~400 TB locally) for processing purposes. The modularity of G-POD components makes it possible to extend its computing resources on the Cloud and to add new data archives to download satellite products on demand, thus making the capacity of the system virtually unlimited. Nevertheless in the era of Big Data, new computing paradigms have to be thought to cope with the obstacles arising from the management and processing of large data volumes: in fact the simple addition of resources federated on a Cloud provider (which is today operated by hand from the operator) to accommodate an increasing processing request, although can solve the problem temporarily, is unnecessary and significantly expensive as a static solution.
A smarter solution would add new processing resources when the processing request increases, and remove the unnecessary extra-resources when the processing demand decreases. Moreover such a tool should be able to take decisions about the amount of resources to add/remove also considering the available budget. This tool, named Automatic Cloud Scaling Engine (ACSE), has been designed and prototyped by RSS and is planned to be tested on several Cloud providers. In the next sections we will give an overview and a technical description of ACSE and will discuss the test performed so far.

2. THE CLOUD SCALING ENGINE

2.1 General Information

The Automatic Cloud Scaling Engine has been developed by RSS to enable the optimal allocation of computing resources on Cloud providers in order to fully utilize all the available resources when it is necessary while keeping low the operational costs.
The resources optimized by the Cloud Scaling Engine are Worker Nodes of TORQUE-based [3] computing clusters, set up as Linux virtual machines (only TORQUE-based clusters are supported and have been tested for now. The ACSE allows anyway to include clusters with other resource management tools). Each virtual machine utilizes a certain amount of resources (CPU, RAM, disk space, etc.) for which the customer has to pay. It is therefore desirable to keep the number of running Worker Nodes low and only

doi: 10.2788/854791

increase their number when there is a need for higher computing power. The ACSE handles this task automatically by dynamically creating new Worker Nodes when it is necessary, from an existing virtual machine template, up to a number obtained from a scaling algorithm and removing existing Worker Nodes when they stay idle for a predefined period of time.

2.2 Architecture

The Automatic Cloud Scaling Engine comprises two essential components:

1. The Cloudscaler daemon for monitoring clusters and dynamically managing their resources. The Cloudscaler daemon can be run on any host on the network as long as there is access from that host to the cluster's head (the Computing Element or CE) and to the cloud platform's API server, since it has to be able to communicate with them. A single instance of the Cloudscaler daemon can manage any number of computing clusters simultaneously with each of them having a different configuration. The Cloudscaler daemon is written in Python and has a modular architecture which allows to manage clusters running on different Cloud platforms. It can also use different scaling algorithms for calculating the number of Worker Nodes to be created or removed. The Cloudscaler provides logs of the operations and can also be run in the one-off mode to check and update the cluster resources just once. The Cloudscaler only runs on Linux and has been tested with the following distributions: Ubuntu 14.04, CentOS 6.6 and 6.7.

2. Customized installations of one or more TORQUE-based computing clusters. Customization of the clusters is the result of an optimization work aimed at making available software and library packages needed to the users for their research/analysis activity. The customizations allow dynamic configuration of the cluster, keep track of the Worker Nodes used as well as provide information about the cluster load and Worker Nodes to the Cloudscaler daemon. These utilities are written as Linux shell scripts in bash, with some of them acting as init scripts launched on virtual machines startup and shutdown. The scripts have been tested on a TORQUE cluster set up on CentOS 6.6 and 6.7.

2.3 Workflow

The Automatic Cloud Scaling Engine workflow, represented schematically in Fig. 1, is described in this section.

1. When a new cluster is created, the CE (cluster's head) must be created first from the CE Image on the Cloud platform. Also the WN Image serves as a template for creating other Worker Nodes on demand.
2. The Cloudscaler daemon checks periodically for load of the cluster by running a script over *ssh* on the CE and reading its output.
3. Based on the information obtained in #2 and the scaling algorithm the Cloudscaler daemon calculates the number of WNs to be created/removed and sends a command to the Cloud platform to create or remove WNs.
4. If new WNs are to be created, the Cloud platform instantiates new virtual machines from the WN Image. If WNs are to be removed, the Cloud Platform shuts down the virtual machines.
5. Newly created WNs report their availability to the CE. WNs being removed report their unavailability to the CE. The cluster configuration on the CE gets updated accordingly.
6. Additionally, CE checks periodically for availability of WNs by trying to *ssh* to them. If the connection fails, the WN is considered unavailable and is removed from the cluster configuration.

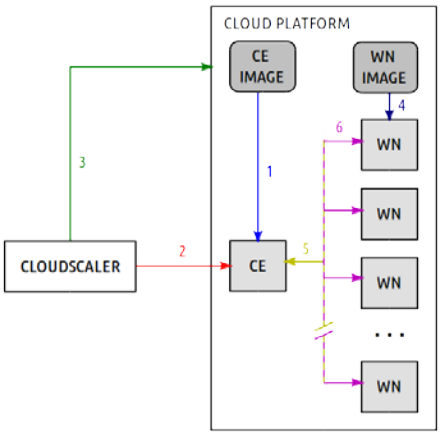

Figure 1: ACSE workflow

2.4 Cluster Configuration

Each computing cluster the Cloudscaler daemon manages has its own configuration within the Cloudscaler. The configuration defines:

– The Cloud platform driver used for the cluster.
– Platform driver specific settings. These depend on the Cloud platform on which the computing cluster is set up. Usually the identifier or name of the virtual machine template used for creating new Worker Nodes is also configured here.

doi: 10.2788/854791

- The scaling algorithm used by the Cloudscaler to manage the cluster resources.
- Parameters of the cluster and its Worker Nodes such as the minimum and maximum number of Worker Nodes, maximum idle time of Worker Nodes and the frequency of obtaining the information about the cluster load.
- Parameters of the cluster's head for accessing it over *ssh*, such as its IP address, login and password or *ssh* key.

2.5 Cloud Platform Drivers

Clusters may run on different Cloud platforms. To manage a cluster on a certain Cloud platform through the Cloudscaler daemon a driver for that platform must be provided.
The Cloudscaler provides an abstract base scaler class. This class must be extended in the platform-specific driver by implementing methods for creating and removing Worker Nodes on the cloud platform. Certain naming conventions for classes and files must be followed when implementing the driver. Drivers typically utilize platforms' Python API bindings.
Currently there are drivers available for the following Cloud platforms:

- ITER Teide-HPC platform managed by OpenNebula [4] http://opennebula.org/
  useing the *python-oca* (Python Bindings for XML-RPC OpenNebula Cloud API) from https://github.com/python-oca/python-oca
- another commercial Cloud provider.

2.6 Scaling Algorithms

The scaling algorithm is the core part of the Cloudscaler which decides when and what number of Worker Nodes can be created or removed. The modular architecture of the Cloudscaler allows to use different scaling algorithms for different clusters.
The Cloudscaler provides an abstract base algorithm class. This class must be extended in the specific algorithm by implementing a method for calculating the number of Worker Nodes to be created or removed. Certain naming conventions for classes and files must be observed when implementing the algorithm.

2.6.1 Budget Scaling Algorithm

Currently the only implemented algorithm is the Budget algorithm. The algorithm adds some abstraction to the way clusters are normally defined. It is not based on the number of Worker Nodes but rather on the total budget time of Worker Nodes and the time already consumed by the Worker Nodes measured in defined time units.

The Cloudscaler daemon contacts the cluster's head in intervals defined in the cluster configuration to obtain the data on cluster load. The data is then passed to the algorithm which performs the following operations:

1. If there are Worker Nodes idle for longer time than the time defined in the cluster configuration, it orders the Cloud platform to shut them down and remove the virtual machines. A number of idle Worker Nodes may be left reserved for future use if it is set in the cluster configuration.
2. Otherwise, if there are new jobs in the cluster's queue it checks if and how many Worker Nodes should be created (see Fig. 2):
   - divides the entire budget period into time intervals
   - checks how many Worker Nodes were used in each interval and sums the amount
   - compares the calculated usage with the theoretical maximum usage, which is the number of time units passed times the number of Worker Nodes in the budget period
   - decides on the number of new Worker Nodes to be created. The number depends on the past usage, the budget number of Worker Nodes and the absolute maximum number of Worker Nodes from the cluster configuration.

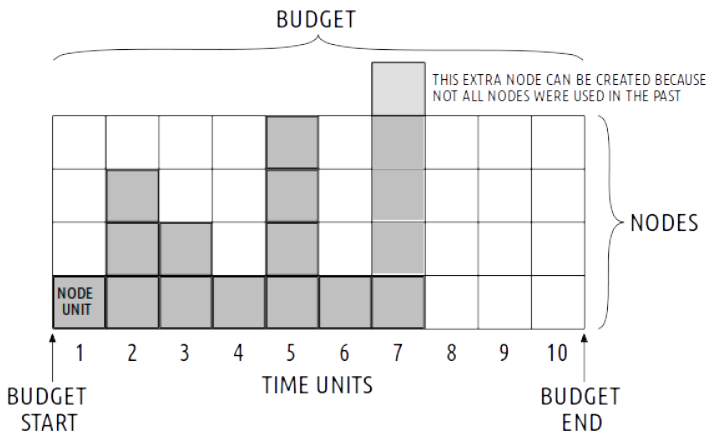

Figure 2. Schematic view of the Worker Nodes usage in the time units as defined in the scaling algorithm configuration.

3. TEST AND RESULTS

Tests for both supported platforms were performed: the Cloud platform provided by ITER Teide-HPC data Centre and another commercial Cloud provider.

doi: 10.2788/854791

Two virtual machines were created initially using the minimal installation of CentOS 6.7: one for the Computing Cluster and one for the Worker Node. The required cluster software was installed and configured on both machines, following the cluster deployment guidelines. Networking was set up on the machines according to the capabilities provided by the platform. Finally, the Cloud Scaling Engine scripts were installed and configured on both machines. The Computing Element was then fully functional and ready to use, while the Worker Node virtual machine served as a template for creating other Worker Nodes on demand by the Cloudscaler.

The Cloudscaler daemon was installed on another virtual machine. Two cluster configurations were created for the Cloudscaler daemon – one for each platform to test.

Several tests were performed on each platform, with the cluster budget period being set each time to different values, the "granularity" or the time unit of the budget period from 1 to 10 minutes and the number of Worker Nodes in the budget period from 2 to 5 (this also depended on the resources available on the platform).

The tests comprised adding a number of computing tasks exceeding the number of existing Worker Nodes (up to 20 simultaneous tasks) to the clusters' queues. The tasks were spawned from a script. A single task was mainly a simple loop which kept the Worker Node occupied for a few to several dozen minutes.

It was observed if the Cloudscaler daemon created new Worker Nodes to meet the increased demand and in accordance with the parameters of the algorithm, i.e. the limits set up in the configuration were not exceeded. Then, it was observed if the Cloudscaler properly shut down and removed the Worker Nodes after the computing tasks were finished. These tests were performed with good results demonstrating that the Worker Nodes were created and removed in full accordance with the cluster configuration without affecting the cluster functionalities.

Initial tests were also conducted with the clusters integrated with the G-POD (Grid Processing On-Demand) environment operated by the ESA Research and Service Support for performing heavy scientific computations. The G-POD automatically assigns tasks to clusters and it was examined if it worked well for ACSE-managed clusters. Preliminary results are promising, yet so far not fully conclusive and further tests will have to be performed.

4. CONCLUSIONS

The ESA Research and Service support has designed and developed an automatic scaling prototype, ACSE, aimed to add/remove processing resources to/from a computing cluster ideally hosted on any IaaS/Cloud provider. The tool takes into account the processing requests and the available budget to expand or reduce the cluster capacity. Tests have been performed on the Cloud platform provided by ITER Teide-HPC data Centre and on another commercial Cloud provider. Test results are promising, although it has still to be tested how ACSE behaves in a real computing environment where the number of Worker Nodes usually reaches and exceeds the range of tens.

In the new era of Big Data ACSE goes in the direction of responding to the increasing processing demand by keeping operational cost as limited as possible.

5. REFERENCES

[1] J. Farres, E. Mathot, S. Pinto, G-POD: A Collaborative Environment for Earth Observation at the European Space Agency, Proceedings of the ESA Living Planet Symposium, Bergen, Norway. Special Publication SP-686 on CD-ROM, ESA Publications Division, European Space Agency, Noordwijk, The Netherlands, 2010.

[2] Marchetti P.G., Rivolta G., D'Elia S., Farres J., Mason G. and Gobron N. (2012). A Model for the Scientific Exploitation of Earth Observation Missions: The ESA Research and Service Support. IEEE Geoscience and Remote Sensing(162): 10-18.

[3] G. Staples, Torque resource manager, in Proceedings of the 2006 ACM/IEEE conference on Supercomputing , ser. SC '06. ACM, 2006.

[4] IaaS Cloud Architecture: From Virtualized Datacenters to Federated Cloud Infrastructures, R. Moreno-Vozmediano, R. S. Montero, I. M. Llorente. IEEE Computer, vol. 45, pp. 65-72, Dec. 2012.

SAR ALTIMETRY PROCESSING ON DEMAND SERVICE FOR CRYOSAT-2 AND SENTINEL-3 AT ESA G-POD

Salvatore Dinardo[(1)], Marco Restano[(1)], Américo Ambrózio[(2)], Jérôme Benveniste [(3)]

[(1)] SERCO/ESRIN, Via Galileo Galilei, Frascati, Italy
[(2)] DEIMOS/ESRIN, Via Galileo Galilei, Frascati, Italy
[(3)] ESA-ESRIN, Via Galileo Galilei, Frascati, Italy, Email: jerome.benveniste@esa.int

ABSTRACT

The scope of this paper is to feature the G-POD SARvatore service to users for the exploitation of CryoSat-2 data, which was designed and developed by the Altimetry Team at ESA-ESRIN EOP-SER. The G-POD service coined SARvatore (SAR Versatile Altimetric Toolkit for Ocean Research & Exploitation) is a web platform that allows any scientist to process on-line, on-demand and with user-selectable configuration CryoSat-2 SAR/SARIN data, from L1a (FBR) data products up to SAR/SARin Level-2 geophysical data products. Several years of CryoSat-2 FBR data are at the disposal of the user, plus the full power of the G-POD's cluster: 350 CPUs and over 500 TB of storage.

Index Terms - SAR ALTIMETRY, CRYOSAT, GPOD, SAMOSA, SENTINEL-3 STM

1. INTRODUCTION

The SAR Versatile Altimetric Toolkit for Ocean Research & Exploitation (SARvatore) takes advantage of the G-POD (Grid Processing On Demand) distributed computing platform (350 CPUs in ~70 Working Nodes) to timely deliver output data products and to interface with ESA-ESRIN FBR data archive (155'000 SAR passes and 41'000 SARin passes). The output data products are generated in standard NetCDF format (using CF Convention), therefore being compatible with the Multi-Mission Radar Altimetry Toolbox (BRAT) and other NetCDF tools. By using the G-POD graphical interface, it is straightforward to select a geographical area of interest within the time-frame related to the Cryosat-2 SAR/SARin FBR data products availability in the service catalogue. The processor prototype is versatile, allowing users to customize and to adapt the processing according to their specific requirements by setting a list of configurable options. After the task submission, users can follow, in real time, the status of the processing, which can be lengthy due to the required intense number-crunching inherent to SAR processing. From the web interface, users can choose to generate experimental SAR data products as stack data and RIP (Range Integrated Power) waveforms. The processing service, initially developed to support the awarded development contracts by confronting the deliverables to ESA's prototype, is now made available to the worldwide SAR Altimetry Community for research & development experiments, for on-site demonstrations/training in training courses and workshops, for cross-comparison to third party products (e.g. CLS/CNES CPP or ESA SAR COP data products), for the preparation of the Sentinel-3 Surface Topography Mission, for producing data and graphics for publications, etc. Initially, the processing was designed and uniquely optimized for open ocean studies. It was based on the SAMOSA model developed for the Sentinel-3 Ground Segment using CryoSat data (Cotton et al., 2008; Ray et al., 2014). However, since June 2015, a new retracker (SAMOSA+) is offered within the service as a dedicated retracker for coastal zone, inland water and sea-ice/ice-sheet. In view of the Sentinel-3 launch, a new flavor of the service will be initiated, exclusively dedicated to the processing of Sentinel-3 mission data products. The scope of this new service will be to maximize the exploitation of the upcoming Sentinel-3 Surface Topography Mission's data over all surfaces. The service is open, free of charge (supported by the ESA SEOM Programme Element) for worldwide scientific applications and available at https://gpod.eo.esa.int/services/CRYOSAT_SAR/. In this paper, we present first the ESA G-POD framework and system. Then we describe in detail the CryoSat-2 SAR Processing services integrated in G-POD and we conclude with the output package description and information on the contacts and references.

2. G-POD SYSTEM

The ESA Grid Processing on Demand (G-POD) system is a generic GRID-based operational computing environment where specific data-handling Earth-Observation services can be seamlessly plugged into system. One of the goals of G-POD is to provide users with a fast computational facility without the need to handle bulky data.

The G-POD system hosts high-speed connectivity, distributed processing resources and large volumes of data to provide scientific and industrial partners with a shared data processing platform fostering the development, validation and operations of new Earth Observation applications.
In particular, the G-POD environment consists of:

- Over 350 CPUs in about 70 Working Nodes
- Over 330 TB of local on-line Storage plus 180 TB of EO data accessed directly from the PACs.
- Access to Cloud processing and data resources on demand (from Interoute and other providers)
- Internal dedicated 1 Gbit LAN at ESA-ESRIN and at UK-PAC archives
- 1 Gbps external connection
- Online software resources: IDL, MATLAB, BEAT, BEAM, BRAT.

Actually, G-POD has more than 300TB of EO data locally stored. EO Data available to G-POD services come either from ESA and

doi: 10.2788/854791

non-ESA missions. The G-POD web portal (http://gpod.eo.esa.int/) is a flexible, secure, generic and distributed web platform where the user can easily manage all own tasks. From the creation of a new task to the output publication, including data selection and job monitoring, the user goes through a friendly and intuitive user interface accessible from everywhere. More detailed information on the G-POD Web Portal and System are available here: http://wiki.services.eoportal.org/tiki-index.php?page=GPOD+User+Manual#Annex

3. CRYOSAT-2 SAR PROCESSING ON DEMAND SERVICE

The ESA G-POD Earth-Observation Service, SARvatore (SAR Versatile Altimetric Toolkit for Ocean Research & Exploitation) for CryoSat-2 is an Earth-Observation application that provides the capability to process remotely and on demand CryoSat-2 SAR data, from L1a (FBR, Full Bit Rate) data products until SAR Level-2 geophysical data products (Jensen and Raney, 1998; Wingham et al., 2006; Martin-Puig et al., 2008; Raney, 2008; Raney, 2012; Raney 2013).

The service works over any kind of surfaces but it has been so far optimized for ocean studies. It has been recently enhanced for inland water, land, sea-ice and ice sheets, implementing the SAMOSA+ model. The service is based on the SAR Processor Prototype that has been developed entirely by the ESA-ESRIN EOP-SER Altimetry Team (the authors) for CryoSat-2 validation purposes and preparation for the Sentinel-3 mission, with the following system features:

- SAR/SARin FBR/L1b DATA Archiving and Cataloguing
- SAR/SARin L1b Processor Prototype (Standard Delay-Doppler Processing)
- SAR/SARin L2 Retracker Prototype with SAMOSA Analytical Model and LEVMAR Least Square Estimator (Cotton et al., 2008; Ray et al., 2014)
- Input: CRYOSAT SAR/SARIN FBR DATA
- Output L1b → Radar Echogram
- Output L2 → SSH, SLA (w/o SSB), SWH, sigma0, wind speed

The ESRIN EOP-SER ALT Team succeeded to compile the processor for a 64-bit Linux platform and delivered to the ESA G-POD team the executable codes, the input archive (CryoSat SAR FBR) and satellite footprints (ASCII tracks).
Now, the toolkit has been fully integrated in the GPOD System for gridded and on-demand computation.

The objectives of the service integration in GPOD are:

- to experiment in-house research themes that will be further matured in the ESA-funded R&D projects;
- to provide expert users with consolidated SAR geo-products to get acquainted with the novelties and specifities of SAR Altimetry;
- to validate CryoSat-2 ocean products and get prepared for the exploitation of the Sentinel-3 mission.

The service is open, free of charge and accessible online from everywhere. In order to be granted the access to the service, you need an EO-SSO (Earth Observation Single Sign-On) credentials (for EO-SSO registration, go to https://earth.esa.int/web/guest/general-registration) and afterwards, you need to submit an e-mail to G-POD team (write to eo-gpod@esa.int), requesting the activation of the service for your EO-SSO user account.
After the registration to EO-SSO, users can freely access the online service at: https://gpod.eo.esa.int/services/CRYOSAT_SAR/. The service is listed under the Marine Theme. You can find it using the search bar as well.

The current GPOD service works only in SAR and SARin Mode (no LRM mode). As of October 2015, in the service catalogue, we have stored 150 thousands of SAR passes over the entire globe for period 2010-2015. This amounts to 20 TB of CryoSat-2 FBR data archived into G-POD storage. They can be all processed on-demand and online at user request.

4. WEB USER INTERFACE

Once you get to the service page (Fig. 1), the first action is to select the zone of interest and the time of interest for the required run. Regarding the selection of the area of interest, the user can simply draw a rectangle on the world map, after clicking on the rectangle icon on the tool bar. Instead, for more precise geo-selection, the user can type directly the geo-coordinates of the area of interest using the geographical bar.

Regarding the time of interest, the user may set the start and stop dates in the calendar bar. By default, the start date is the time of CryoSat-2 launch and the stop date is set at 2 months prior to the current date. The GUI embeds all the standard buttons for image browsing as panning, zoom-in zoom-out, centering, undo, redo, reset, etc.

Once the time and geo selection is done, clicking on the "QUERY" button, the service lists all the CryoSat-2 passes matching the time and space requirements. The CryoSat-2 SAR tracks, crossing the area of interest, are then shown on the world map in overlay. The graphical interface lists a maximum of 100 passes per page and informs users of the total number of found passes. The user can decide which passes to select by clicking on the passes, select all, or delete some specific passes from the list.

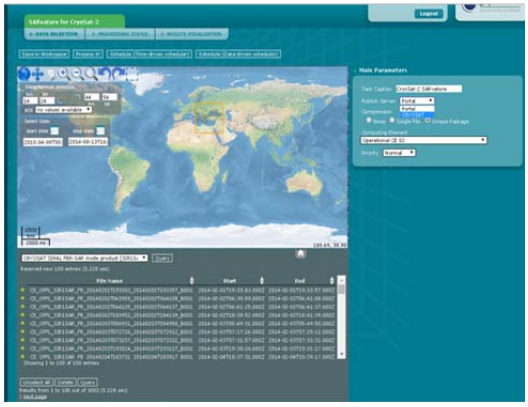

Figure 1: G-POD CryoSat-2 Service Main Interface.

On the top right, user finds a preference panel wherein user can set:

- Name of the current task
- Ftp Server where to publish the results (portal or personal)
- Data compression (tgz, none, single file)
- Grid Computing Resources
- Task Priority

The last step, before submitting the task, is to set the list of processing options. Indeed, the processor prototype is versatile in the sense that the users can customize and adapt the processing algorithms with flags and parameters, according their specific requirements, acting upon a list of configurable options. In the G-POD interface, users can easily enter this list of processing options via a series of drop-down menus. The configurable options are divided according to the processing level they refer to (L1b and L2). Starting from the first SARvatore release in 2014, the following upgrades have been introduced:

- Support for CryoSat-2 SARin Data.
- Enhancement of re-tracking capabilities in coastal zone and inland water by means of an advanced SAMOSA algorithm (SAMOSA+).
- Added support for posting rate at 80 Hz in delivering the output geophysical parameters.
- New Tide Model (TPXO8) and Geoid (EGM2008).

Moreover, by selecting the processing options properly, users can mimic the CryoSat-2 or the Sentinel-3 processing baseline for an easy cross-comparison between missions. Once the user has selected his processing options, in order to submit the task to G-POD Computing Elements, remains to click on the "PROCESS IT" button. After submission of a job, users will be directed to the workspace page where they can monitor in real time the status of the run and can be notified on the run status. The color code is:

- **Orange** → run under processing
- **Green** → run completed
- **Red** → run failed

Furthermore, by clicking on the task, the user can have more information on the processing task, such as:

- Task Id
- Task Creation Time
- Processing Id
- Grid Working Node Id
- Task Progress (retrieving, processing, publishing)

After run completion, by clicking on the button "Jobs Information", the user can inspect:

- the GPOD log file (.stdout or .stderr) where eventual errors on data retrieving or data storing are reported;
- the prototype configuration file (L1b_CONFIG_FILE.log and L2_CONFIG_FILE.log) where are reported all the processing options;
- the prototype log files (L1b_start.log and L2_start.log) where are reported eventual prototype processing errors.

Users can also decide to change one or more processing options and then re-submit the task. In case of successful run completion (green status), the portal will provide an http link from where to download the output package on the user's own local drive. The users can order to post the package directly on a personal ftp server after having communicated to the web platform the ftp server credentials (through the "publish servers" sub-menu). This is the recommended option in case of processing of large amount of data.

Future releases will:

- Support the UPorto GPD wet correction.
- Support the new Geoid (EIGEN-4C6) and the Tide Model (FES 2012).
- Provide a sea state bias solution.

5. OUTPUT PACKAGE & BRAT TOOLBOX COMPATIBILITY

The output package consists of:

- Satellite Pass Ground-Track in KML format
- Radar Echogram Picture in PNG format
- L2 Data Product in NetCDF format containing all the scientific results

The NetCDF format is self-explanatory with all the data field significance described in the attributes. The NetCDF Data Product follows the CF (Climate&Forecast) 1.6 Convention and can be opened with any standard NetCDF tools (ncdump, HDFview, etc.).

The following upgrades have been introduced for NetCDF Data Products:

- Inclusion of SAR echo and SAR RIP (Range Integrated Power) waveforms in the NetCDF files.
- Inclusion of STACK Data in the NetCDF files.

The recommended option is to ingest the NetCDF Data Products in BRAT Toolbox in order to exploit all the BRAT functionalities to browse and visualize the output content (Fig. 2). The Broadview Radar Altimetry Toolbox (BRAT) is a software suite designed to facilitate the use of radar altimetry data. It is able to read most distributed radar altimetry data, from ERS-1, ERS-2, TOPEX/Poseidon, Geosat Follow-On, Jason-1, Jason-2, Envisat, CryoSat-2, Jason-3 and Sentinel-3, to perform some processing, data editing and statistics, and to visualize the results. As part of the Toolbox, a Radar Altimetry Tutorial provides information about radar altimetry, the technique involved and its applications, as well as an overview of past, present and future missions, including information on how to access data and additional software and documentation. It also presents a series of data use cases, covering all uses of altimetry over ocean, cryosphere, inland water and land, showing the basic methods for some of the most frequent manners of using altimetry data. BRAT has been developed under contract with ESA and CNES (http://www.altimetry.info and http://earth.esa.int/brat/).

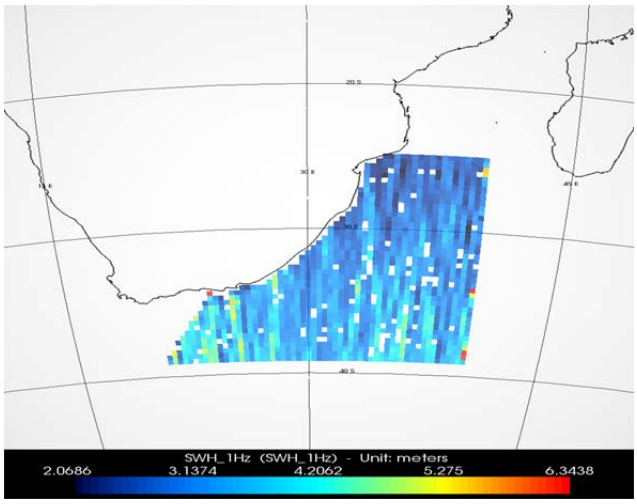

Figure 2: G-POD SAR CryoSat-2 Data products (Wave Height) opened in BRAT

6. CONCLUSIONS

To foster a new generation of SAR altimeter specialists and to get prepared for the Scientific Exploitation of Operational Missions (SEOM), a configurable versatile SAR processor has been developed and hosted in the ESA G-POD infrastructure. The G-POD Service coined SARvatore (SAR Versatile Altimetric Toolkit for Ocean Research & Exploitation) is a web platform that provides the capability to process on-line and on-demand CryoSat-2 SAR data, from L1a (FBR) data products until SAR Level-2 geophysical data products, with a suite of selectable configuration parameters. The processing algorithms are the ones used in the Sentinel-3 Ground Segment, which mathematical model, SAMOSA, is described in Ray et al. (2014). By selecting the processing options properly, users can mimic the CryoSat-2 or the Sentinel-3 processing baseline for an easy cross-comparison between missions. The Broadview Radar Altimeter Toolbox can display the output of SARvatore. The service is open, free of charge and accessible online from everywhere.

7. FURTHER INFORMATION

For any question, bug report and support, please contact us at: altimetry.info@esa.int

For G-POD platform specific questions please contact: eo-gpod@esa.int

Service Manual is available at: http://wiki.services.eoportal.org/tiki-index.php?page=GPOD+CryoSat-2+SARvatore+Software+Prototype+User+Manual

SARvatore is available at: https://gpod.eo.esa.int/services/CRYOSAT_SAR/

BRAT is available at: http://earth.esa.int/brat

8. REFERENCES

Cotton, D. et al., 2008, Development of SAR Altimetry Mode Studies over Ocean, Coastal Zones and Inland Water - State of the Art Assessment, http://www.satoc.eu/projects/samosa/docs/SAMOSA-TN01-V1.0full.pdf/

CryoSat User Workshop proceedings, SP-717, 2014, http://www.spacebooks-online.com/product_info.php?products_id=17581&osCsid=sscldrhw/

Dinardo, S. and J. Benveniste, Guidelines for the SAR (Delay-Doppler) L1b Processing, ESA XCRY-GSEG-EOPS-TN-14-0042, Is. 2.3, 29/05/2013.

https://wiki.services.eoportal.org/tiki-download_wiki_attachment.php?attId=2540/

Jensen, J. R., and R. K. Raney, "Delay Doppler radar altimeter: Better measurement precision," in Proceedings IEEE Geoscience and Remote Sensing Symposium IGARSS'98. Seattle, WA: IEEE, 1998, pp. 2011-2013.

Martin-Puig, C., Marquez, J., Ruffini, G., Raney, K.R., and Benveniste, J, SAR Altimetry Applications over Water, ESA SeaSAR Workshop, 21-25 January, SP-656, 2008.

Raney, R. K., "CryoSat SAR-Mode Looks Revisited," Proceedings, ESA Living Planet Symposium, Bergen, Norway, 2010.

Raney, R. K., "CryoSat SAR-Mode Looks Revisited," IEEE Geoscience and Remote Sensing Letters, vol. 9, pp. 393-397, 2012.

Raney, R. K., "Maximizing the intrinsic precision of radar altimetric measurements," IEEE Geoscience and Remote Sensing Letters, vol. 10, pp. 1171-1174, 2013.

Ray, C., C. Martin-Puig, M. P. clarizia, G. Ruffini, S. Dinardo, C. Gommenginger and J. Benveniste, SAR Altimeter Backscattered Waveform Model, IEEE Trans. GeoSci. And Rem. Sens., Vol. 53, Iss. 2., pp 911 – 919, 2014. DOI: 10.1109/TGRS.2014.2330423.

Wingham D. J., Francis C.R., Baker S., Bouzinac C., Cullen R., de Chateau-Thierry P., Laxon S.W., Mallow U., Mavrocordatos C., Phalippou L., Ratier G., Rey L., Rostan F., Viau P., and Wallis D., CryoSat: A Mission to Determine the Fluctuations in Earth's Land and Marine Ice Fields. *Advances in Space Research* 37 (2006) 841-871.

9. UNIVERSAL RESOURCE LOCATORS (URL)

SEOM web site	http://seom.esa.int/
ESA Earth Online	http://eopi.esa.int/
Sentinels Online	http://sentinel.esa.int/
Copernicus	http://www.copernicus.eu
CP4O	http://www.satoc.eu/projects/CP4O/
CRUCIAL	https://research.ncl.ac.uk/geodesy/research/projects/altimetricenvironmentalmonitoring/crucial/
Coastal Altimetry community	http://www.coastalt.eu
Coastal Altimetry Workshops	http:////www.coastalaltimetry.org
RADS	http://rads.tudelft.nl
SAMOSA	http://www.satoc.eu/projects/samosa/
AVISO+	http://www.aviso.altimetry.fr/

AN OMNIBUS LIKELIHOOD RATIO TEST STATISTIC AND ITS FACTORIZATION FOR CHANGE DETECTION IN TIME SERIES OF POLARIMETRIC SAR DATA

Allan Aasbjerg Nielsen[a], Knut Conradsen[a], and Henning Skriver[b]

Technical University of Denmark
[a]DTU Compute – Applied Mathematics and Computer Science
[b]DTU Space – National Space Institute
DK-2800 Kgs. Lyngby, Denmark

ABSTRACT

Based on an omnibus likelihood ratio test statistic for the equality of several variance-covariance matrices following the complex Wishart distribution with an associated p-value and a factorization of this test statistic, change analysis in a short sequence of multilook, polarimetric SAR data in the covariance matrix representation is carried out. The omnibus test statistic and its factorization detect if and when change(s) occur. The technique is demonstrated on airborne EMISAR L-band data but may be applied to Sentinel-1, Cosmo-SkyMed, TerraSAR-X, ALOS and RadarSat-2 or other dual- and quad/full-pol, and even single-pol data also.

1. INTRODUCTION

In earlier publications we have described a test statistic for the equality of two variance-covariance matrices following the complex Wishart distribution with an associated p-value [1]. We showed their application to bitemporal change detection and to edge detection [2] in multilook, polarimetric synthetic aperture radar (SAR) data in the covariance matrix representation. The test statistic and the associated p-value is described in [3] also. In [4] we focused on the block-diagonal case, we elaborated on some computer implementation issues, and we gave examples on the application to change detection in both full and dual polarization bitemporal, bifrequency, multilook SAR data.

In [5] we described an omnibus test statistic Q for the equality of $k \geq 2$ variance-covariance matrices following the complex Wishart distribution. We also described a factorization of $Q = \prod_{j=2}^{k} R_j$ where Q and R_j determine if and when a difference occurs. Additionally, we gave p-values for Q and R_j. Finally, we demonstrated the use of Q and R_j and the p-values to change detection in truly multitemporal, full polarization SAR data.

For more references to change detection in polarimetric SAR data, see [5].

The methods may be applied to other polarimetric SAR data also such as data from Sentinel-1, COSMO-SkyMed, TerraSAR-X, ALOS, and RadarSat-2 and also to single-pol data.

2. TEST STATISTICS AND THEIR DISTRIBUTIONS

This section gives the main results from [5]. The average covariance matrix for multilook polarimetric SAR is defined as [6]

$$\langle \boldsymbol{C} \rangle = \begin{bmatrix} \langle S_{hh}S_{hh}^* \rangle & \langle S_{hh}S_{hv}^* \rangle & \langle S_{hh}S_{vv}^* \rangle \\ \langle S_{hv}S_{hh}^* \rangle & \langle S_{hv}S_{hv}^* \rangle & \langle S_{hv}S_{vv}^* \rangle \\ \langle S_{vv}S_{hh}^* \rangle & \langle S_{vv}S_{hv}^* \rangle & \langle S_{vv}S_{vv}^* \rangle \end{bmatrix}$$

where $\langle \cdot \rangle$ denotes ensemble averaging and * denotes complex conjugation. S_{rt} denotes the complex scattering amplitude for receive and transmit polarization ($r, t \in \{h, v\}$ for horizontal and vertical polarization).

2.1. Test for equality of several complex covariance matrices

To test whether a series of $k \geq 2$ complex variance-covariance matrices $\boldsymbol{\Sigma}_i$ are equal, i.e., to test the null hypothesis

$$H_0 : \boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_2 = \cdots = \boldsymbol{\Sigma}_k$$

against all alternatives, we use the following omnibus test statistic (for the real case see [7]; for the case with two complex matrices see [1, 2]; $|\cdot|$ denotes the determinant)

$$Q = \left\{ k^{pk} \frac{\prod_{i=1}^{k} |\boldsymbol{X}_i|}{|\boldsymbol{X}|^k} \right\}^n .$$

Here the $\boldsymbol{\Sigma}_i$ (and the $\boldsymbol{X}_i$) are p by p ($p = 3$ for full pol data, $p = 2$ for dual pol data, and $p = 1$ for single channel power data), and the $\boldsymbol{X}_i = n\hat{\boldsymbol{\Sigma}}_i = n\langle \boldsymbol{C} \rangle_i$ follow the complex Wishart distribution, i.e., $\boldsymbol{X}_i \sim W_C(p, n, \boldsymbol{\Sigma}_i)$. n is the equivalent number of looks. Further, $\boldsymbol{X} = \sum_{i=1}^{k} \boldsymbol{X}_i \sim W_C(p, nk, \boldsymbol{\Sigma})$. If the hypothesis is true ("under H_0" in statistical parlance), $\hat{\boldsymbol{\Sigma}} = \boldsymbol{X}/(kn)$. $Q \in [0, 1]$ with $Q = 1$ for equality.

doi: 10.2788/854791

Fig. 1. RGB images of diagonal elements of the L-band data March, April, May (top row, left to right), June, July, August (bottom row, left to right).

For the logarithm of the test statistic we get

$$\ln Q \quad = \quad n\left\{pk\ln k + \sum_{i=1}^{k}\ln|\boldsymbol{X}_i| - k\ln|\boldsymbol{X}|\right\}.$$

A simple expression for the probability of finding a smaller value of $-2\ln Q$ is ($z = -2\ln q_{\text{obs}}$)

$$P\{-2\ln Q \le z\} \quad \simeq \quad P\{\chi^2((k-1)p^2) \le z\}.$$

A better a approximation for P can be obtained. Setting

$$\begin{aligned} f &= (k-1)p^2 \\ \rho &= 1 - \frac{(2p^2-1)}{6(k-1)p}\left(\frac{k}{n} - \frac{1}{nk}\right) \\ \omega_2 &= \frac{p^2(p^2-1)}{24\rho^2}\left(\frac{k}{n^2} - \frac{1}{(nk)^2}\right) - \frac{p^2(k-1)}{4}\left(1-\frac{1}{\rho}\right)^2 \end{aligned}$$

the probability of finding a smaller value of $-2\rho\ln Q$ is ($z = -2\rho\ln q_{\text{obs}}$)

$$\begin{aligned} P\{-2\rho\ln Q \le z\} \;\simeq\; & P\{\chi^2(f) \le z\} \\ & + \omega_2[P\{\chi^2(f+4) \le z\} - P\{\chi^2(f) \le z\}]. \end{aligned}$$

$P\{-2\rho\ln Q \le -2\rho\ln q_{\text{obs}}\} = P\{Q \ge q_{\text{obs}}\}$ is the change probability, $1-P\{-2\rho\ln Q \le -2\rho\ln q_{\text{obs}}\} = P\{Q < q_{\text{obs}}\}$ is the no-change probability.

2.2. Test for equality of first $j < k$ complex covariance matrices

If the above test shows that we cannot reject the hypothesis of equality, no change has occurred over the time span covered by the data. If we can reject the hypothesis, change has occurred at some time point. To test whether the first j complex variance-covariance matrices $\boldsymbol{\Sigma}_i$ are equal, i.e., given that

$$\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_2 = \cdots = \boldsymbol{\Sigma}_{j-1}$$

then the likelihood ratio test statistic R_j for testing the hypothesis

$$H_{0,j} : \boldsymbol{\Sigma}_j = \boldsymbol{\Sigma}_1 \;\text{ against }\; H_{1,j} : \boldsymbol{\Sigma}_j \ne \boldsymbol{\Sigma}_1$$

is

$$R_j \quad = \quad \left\{\frac{j^{jp}}{(j-1)^{(j-1)p}}\frac{|\boldsymbol{X}_1 + \cdots + \boldsymbol{X}_{j-1}|^{(j-1)}|\boldsymbol{X}_j|}{|\boldsymbol{X}_1 + \cdots + \boldsymbol{X}_j|^j}\right\}^n$$

doi: 10.2788/854791

Table 1. Part of the change analysis structure for an example with data from six time points.

	$t_1 = \cdots = t_6$	$t_2 = \cdots = t_6$	$t_3 = \cdots = t_6$	$t_4 = \cdots = t_6$	$t_5 = t_6$
Omnibus	$Q^{(1)}$: $P\{Q^{(1)} < q^{(1)}_{\rm obs}\}$	$Q^{(2)}$: $P\{Q^{(2)} < q^{(2)}_{\rm obs}\}$	$Q^{(3)}$: $P\{Q^{(3)} < q^{(3)}_{\rm obs}\}$	$Q^{(4)}$: $P\{Q^{(4)} < q^{(4)}_{\rm obs}\}$	$Q^{(5)}$: $P\{Q^{(5)} < q^{(5)}_{\rm obs}\}$
$t_1 = t_2$	$R^{(1)}_2$: $P\{R^{(1)}_2 < z^{(1)}_{2,\rm obs}\}$				
$t_2 = t_3$	$R^{(1)}_3$: $P\{R^{(1)}_3 < z^{(1)}_{3,\rm obs}\}$	$R^{(2)}_2$: $P\{R^{(2)}_2 < z^{(2)}_{2,\rm obs}\}$			
$t_3 = t_4$	$R^{(1)}_4$: $P\{R^{(1)}_4 < z^{(1)}_{4,\rm obs}\}$	$R^{(2)}_3$: $P\{R^{(2)}_3 < z^{(2)}_{3,\rm obs}\}$	$R^{(3)}_2$: $P\{R^{(3)}_2 < z^{(3)}_{2,\rm obs}\}$		
$t_4 = t_5$	$R^{(1)}_5$: $P\{R^{(1)}_5 < z^{(1)}_{5,\rm obs}\}$	$R^{(2)}_4$: $P\{R^{(2)}_4 < z^{(2)}_{4,\rm obs}\}$	$R^{(3)}_3$: $P\{R^{(3)}_3 < z^{(3)}_{3,\rm obs}\}$	$R^{(4)}_2$: $P\{R^{(4)}_2 < z^{(4)}_{2,\rm obs}\}$	
$t_5 = t_6$	$R^{(1)}_6$: $P\{R^{(1)}_6 < z^{(1)}_{6,\rm obs}\}$	$R^{(2)}_5$: $P\{R^{(2)}_5 < z^{(2)}_{5,\rm obs}\}$	$R^{(3)}_4$: $P\{R^{(3)}_4 < z^{(3)}_{4,\rm obs}\}$	$R^{(4)}_3$: $P\{R^{(4)}_3 < z^{(4)}_{3,\rm obs}\}$	$R^{(5)}_2$: $P\{R^{(5)}_2 < z^{(5)}_{2,\rm obs}\}$

Table 2. Average no-change probabilities for the grass field.

	$t_1 = \cdots = t_6$	$t_2 = \cdots = t_6$	$t_3 = \cdots = t_6$	$t_4 = \cdots = t_6$	$t_5 = t_6$
Omnibus	0.0003	0.0010	0.0210	0.0653	0.0791
$t_1 = t_2$	0.2753				
$t_2 = t_3$	0.0171	0.0784			
$t_3 = t_4$	0.0341	0.0895	0.2688		
$t_4 = t_5$	0.0015	0.0048	0.0309	0.1287	
$t_5 = t_6$	0.3565	0.3016	0.2184	0.1521	0.0791

or

$$\begin{aligned}\ln R_j &= n\{p(j\ln j - (j-1)\ln(j-1)) \\ &\quad + (j-1)\ln|\sum_{i=1}^{j-1}\boldsymbol{X}_i| + \ln|\boldsymbol{X}_j| - j\ln|\sum_{i=1}^{j}\boldsymbol{X}_i|\}.\end{aligned}$$

Furthermore, the R_j constitute a factorization of Q

$$Q = \prod_{j=2}^{k} R_j$$

or $\ln Q = \sum_{j=2}^{k}\ln R_j$. If H_0 is true the R_j are independent. A simple expression for the probability of finding a smaller value of $-2\ln R_j$ is ($z_j = -2\ln r_{j,\rm obs}$)

$$P\{-2\ln R_j \le z_j\} \simeq P\{\chi^2(p^2) \le z_j\}.$$

A better a approximation for P can be obtained. Letting

$$\begin{aligned} f &= p^2 \\ \rho_j &= 1 - \frac{2p^2-1}{6pn}\left(1+\frac{1}{j(j-1)}\right) \\ \omega_{2j} &= -\frac{p^2}{4}\left(1-\frac{1}{\rho_j}\right)^2 \\ &\quad + \frac{1}{24n^2}p^2(p^2-1)\left(1+\frac{2j-1}{j^2(j-1)^2}\right)\frac{1}{\rho_j^2}\end{aligned}$$

we get ($z_j = -2\rho_j\ln r_{j,\rm obs}$)

$$\begin{aligned}P\{-2\rho_j\ln R_j \le z_j\} &\simeq P\{\chi^2(f) \le z_j\} \\ &\quad + \omega_{2j}[P\{\chi^2(f+4) \le z_j\} - P\{\chi^2(f) \le z_j\}].\end{aligned}$$

3. CHANGE VISUALIZATION EXAMPLES

To illustrate the above we use full polarimetry EMISAR [8,9] L-band data acquired in 1998 over a Danish agricultural test site on $t_1 = 21$ March, $t_2 = 17$ April, $t_3 = 20$ May, $t_4 = 16$ June, $t_5 = 15$ July, and $t_6 = 16$ August. Figure 1 shows the diagonal elements of the covariance matrix. $\langle S_{hv}S^*_{hv}\rangle$ (red) is stretched linearly between –36 dB and –6 dB, $\langle S_{hh}S^*_{hh}\rangle$ (green) between –30 dB and 0 dB and $\langle S_{vv}S^*_{vv}\rangle$ (blue) between –24 dB and 0 dB. The darker areas in the March and April images are bare surfaces corresponding to spring crops, and the very bright areas in all images are forest areas, primarily coniferous forest. The development of the crops during the growing season is clearly seen in the series of images from March to August.

Table 1 shows the change structure built (for each pixel) for an example with data from six time points. The first column indicates which tests are performed for the row in question. The second column shows $Q^{(1)}$ and $P\{Q^{(1)} < q_{\rm obs}\}$ ("Omnibus" row), or $R^{(1)}_j$ and $P\{R^{(1)}_j < r_{j,\rm obs}\}$, $j = 2,\ldots,6$ for all time points t_1 through t_6. The third column shows $Q^{(2)}$ and $P\{Q^{(2)} < q_{\rm obs}\}$ ("Omnibus" row), or $R^{(2)}_j$ and $P\{R^{(2)}_j < r_{j,\rm obs}\}$, $j = 2,\ldots,5$ for time points t_2 through t_6. The fourth column shows $Q^{(3)}$ and $P\{Q^{(3)} < q_{\rm obs}\}$ ("Omnibus" row), or $R^{(3)}_j$ and $P\{R^{(3)}_j < r_{j,\rm obs}\}$, $j = 2,\ldots,4$ for time points t_3 to t_6, etc. Remember, that for a test for $R^{(\ell)}_j$ to be valid, all previous tests for $R^{(\ell)}_i$, $i = 2,\ldots,j-1$ must show equality, see hypothesis $H_{0,j}$ in Section 2.2.

Note, that $R^{(\ell)}_2$ are the (marginal, non-omnibus) pairwise tests for equality.

doi: 10.2788/854791

(a) $-2\rho \ln Q$

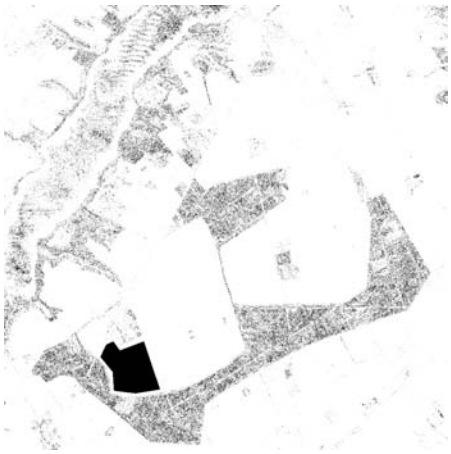

(b) p-value

Fig. 2. Test statistic (a) and p-value with grass field marked as black (b). Dark areas are no-change. p is approximately 1 in the grass field.

3.1. Per pixel change visualization

As examples of per pixel change visualization, Figure 2 shows the quantity $-2\rho \ln Q$ and the corresponding p-value, i.e., the change probability. Figure 3 shows changes from t_1 to t_2 as blue, from t_3 to t_4 as green, and from t_5 to t_6 as red after applying a 3 by 3 mode filter. Black areas have not changed.

3.2. Per field change visualization

Table 2 shows the average no-change probabilities for the grass field shown in Figure 2. Table 2 shows that the pairwise tests reveal no change over time for the grass field (p-values are 0.2753, 0.0784, 0.2688, 0.1287 and 0.0791, respectively). The omnibus test statistic Q indicates change at some time point between March and August ($P\{Q^{(1)} < q^{(1)}_{\text{obs}}\} = 0.0003$), and the R_j show that the first change for this field occurs between April and May

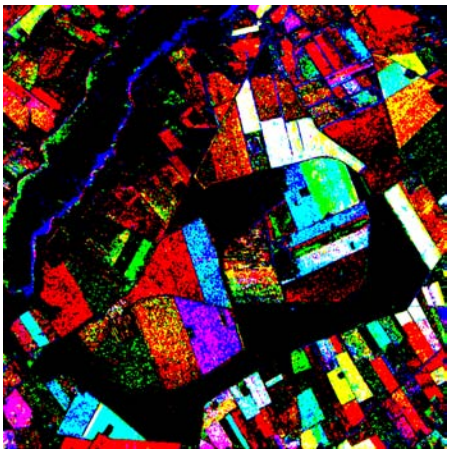

Fig. 3. Shows changes from t_1 to t_2 as blue, from t_3 to t_4 as green, from t_5 to t_6 as red (after application of a 3 by 3 mode filter); change probability significance level is 99.99%.

($P\{R^{(1)}_3 \leq r^{(1)}_{3,\text{obs}}\} = 0.0171$). The second and last change for this field occurs between June and July ($P\{Q^{(2)} < q^{(2)}_{\text{obs}}\} = 0.0010$ and $P\{R^{(2)}_4 \leq r^{(2)}_{4,\text{obs}}\} = 0.0048$).

4. REFERENCES

[1] K. Conradsen, A. A. Nielsen, J. Schou, and H. Skriver, "A test statistic in the complex Wishart distribution and its application to change detection in polarimetric SAR data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, no. 1, pp. 4–19, Jan. 2003, http://www.imm.dtu.dk/pubdb/p.php?1219.

[2] J. Schou, H. Skriver, A. A. Nielsen, and K. Conradsen, "CFAR edge detector for polarimetric SAR images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, no. 1, pp. 20–32, Jan. 2003, http://www.imm.dtu.dk/pubdb/p.php?1224.

[3] M. J. Canty, *Image Analysis, Classification and Change Detection in Remote Sensing,with Algorithms for ENVI/IDL and Python*, Taylor & Francis, CRC Press, third revised edition, 2014.

[4] A. A. Nielsen, K. Conradsen, and H. Skriver, "Change detection in full and dual polarization, single- and multi-frequency SAR data," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 8, pp. 4041–4048, Aug. 2015, http://www.imm.dtu.dk/pubdb/p.php?6827.

[5] K. Conradsen, A. A. Nielsen, and H. Skriver, "Determining the points of change in time series of polarimetric SAR data," *Accepted for IEEE Transactions on Geoscience and Remote Sensing*, 2015, http://www.imm.dtu.dk/pubdb/p.php?6825.

[6] J. J. van Zyl and F. T. Ulaby, "Scattering matrix representation for simple targets," in *Radar Polarimetry for Geoscience Applications*, F. T. Ulaby and C. Elachi, Eds. Artech, Norwood, MA, 1990.

[7] T. W. Anderson, *An Introduction to Multivariate Statistical Analysis*, John Wiley, New York, third edition, 2003.

[8] S. N. Madsen, E. L. Christensen, N. Skou, and J. Dall, "The Danish SAR system: Design and initial tests," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 29, pp. 417–476, 1991.

[9] E. L. Christensen, N. Skou, J. Dall, K. Woelders, J. H. Jørgensen, J. Granholm, and S. N. Madsen, "EMISAR: An absolutely calibrated polarimetric L- and C-band SAR," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 36, pp. 1852–1865, 1998.

doi: 10.2788/854791

BIG DATA FROM MARS: DESIGN OF GLOBAL PLANETARY DATA ANALYSIS TOOLS

Panagiotis Sidiropoulos and Jan-Peter Muller

Imaging Group, Mullard Space Science Laboratory, Dept. of Space and Climate Physics, University College London, Holmbury St. Mary, Dorking, Surrey, RH5 6NT, UK
p.sidiropoulos@ucl.ac.uk, j.muller@ucl.ac.uk

ABSTRACT

As data recording, storage and transmission capabilities increase and space exploration involve more and more space agencies, the one-at-a-time processing of planetary data from expert scientists is becoming more and more obsolete, since it unnecessarily limits the amount of data that can be analysed. In reality, large-scale batch processing of planetary data is expected to significantly increase our understanding of planetary phenomena through the detection of data features that become apparent only when processing takes place on a global scale. In this work, we report on ongoing efforts to build software tools and design strategies in this direction. Our analysis is focused on Mars, which is the planetary object with the highest data abundance for the time being, thus being an excellent case study for the new planetary science data processing paradigm shift. Preliminary results are demonstrated here to show the feasibility of this approach as well as the scientific merit of such an approach.

Index Terms— Mars, Orbital data, image registration, change detection, batch processing

1. INTRODUCTION

Until recently, the paradigm for processing planetary data coming from solar system objects was the one-at-a-time processing of incoming products by expert scientists, who were interested in features that have scientific implications that are present in the data. Such an approach was justifiable at a time when planetary data was sparse, and new data was acquired at a speed that a single expert scientist could thoroughly examine a large part of input data without having to resort to any automated data analysis tools.

However, recent advances in data acquisition, recording, storage and transmission have led to a significant increase in the volume of planetary data, their variety as well as the rate that they are acquired. In order to illustrate an example we focus on high-resolution orbital data from Mars. It is reported that by the end of 2014 more than 400,000 high-resolution images were acquired, covering an area larger than 5 times the overall area of Mars [1]. If low-quality images and other types of data are taken into account, the amount of current Mars orbital data reaches 150Tb, with the 5 present-day missions from 3 different space agencies adding more than 20 Tb per year.

While this data abundance generates great opportunities for new scientific discoveries, the community seems to struggle to catch up with the technological progress (still mainly using the traditional one-at-a-time processing), largely due to a number of unresolved issues that the "batch-mode" planetary data analysis presents. In this work, a planetary science big-data case study is presented, focusing on high-resolution orbital image data from Mars. Our approach aims on two main objectives. The first objective is to systematically organise all multi-instrument high-resolution images, thus allowing the full exploitation of the complete available dataset, while the second is to use this dataset to perform automatic detection of the changes on the Martian surface at a global scale.

The rest of the paper is structured as follows. In Section 2 we discuss the multi-instrument input data processing that is required before global-scale change detection, which is presented in Section 3. Experimental results in the MC11-E quadrangle are presented in Section 4, while, finally, in Section 5 the current status and the work that will be done in the future is reported.

2. CO-REGISTRATION OF MULT-INSTRUMENT IMAGE DATASETS TO A COMMON BASELINE

The high-resolution Mars orbital imagery consists of more than 400,000 images [1], acquired by 5 NASA and 1 ESA missions. Even though all missions record and release housekeeping data, including the calculated location of Mars that the high-resolution camera is pointing towards, small pointing errors may trigger large discrepancies between the nominal and the actual location, which may exceed in some cases 1,000 pixels. The result is that each image is actually in its own coordinate system. This fact hinders the straightforward comparison between images that supposedly map the same Mars region, thus preventing the development of methods re-

This research has received funding from the STFC MSSL Consolidated Grant ST/K000977/1 and from the European Unions Seventh Framework Programme (FP7/2007-2013) under iMars grant agreement n 607379.

doi: 10.2788/854791

lated to content-based image retrieval (CBIR) [2], change detection [3], etc.

The state-of-the-art way to deal with this discrepancy when more than one image of the same area is examined is to perform a tedious, manual and not systematic co-registration of the input images into the same coordinate system. For example, such a process was followed for the production of the mosaic of MC11-E rectangle of Mars [4] from 89 input images from the HRSC camera [5] on board ESA's Mars Express mission. The fact that the total amount of high-resolution Martian data is 3 orders of magnitude larger than this illustrates the limitations of such an approach.

Therefore, we have developed an automated processing technique that can be used in a batch-mode to project large numbers of images into a common coordinate system (i.e. a common baseline), which on Mars is determined by the HRSC orthorectified dataset. ESA's HRSC level-4 dataset is used to define the baseline, since it is so far the only high-resolution camera that acquires stereo photogrammetric images from Mars.

This pipeline aims at achieving batch-mode co-registration of high-resolution images, therefore it incorporates methods that minimise the required time, without compromising the spatial accuracy. Most importantly, it avoids the use of computationally too demanding co-registration methods based on pixel cross-correlation (e.g. [6]), which may exhibit good accuracy but are very slow. Instead, image matching is based on the extraction of feature points from both the high-resolution input image and the HRSC nadir image and their subsequent matching.

Even though the descriptor used for this purpose is the original Scale Invariant Feature Transform (SIFT) [7], a number of adjustments were made in order to allow the pipeline to handle such a task. One of the most important was the introduction of geometrical constraints to the SIFT point matching process, as discussed in [8], which reduced the algorithm's computational time while at the same time increasing the number of correct matches between the images. Additionally, the pipeline makes use of multiple copies of the input high-resolution images, which are created by varying the sampling rate, the level of processing, the image orientation, etc., in order to maximise the number of matching points that are finally extracted.

After matching points between the input high-resolution image and the corresponding HRSC level-4 nadir image are extracted, these are transformed from 2D pixel coordinates to 3D world coordinates, using the areo-reference information of the HRSC nadir image and the corresponding HRSC DTM. This leads to a number of correspondences between the input high-resolution pixels and their position in 3D world coordinates, which are used to estimate a camera model for the input image, which subsequently determines the projection of the input image into the common coordinate system. The employed camera model is a combination of a rigorous camera model and of a polynomial model. The rigorous camera model depends on the type of the instrument that acquired the image, i.e. whether it is a frame camera (as in Viking Orbiter images) or a pushbroom camera (as in most Mars high-resolution images).

3. BATCH PROCESSING OF MARS DATA FOR CHANGE DETECTION

After the input high-resolution data are projected into a common coordinate system, multi-instrument and multi-temporal images mapping the same area can be straightforwardly compared. Image comparison is mainly performed in order to identify changes that have taken place on the surface of Mars in the time interval between the times that they were acquired. Even though image comparison is a widespread approach used to identify surface changes associated with planetary natural phenomena (e.g. [9]), once again its potential is limited by the fact that image comparison is performed one-pair-at-a-time and not in batch mode.

Therefore, change detection could also benefit from a big-data approach, that will allow the identification of surface changes on a global scale. However, in order to establish such a pipeline it is imperative to resolve a number of issues that are related to the different internal characteristics of each camera, the different lighting conditions that each image was acquired with, the different image resolution, etc. This is accomplished through a cascaded learning approach.

Keeping in mind that the change detection pipeline conducts a pairwise comparison (i.e. it takes as an input two images of the same location that were taken at different time and examines whether some semantic surface change has taken place in the time between their acquisition), the first stage of the algorithm extracts low-level image descriptors from both images. Because the straightfoward matching of low-level descriptors would be undermined by the differences in the image appearance, these are used to detect primitive features in the two images, i.e. intermediate level characteristics such as lines, circles, etc.

The latter are expected to be robust to image differences that cause differences on pixel values but not on the semantic content of the image. Therefore, they can be compared to identify changes on the images. Two different types of changes are identified, movement (e.g. dune migration [10]) and semantic shift (e.g. the creation of a new crater or the disappearance of a slope streak), using distinct algorithms for each type.

4. MC11-E PROCESSING PRELIMINARY RESULTS

The algorithms that are going to be used within the EU FP7 project iMars to process as many of the high-resolution images of the Martian surface as possible, demonstrate for the first time a big-data approach for multimedia data acquired of

doi: 10.2788/854791

Mars. Before passing on full-scale processing, these pipelines are initially tested on a new mosaic generated over the MC11-E quadrangle. MC11-E is the East half of the Oxia Palus quadrangle, extending between 0 and 30 degrees North and 0 and 22.5 West (or 337.5 to 360 degrees East) and it contains a number of the most geologically interesting regions of Mars, such as Chryse Planitia, Xanthe Terra, Mawrth Vallis, Meridiani Planum, etc.

In 2015, the HRSC team released a mosaic of MC11-E, with a resolution of 12.5m/pixel, while the corresponding DTM has a resolution of 50m/pixel [4]. This mosaic is used as a baseline to co-register all high-resolution visible images, i.e. images from Viking Orbiter, MOC-NA, THEMIS-VIS, CTX and HiRISE cameras (Table 1). In total, this processing involved 7,920 high-resolution orbital images overlapping with MC11-E, if a resolution cut-off of 100m/pixel is used to discriminate high-resolution from low-resolution images. MC11-E imagery constitute approximately 2% of all high-resolution images of Mars, a percentage that is increased to 4% if we ignore the regions of Mars for which there is currently no HRSC 3D-model (i.e. there is no basemap available).

Instrument	Images in MC11-E	Resolution (m)
VO	504	8-100
MOC-NA	1,558	1.5-12
THEMIS	3,629	17.5-75
CTX	1,365	5-6
HiRISE	864	0.25-0.5

Table 1. High-resolution orbital images of MC11-E.

Currently, the processing of CTX, MOC-NA and THEMIS-VIS products has been finished, while Viking Orbiter and HiRISE processing is still pending. The results for the datasets that have finished can be found in Table 2. Three statistics are reported:

- Failure Rate: The perecentage of images that the automatic pipeline failed to produce any results
- Accuracy: The median average accuracy in X and Y dimensions, i.e. the expected average mis-registration error for an image
- Computational Time: The average processing time, using single-core threads on machines with 1.6GHz CPU and 48Gb of RAM

Table 2 shows that the developed algorithm could handle the batch-mode processing of large amounts of input data, with an accuracy that is close both to the accuracy achieved by the tedious manual co-registration of each individual image and by the computationally demanding cross-correlation automatic techniques. For example, by extrapolating the MOC-NA processing time it can be deduced that the whole dataset of 95,966 MOC-NA images could be processed by a single 16-core machine in 4 months.

	MOC-NA	THEMIS	CTX
Failure Rate (%)	34.55	19.83	7.25
ErrX (m)	5.334	7.012	6.487
ErrY (m)	4.851	6.849	6.081
Time (min)	29	26	331

Table 2. Statistics of the batch-mode automatic co-registration of high-resolution image to HRSC baseline.

Finally, we should note that the failure rate seems to be aligned with the product quality of different instruments. MOC-NA is the oldest camera of the three while CTX the newest and the one that produces the images with the highest quality. This correlation is an early indication that the failure of the system is not caused by design faults in the pipeline but by poor quality input high-resolution images

5. CURRENT STATUS AND FUTURE WORK

As has already been mentioned, this is an ongoing work, therefore, there are parts of the discussed pipeline that are still under development. Actually, from the two processing chains, the co-registration one is close to being finalised at the time of writing, assuming that the MC11-E co-registration of Viking Orbiter and HiRISE products will not reveal currently unknown shortcomings. On the other hand, the change detection pipeline is currently under development. However, preliminary change detection results seem to be very promising about the potential of this approach (Figure 1). In the future, apart from the above pipelines, we are planning to include more data in our approach, including spectral data and infrared cameras data.

Fig. 1. Example of Mars surface change. (a) Part of HRSC image H2027_0000, acquired at 13-8-05 (b) Part of CTX image B06_011856_1834_XI_03N329W, acquired at 20-07-09.

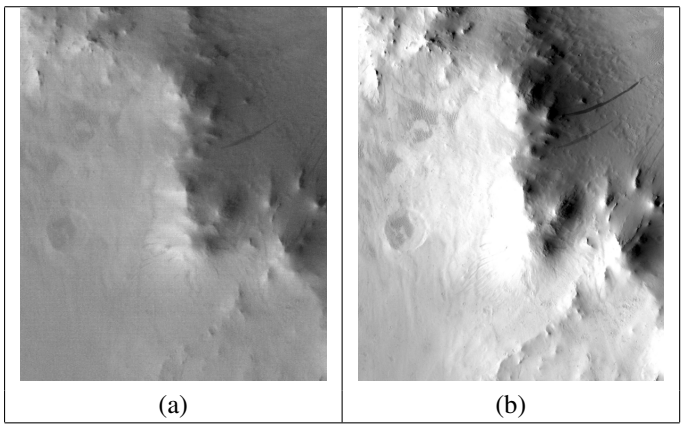

6. REFERENCES

[1] P. Sidiropoulos and J.-P. Muller, "On the status of orbital high-resolution repeat imaging of mars for the observation of dynamic surface processes," *Planetary and Space Science*, vol. 117, pp. 207–222, November 2015.

[2] A. W. M. Smeulders, M. Worring, et al., "Content-based image retrieval at the end of the early years," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 12, pp. 1349–1380, 2000.

[3] K. L. Wagstaff, J. Panetta, et al., "Dynamic landmarking for surface feature identification and change detection," *ACM Transactions on Intelligent Systems and Technology*, vol. 3, no. 3, 2012.

[4] K. Gwinner, R. Jaumann, et al., "The first Quadrangle of the Mars Express HRSC multi-orbit data products (mc-11-e)," in *European Planetary Science Congress*, 2015, vol. 10.

[5] R. Jaumann, G. Neukum, et al., "The high-resolution stereo camera (HRSC) experiment on Mars Express: Instrument aspects and experiment conduct from interplanetary cruise through the nominal mission," *Planetary and Space Science*, vol. 55, no. 7-8, pp. 928–952, 2007.

[6] G. P. Otto and T. K. Chau, "A region-growing algorithm for matching of terrain images," *Image and Vision Computing*, vol. 7, no. 2, pp. 83–94, 1989.

[7] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.

[8] P. Sidiropoulos and J.-P. Muller, "Matching of large images through coupled decomposition," *IEEE Transactions on Image Processing*, vol. 24, no. 7, pp. 2124–2139, 2015.

[9] A. McEwen, L. Ojha, et al., "Seasonal flows on warm martian slopes," *Science*, vol. 333, no. 6043, pp. 740–743, 2011.

[10] M. C. Bourke, K. S. Edgett, and B. A. Cantor, "Recent aeolian dune change on mars," *Geomorphology*, vol. 94, pp. 247–255, 2008.

FLOOD FORECASTING AND MONITORING USING SENTINEL-1 AND SMOS SATELLITES: A SUPERVISED PREDICTIVE MODELING

David Bastide [1], Jérôme Gasperi [2], Yann Kerr [3], Richard Moreno [2], Pierre Villard [1]

[1] Capgemini Aerospace and Defense, Toulouse, France – pierre.villard@capgemini.com, david.bastide@capgemini.com
[2] CNES, Toulouse, France – richard.moreno@cnes.fr, jerome.gasperi@cnes.fr
[3] CESBIO, Toulouse, France – yann.kerr@cesbio.cnes.fr

ABSTRACT

Increasingly intense downpours driven by climate change will see flood damages rise from €4.5bn to 23€bn a year by 2050 according to a research published in the journal Nature Climate Change [1]. At this time, a lot of papers are dealing about the possibilities to forecast flood events and to monitor such events as well. Our objective is to combine the two approaches in a learning system in order to have a supervised predictive modeling which is automatically improving itself over time. This paper describes our current work regarding the use of Sentinel-1 and SMOS data to build such a system using the Data Cube developed by Geoscience Australia.

Index Terms— Flood, forecasting, monitoring, Sentinel-1, SMOS, machine learning, data cube

1. INTRODUCTION

Nowadays EO data is easily and quickly accessible from a growing number of satellites. The question is how to deal with an increasing amount of data in an efficient way and to produce some interesting results that can help people taking decisions. In these proceedings we are focusing on some ideas and concepts that can be united to fulfill one objective: efficiently forecasting flood events using an algorithm that is able to learn from its mistakes thanks to a parallel flood monitoring system.

2. THE AUSTRALIAN GEOSCIENCE DATA CUBE

Geoscience Australia (GA) is leading the development of the Australian Geoscience Data Cube (AGDC) to support the management and quantitative analysis of massive volumes of Earth observation (EO) and other geo-scientific data. One of the features of the AGDC approach is that all of the original observations (pixels) are retained for analysis; the data are not mosaicked, binned, or filtered in any way and the source data for each pixel can be traced through the metadata.

The AGDC approach allows us to store heterogeneous data from multiple sensors and sources in one single place and to stack this data over time. It eases the access to the data when focusing on a specific region of interest and frees researchers so they can focus on the direct scientific application and exploitation of the data [2].

3. FLOOD RISK FORECASTING USING SMOS

On one hand, two years ago, Capgemini initiated the implementation of an algorithm developed by the CESBIO (center for the study of the biosphere from space) used to evaluate the probability of a flood event based on SMOS data and weather forecasting [3]. The idea is to compute a flood risk map based on precipitation predictions for the next 5 days merged with information on actual soil moisture conditions. The precipitation forecasts are obtained from the NCEP Forecast data and the soil moisture measurements are acquired from SMOS CATDS L3 Soil Moisture data produced at Ifremer for CNES.

The algorithm behind the risk map is purely statistical: the precipitation predictions are compared to precipitation percentiles from 40 years daily reanalysis data from NCEP. The soil moisture is used to increase or decrease the flood risk from precipitation maps based on initial soil moisture conditions. Wet to Dry initial soil moisture conditions are defined based on the comparison of actual soil moisture from SMOS to the SMOS 4-years dataset archive. So far the algorithm was validated against flood events archives from the Dartmouth Flood Observatory.

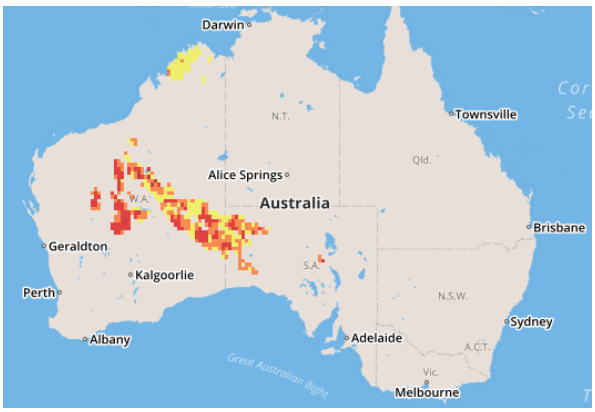

Fig. 1: Flood risk forecasting map based on SMOS data

doi: 10.2788/854791

4. FLOOD MONITORING USING SENTINEL-1

On the other hand, SAR data from Sentinel-1 satellite is easily available from multiple web portals including the French Sentinel data portal (PEPS) and is an excellent product to perform flood mapping operations.

The products are first calibrated, filtered using speckle filtering, and orthorectified. Then a split-based automatic thresholding procedure is used to classify pixels in a "water" class or in a "non-water" class [4]. At the end, an additional segmentation analysis is performed to remove potential false-positive water detections.

Eventually, we identify a new class "flooded area" by comparing the obtained results to the permanent water identified on previous acquisitions. This temporal approach provides a way to produce maps of flooded areas. It also allows the exploitation of the value of the dense Sentinel-1 time series (there will be four Sentinel-1 satellites in a close future).

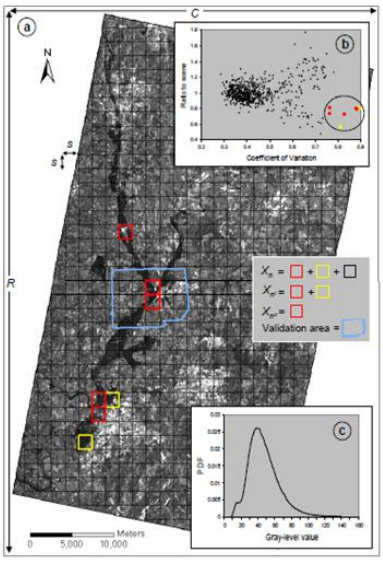

Fig. 2: NRT flood detection using a split-based automatic thresholding procedure based on Sentinel-1 data

5. APACHE SPARK

Apache Spark is an open source cluster computing framework originally developed in the AMPLab at University of California, Berkeley but was later donated to the Apache Software Foundation where it remains today.

In contrast to Hadoop's two-stage disk-based MapReduce paradigm, Spark's multi-stage in-memory primitives provides performance up to 100 times faster for certain applications. By allowing user programs to load data into a cluster's memory and query it repeatedly, Spark is well-suited to machine learning algorithms.

MLlib is Spark's scalable machine learning library consisting of common learning algorithms and utilities, including classification, regression, clustering, collaborative filtering, dimensionality reduction, as well as underlying optimization primitives.

6. SUPERVISED PREDICTIVE MODELING

So far the two systems were running operationally in a Hadoop-based infrastructure to produce both flood risk events maps and flooded areas maps. Recently, we moved from this infrastructure to the Data Cube approach to ease data access and to facilitate the interactions between the two systems.

The underlying idea is quite simple (some of our explanations are illustrated on the next pages): we developed two automatic processing chains to compute both the global flood risk forecasting map and the flood detection map; then the results are stored in a data lake which is the input of our in-memory cluster computing based on Apache Spark to run machine learning algorithms.

In this way, we are able to generate a model which gets better over time and which is able to predict flood events at a Sentinel-1 resolution instead of the SMOS resolution.

On one hand we have flood risk forecasting maps that are computed daily, each time a new L3 SMOS product is made available. On the other hand we have an automatic processing chain downloading all Sentinel-1 products related to our ROI and computing flood detection maps. The combination of the two sources allows detecting false-positive events and false-negative events in order to improve the computed model. Besides, this approach deals with all the Sentinel-1 products and, consequently, exploits the value of the dense Sentinel-1 time series.

Moreover, the use of Sentinel-1 data is a perfect solution to get a flood risk forecasting map at Sentinel-1 resolution instead of the original SMOS resolution. This way, the combination of the products gives the opportunity to have a model taking into account local characteristics of the ground, of the relief, of the land use and so on.

Consequently, this two-side system is an automatically supervised predictive modeling system which is giving more accurate results over time. The accuracy of the predictions is a key factor in decision-making: when a high risk of a flood event is detected, it gives the assurance to correctly take actions in the field, to warn people and to anticipate search and rescue operations.

This on-going project is currently supervised by CNES and the first conclusive results will be published by the end of March 2016.

doi: 10.2788/854791

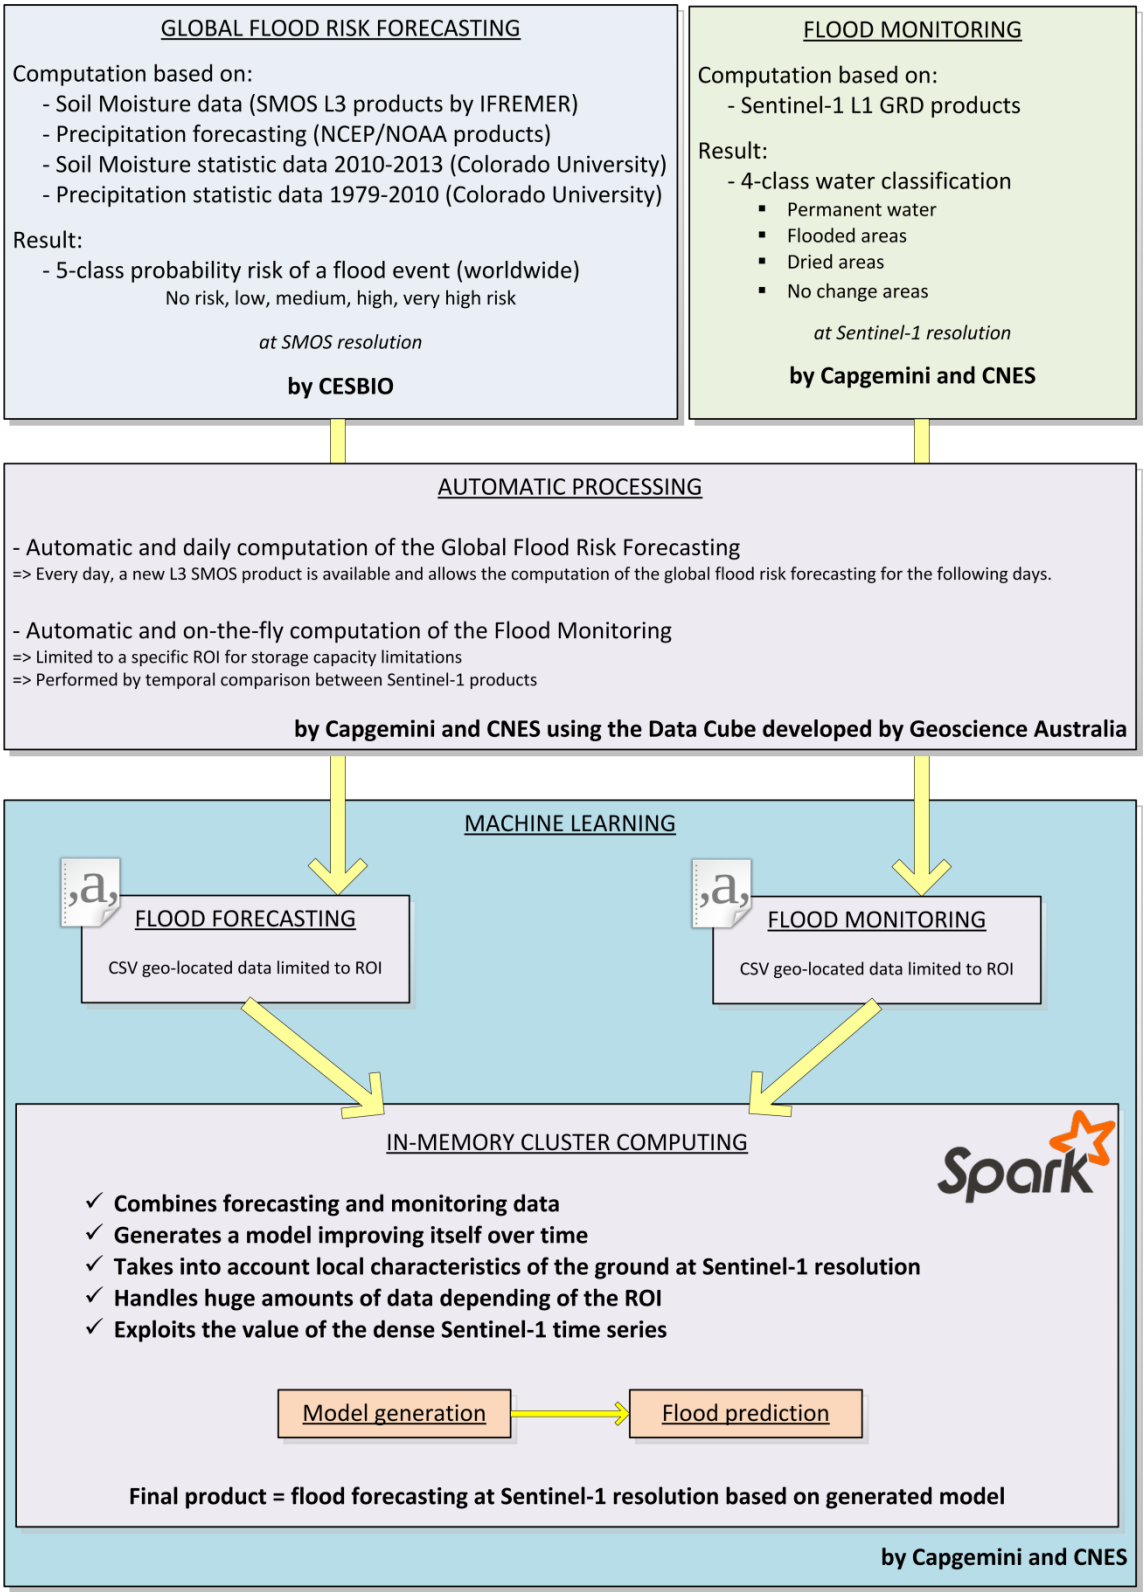

Fig. 3: Flood forecasting and monitoring using Sentinel-1 and SMOS satellites: a supervised predictive modeling

doi: 10.2788/854791

6. PRELIMINARY RESULTS

The main reasons we moved to the Data Cube is the ability to access the data very easily (pictures projected on the same grid and decomposed into smaller tiles), and the possibility to combine it with the Jupyter Notebook. It is a web application that allows you to create and share documents that contain live code (for over 40 programming languages, including those popular in Data Science such as Python, R, Julia and Scala), equations, visualizations and explanatory text. A notebook may also be used to interact with a Spark cluster. This way it allows, very easily, accessing all the data stored in the cube, to perform all the work we want and also to interact with our Spark cluster, all in one place, just using a web browser. To go further, we used Docker instances to allow the deployment of a Data Cube cluster embedding Spark and Jupyter in few minutes.

On one hand, here is an example of the detection of flooded area in India during the monsoon a few months ago:

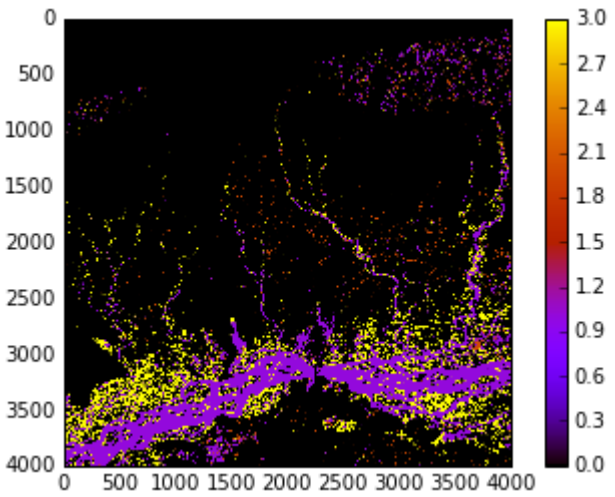

Fig. 4: Flood monitoring using Sentinel-1 data

On the other hand, here is an example of a flood risk forecasting map at global scale using the algorithm developed by CESBIO:

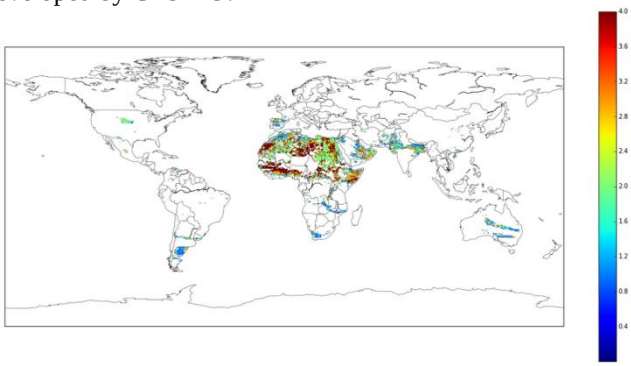

Fig. 5: Flood forecasting using SMOS data

At the moment, the process has been limited to a specific region of interest. The use of this process at a global scale is only a question of infrastructure and available resources in terms of storage, CPU and memory.

The current ROI is limited to the region of Assam in India/Bangladesh. Besides, the quality of the generated model is directly linked to the quantity of data available at the moment the model is trained. Our preliminary results and conclusions will be presented during the conference on Big Data from Space in March.

7. REFERENCES

[1] B. Jongman, S. Hochrainer-Stigler, L. Feyen, J. Aerts, R. Mechler, W.J. Botzen, L. Bouwer, G. Pflug, R. Rojas, P. Ward, “Increasing stress on disaster-risk finance due to large floods”, Nature Climate Change, March 2014, http://bit.ly/208CfoA

[2] A. Lewis, S. Oliver, A. Ip, N. Mueller, S. Ring, D. Roberts, M. Thankappan, M. Purss, “Iterating petabyte-scale earth observation processes in the Australian Geoscience Data Cube”, Proceedings of the 2014 conference on Big Data from Space (BiDS'14), ESA, http://bit.ly/1H7zLLq

[3] A. Al Bitar, A. Chone, S.K. Tomer, F. Cabot, Y. Kerr, SMOS CATDS Flood Risk, Internal report, CESBIO, 2014

[4] S. Martini, “Towards operational near real-time flood detection using a split-based automatic thresholding procedure on high resolutionTerraSAR-X data”, Nat. Hazards Earth Syst. Sci., 9, 303–314, DLR, 2009

doi: 10.2788/854791

INTERACTIVE SUB-SETTING OF 100S OF TBYTES OF LAND ECV PRODUCTS

Said Kharbouche & Jan-Peter Muller

Imaging Group, Mullard Space Science Laboratory
UCL Department of Space and Climate Physics
Holmbury St. Mary, Dorking, Surrey, United Kingdom, RH5 6NT
(s.kharbouche@ucl.ac.uk; j.muller@ucl.ac.uk)

ABSTRACT

We describe a web based system that has been developed and implemented at MSSL within the EU-FP7 QA4ECV project to facilitate the extraction of time series from our internally generated and third party big data sets by selecting a region of interest. Using a dedicated web interface, users can delimit a geographical region of interest (point, line, rectangle or polygon), select a product, specify which product layer/parameter, the spatial resolution and temporal resolution, and the time period of interest. After submission, a very large linux cluster processes the query using scheduled and simultaneous jobs. Once all jobs are completed the system merges all outputs of the jobs into a single ascii-csv file, extracting some plots from the latter and then sending a web-link to the user by email. More than 500 TBytes of data (e.g. surface reflectance, albedo, BRDF, fAPAR, LAI) are accessible via this tool. The subsetting tool has been designed from the outset to be an extensible system, so that adding a new data set only requires the creation of a set of symbolic links, the editing of a configuration file and adding the product name and its parameters in the web interface (HTML file). The system allows our end users and project partners who are interested in testing algorithms for single pixels or any small regions an effortless access to our data, and makes much easier the tasks of inter-product comparison and validation against in-situ data

Index Terms— ECVs, big data, time series, sub-setting

1. INTRODUCTION

ECVs (Essential Climate Variables) derived from Earth Observation satellites are usually packaged into files that cover wide area (global, full orbital swath or tiles), which are often considered memory consuming files, especially for the higher spatial resolution products. As a result, users interested in time series of a product over a small region, find themselves obliged to download and process voluminous files, which requires significant resources both in data transfer, local CPU and memory storage, and the complexity of any product reader software. That complexity grows with the number of RoIs (Regions of Interest). Usually, climate data users do not want to spend much time dealing with all these complexities for small regions, and prefer to get the data as quickly, and simply as possible.

The time series of a small region is needed to perform many tasks, such as validation against ground measurements, vicarious calibration, testing and validating climate model over some key sites, etc. Therefore, some climate data providers offer web-based subsetting tool to extract time series by RoI such as NASAs Global Subsetting Tool [1]. However, this tool does not operate over all available data, and the region of interest must be a rectangle with a maximum length of 201km. NASA also hosts another tool [2] with live response (no email address is required) but it is for pre-selected RoIs (i.e. countries and key sites). NOAA and ESA also provided some tools for subsetting [3, 4] but they provide subsetted products (with the same format and structure as the original) that overlap a specified users RoI.

Within the EU-FP7 QA4ECV project [5], we have implemented a web-based subsetting tool [1] with the aim of offering end users of our datasets and project partners a straightforward access to our final products, including input and intermediate products that have been used in the production, and to some third party comparable products. Hence, the user needs only to: 1) select a product from a list; 2) select all or some layers of that selected product; 3) specify a region of interest, which can be a single point, line, rectangle or polygon; 4) provide a time period of interest; 5) select a spatial resolution; 6) select a temporal resolution; and 7) provide a valid email address to receive the results. There are also other optional inputs, such as buffer values to enlarge a RoI, and specifying a name for a RoI for user convenience as well as provide titles to any of the plots. After submission of this web-form, the user receives a time series (average, standard deviation, number of valid values) for all selected layers in a single ascii-csv file, which will be accompanied by a quick-look plot for each

We acknowledge support from the EU-FP7 programme under QA4ECV grant agreement n° 607405 and partial support from the ESA GlobAlbedo project (ESA/ESRIN contract 22390/09/I-OL). Partial support was also provided from the NERC-NCEO under R8/H12/82.

[1] http://globalbedo.org/roi.php

selected layer. Furthermore, upon request, users also have the possibility to submit their queries using scripts. The next section will describe in more detail the functionality of our tool and main phases of the processing, whilst in the last section we will discuss its performance, limitations and future perspectives.

2. FUNCTIONALITY

The tool is dispersed between two locations: MSSL[2] and CEDA[3] which physically are some 90km apart. The GUI is implemented within the website www.globalbedo.org, which is hosted by MSSL. Once a user submits a query, a server at MSSL re-formats this query and sends it to CEDA where all the processing takes place (and where all our raw and final products are archived). Once the query is processed and the results (csv file and plots) are placed on a publicly accessible folder (on CEDA), a notification email is sent to the user (by the MSSL server) indicating the status of the query and the URL link to download the results. Figure 1 shows a schematic overview of the main phases from query submission to result transmission to the requester.

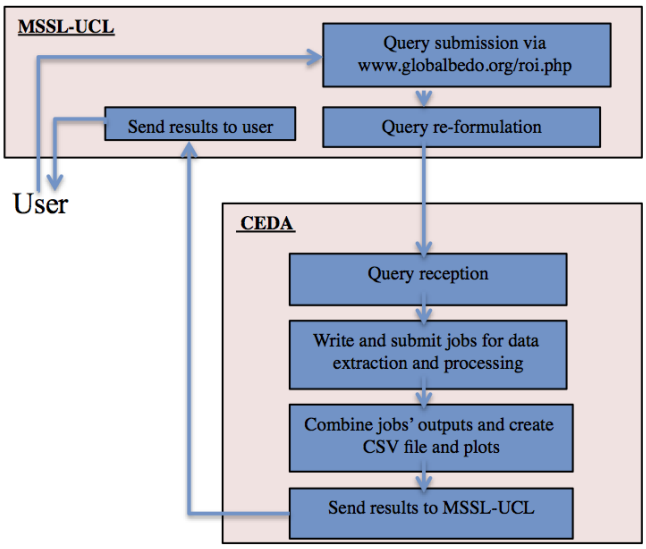

Fig. 1. Main phases of query processing.

2.1. Query submission

Figure 2 displays the web-based interface of the tool. It accepts as input:

- Product: to be selected from a list of available products. The actual list contains the following products
  - Albedos from GlobAlbedo [6];

[2]http://www.ucl.ac.uk/mssl
[3]http://www.ceda.ac.uk/

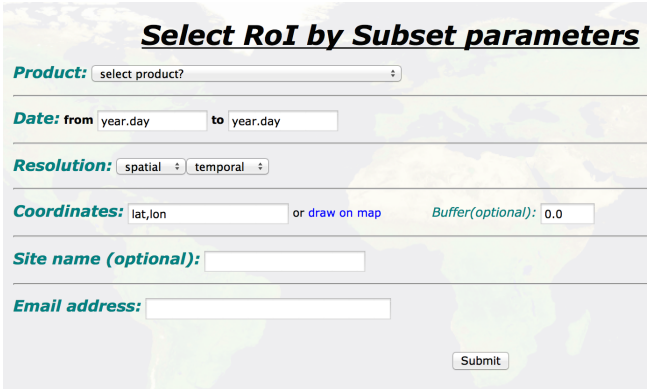

Fig. 2. Web-based form of tool (globalbedo.org/roi.php).

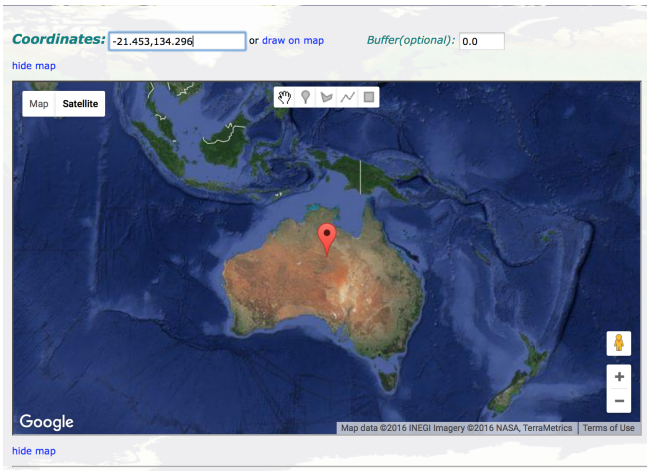

Fig. 3. User has the possibility of drawing their RoI (point, line, rectangle or poylygn) over a google map mashup window.

  - Albedos from MODIS collection 5 [7];
  - fAPAR/LAI derived from GlobAlbedo [8];
  - fAPAR/LAI of MODIS collection 6 [9];
  - BRDF prior climatology extracted from BRDF MODIS collection 5[10];
  - Surface reflectance from MODIS, Terra, collection 5 [11];
  - Surface reflectance from MODIS, Terra, collection 6 [12];

  In future we plan to add further products that are used by QA4ECV[5] either as inputs for processing or for validation:

  - Albedos from MODIS collection 6;
  - Albedos from MISR2AS;
  - Albedos from METEOSAT;

doi: 10.2788/854791

  - Surface reflectance from MERIS;
  - Surface reflectance from VGT;
  - Surface reflectance from MISR;
  - Surface reflectance from METEOSAT;

- Layer(s): to be selected or deselected from a suggested list that related to the selected product (default is all layers);
- Start and End Dates: The time period of interest. The entered date must be in the format YYYY.DDD, where YYYY denotes year (i.e. 2004) and DDD denotes day of the year (i.e. 041 = 10th of February) . The default dates are the whole time period of the selected data. The desired start/end date must be within the available period;
- Temporal resolution: this is selected from a list that specifies the selected product time-step (i.e. 8-daily);
- Spatial resolution: this is selected from a list that specifies the selected product spatial grid (i.e. 1km);
- Region of Interest (RoI): this represents the geographical region of interest. It can include geometry of :
  - Point: lat, lon
  - Line: $line : (lat_1, lon_1)...(lat_i, lon_i)...$
  - Polygon: $polygon : (lat_1, lon_1)...(lat_i, lon_i)...$

    Where (lat_i, lon_i) denotes geographic coordinates (latitude, longitude) in WGS-84 of a location for point, vertex if line, and corner if polygon. The user also has the possibility of drawing their geometry over a google map mashup window (see figure 3);
- Buffer : optional, RoIs can be used to widen a RoI by a buffer value, whose value depends on selected spatial resolution (deg or km);
- Site name: optional, this will be included in the resultant csv files and plots;
- Email Address: to whom to send the results via web URLs.

Note that, as almost all listed products have a single spatial and temporal resolution, the user often has only to select a product and define a RoI and provide a valid address for getting the whole time series of all layer(s) of the selected product,.

After submission, a MSSL server process re-formats the query and sends it to a server process at CEDA for subsequent processing.

2.2. Query Processing

Once the reformatted query arrives at the CEDA process server, a job is launched to select tiles (of selected products) that intersect or overlap the RoI, then create an ascii file of coordinates for each selected tile such that it contains all image coordinates of cells intersecting the RoI. Next, after creation of files of coordinates, multiple jobs are launched such that each job uses the files of coordinates (of previous stages) to compute averages, standard deviation, and the number of valid values for each date in a given subinterval time within the given time period of interest. In other words, the time period of interest is subdivided into several non-overlapping subintervals, and a single job is launched for any given subinterval. When all jobs are completed, another job is launched to merge all ascii files that resulted from previous jobs into a single ascii csv file, create time series plots per layer, compress all files (csv file and plots) and place them in a publicly accessible folder. Then a notification is sent to a MSSL server which then sends an email containing the URL of the results to the user (query submitter) about the results. If the query is successfully processed, the email contains the URL link to download the resulting csv file (and its plots), otherwise it includes the cause of the failure (e.g. non availability of data, too large RoI, etc). Figure 4 shows an overview of the jobs that are scheduled and launched by our tool to process each query.

Figure 5 shows the resulting csv file and plots for a query about Directional-Hemispherical diffuse Reflectance –DHR (i.e. black sky albedo) over a region of interest of whole Thetford Forest Park, UK, and a period of interest of [2000, 2011] (11 years), a temporal resolution of 8-daily, and a spatial resolution of 1 x 1 km. Without the subsetter tool the user must download 10s of GBytes to extract a few Kbytes of target data; however, it is a straightforward operation via our tool.

A bash script is available to any interested user wanting to perform the same analysis over many RoIs. It takes only one variable argument of RoI, and all other arguments (product, layers. email address...) can be fixed once in the script. Nevertheless, it is available upon request. The script is a shell linux file but it can be straightforwardly converted to any other programming language (python, java, etc).

3. DISCUSSION

At the beginning, this tool was only intended for internal use for our project partners to enable them to get speedy access to our final products for small regions of interest. Now, it is publicly accessible but with some restrictions (i.e. limitation on number of daily queries per user). However, as most users of this tool are usually experts on ECVs and they are more interested in the final csv files to feed their software for data analysis or climate modelling, we have put significant effort into reducing the time consumed rather than in the ameliora-

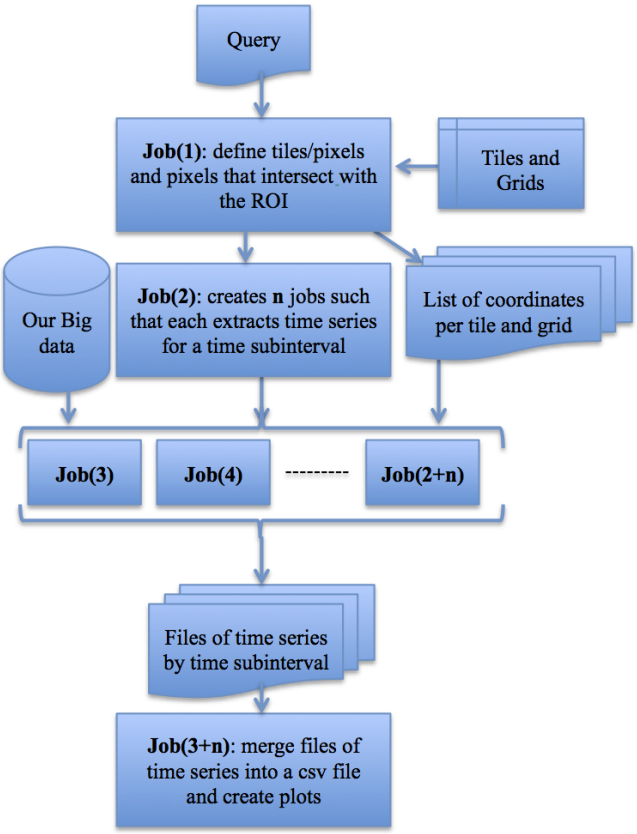

Fig. 4. Overview of the scheduled jobs under our distributed system.

tion of web-interface and data visualization.

The tool is highly appreciated by many expert users, as it allows them an effortless access to our data (now more than 500TB). But, as the tool runs on the same platform as our production line, it runs at low priority. As a consequence, the processing time depends also on the availability of current computing resources which can vary a lot from time to time. Another good thing is that the tool can easily accept any new gridded data, as we just have to create symbolic links respecting a predefined structure and by adding some parameters to a configuration file and to the web interface (product name, resolutions...). Furthermore, the processing part of the tool is based completely on java, which make it easy to import to any other distributed system. Currently, the tool is implemented primarily for GlobAlbedo products, but it will be ported for all of land products from the QA4ECV (Quality Assessment for Essential Climate Variables) project [5].

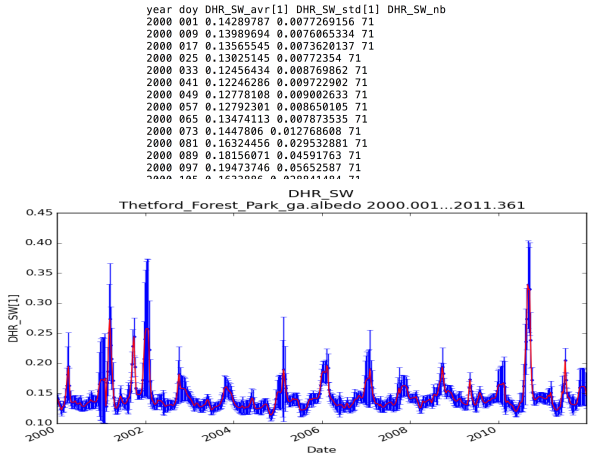

Fig. 5. Resultant time series (csv file and its plot) of DHR (black sky albedo) over a RoI of whole Thetford Forest Park, UK, during the time period 2000-2011. Tens of Gbytes of data were deployed to extract few Kbytes data of the RoI.

4. REFERENCES

[1] http://daacmodis.ornl.gov/cgi-bin/MODIS/GLBVIZ_1_Glb/modis_subset_order_global_col5.pl

[2] http://daac.ornl.gov/cgi-bin/MODIS/GR_col5_1/mod_viz.html

[3] http://merisfrs-merci-ds.eo.esa.int/merci/

[4] www.nsof.class.noaa.gov/saa/products/catSearc

[5] www.qa4ecv.eu

[6] www.globalbedo.org/docs/GlobALBEDO_Product_User_Guide_D17_v16.pdf

[7] http://lpdaac.usgs.gov/dataset_discovery/modis/modis_products_table/mcd43a3

[8] http://aramis.obspm.fr/~jimenez/Docs/WACMOSET/WACMOSET_WP2230_approved.pdf

[9] http://lpdaac.usgs.gov/dataset_discovery/modis/modis_products_table/mod15a2h_v006

[10] S. Kharbouche, J-P. Muller, and P. Lewis;, "A 15 year climatology of spectral brdf derived from modis for a priori optimal estimation of global surface albedo within the EU-FP7 QA4ECV project," in *ISRSE36-256, Berlin*, May 2015.

[11] http://lpdaac.usgs.gov/dataset_discovery/modis/modis_products_table/mod09ga_v006

[12] http://lpdaac.usgs.gov/dataset_discovery/modis/modis_products_table/myd09ga_v006

doi: 10.2788/854791

TEMPORAL CHARACTERIZATION OF THE REMOTE SENSORS RESPONSE TO RADIATION DAMAGE IN L2

Ruben De March[1], Deborah Busonero[2], Rosario Messineo[1], Alessandro Bemporad[2], Francesco Vaccarino[3], Angelo Fabio Mulone[1], Andrea Fonti[1], Mario Lattanzi[2]

[1]ALTEC S.p.A., Corso Marche 79, 10146 Torino, Italy
[2]INAF-OATo, Via Osservatorio 20, 10025 Pino Torinese, Torino, Italy
[3]Politecnico di Torino, Dept. Math. Sciences, Corso Duca degli Abruzzi 24, 10129 Torino, Italy

ABSTRACT

Remote sensors on spacecrafts acquire huge volumes of data that can be processed for other purposes in addition to those they were designed for. The project TECSEL2 was born for the usage of the Gaia AIM/AVU daily pipeline output and solar events data to characterize the response of detectors subjected to strong radiation damage within an environment not protected by the terrestrial magnetic field, the Lagrangian point L2, where Gaia operates. The project also aims at identifying anomalies in the scientific output parameters and relate them to detectors malfunctioning due to radiation damage issues correlating with solar events occurred in the same time range. TECSEL2 actually designs and implements a system based on big data technologies which are the state of art in the fields of data processing and data storage. The final goal of TECSEL2 is not only related to the Gaia project, because it provides useful analysis techniques for generic and potentially huge time series datasets.

Index Terms— TECSEL2, Gaia, DPCT, AIM, CCD, CTI, radiation, big data, time series, PCA, ICA, DFA, data processing, data storage

1. INTRODUCTION

Gaia is the ESA space mission launched on December 19th, 2013 from the French Guyana, aiming at Global Astrometry at few μas, scanning continuously the whole sky in order to build the largest, most precise three-dimensional map of our Galaxy by surveying more than a thousand million stars [1], [2]. It is the first mission to operate within the Charge Couple Devices bandwidth in L2. Gaia focal plane array with its 106 CCDs is therefore an invaluable source of information about the CCDs behaviour within a strong radiation environment.

The Astrometric Instrument Model (AIM) is one of the crucial components of the Astrometric Verification Unit (AVU), the verification counterpart operating independently from the Gaia baseline data reduction chain. The AIM system is devoted to the monitoring, diagnostic and calibration of the Gaia astrometric instrument response over the mission lifetime [3], [4]. AVU has its own dedicated Data Processing Centre in Turin, the DPCT, which is one of the six Gaia DPCs spread across Europe. DPCT is also designed to provide computation, storage, data access and operations services to Italian Gaia Science Community [5].

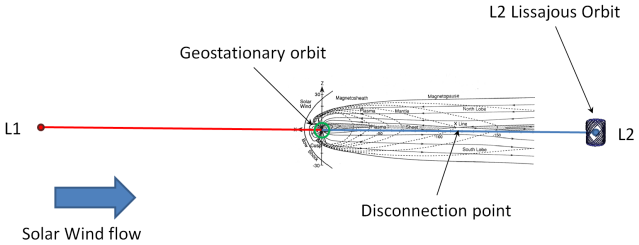

Fig. 1. L2 location vs. Earth's magnetosphere.

The availability of new big data technologies opens new scenarios in which science data collected inside a specific mission can be used to find new information and correlate it with datasets coming from different sources. The goal of TECSEL2 project is to study the possibility to use data acquired by remote sensors like CCDs to describe the environment in which the CCDs operate. The project starts with the preparation of data coming from Gaia CCDs in order to compare them with events related to the Sun.

2. RADIATION DAMAGE AND SCIENCE DATA

Solar flares and/or Coronal Mass Ejections (CMEs), quite often associated with the acceleration of Solar Energetic Particles (SEPs) released by the Sun [6], are the major responsible for disturbances on Earths magnetosphere and lead to Geomagnetic Storms. The explanations for these physical processes are far from being fully understood and are the main topic of the recently new-born Space Weather discipline.

In particular, one of the less known regions of Earth's magnetosphere is the magnetotail. The reason is that only a few satellites visited this far region and sampled the evolution

doi: 10.2788/854791

of plasma parameters before, during and after Geomagnetic Storms, so further studies of the interplanetary environment in the far magnetotail are needed.

We need to keep in mind that the main contribution to the radiation damage is due to SEPs. During the maximum solar activity period of the solar cycle, solar eruptions produce large fluxes of solar energetic particles, protons and heavy ions, which reach L2 when the eruptions are directed towards Earth. Heavy ions will not contribute significantly, but may leave traces as strong transient events. The main contributor to transient events (detected cosmic ray traces) and accumulated radiation damage over the mission lifetime will be solar protons. The expectation is that transient effects and radiation damage increase can be correlated to solar flare events.

Further information on the radiation damage effects and evolution, and its impact on the read-out images will come from nominal processing of star images themselves during the mission lifetime. Science data can be used to trace directly the instrument response, taking advantage of the repeated measurements of stars over the CCDs field, and in particular to characterize the CTI effects.

Indeed, the radiation damage changes the population traps and the charge release, degrading the charge transfer efficiency. The consequent deformation of the PSF/LSF profile introduces a centroid bias and a potential flux loss, if not calibrated. Variation in time of critical parameters like centroid, background, flux and secondary processing outputs like diagnostic image profile parameters can be related to radiation damage effects.

3. DATA PROCESSING

3.1. Datasets

The input Gaia dataset consists of observations of stars that are processed by AIM module. Through the analysis we can subset by the CCD row data are coming from, the magnitude and/or the star object type. These data are thereafter processed (averaging values of different measures inside specific time intervals) to obtain time series from the several instant-provided values. The considered datasets are derived from observations of either stars brighter than magnitude 13 (for which bi-dimensional images are acquired) or stars whose magnitude ranges between 15 and 16, i.e. stars neither too bright nor too weak, in order to maximize the detection of the effect of radiation damage. Another choice is taking data from the first row of Gaia CCDs to have less disturbance of other types, like CCDs' background level ([7]).

The space environment measurements are provided by GOES-13 and GOES-15 satellites (placed in geostationary orbit) and by ACE and WIND spacecrafts (orbiting around L1 point). These data consist of time series of particles fluxes, subdivided by energy level and type of particle, and other measures, flow pressure for example. Moreover, the

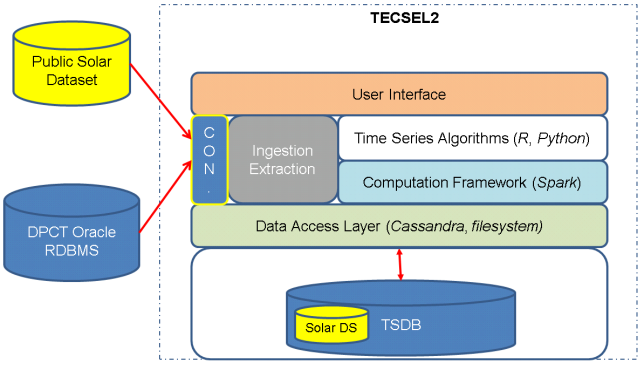

Fig. 2. TECSEL2 High Level Architecture

GOES-15 satellite provides X-ray flux data, useful to detect the occurrence of solar flares. These data are sampled by 1-minute, 5-minute or 1-hour intervals depending on the type of the series and on the satellite.

By the way, in the next Sec. 4 we will use the OMNI dataset, a dataset built by NASA with data from the spacecrafts mentioned above.

3.2. Architecture

TECSEL2 architecture is designed considering the described algorithms but also possible future integrations and other implementations. A high level description of the architecture can be seen in Fig. 2.

TECSEL2 system storage is populated through dedicated extraction/ingestion processes with solar datasets stored in files or with Oracle RDBMS used at DPCT for Gaia's data access and repository.

Since the algorithms belong to time series area, the system has to foresee a data storage component that uses a structure suitable for time series and provides scalable data access services independent from the number of time series and the data volume per series. In big data area this kind of storage system is called "time series database" (TSDB). The TSDB has been constructed comparing the available solutions and taking benefit from studies already conducted on the space sector [8].

TECSEL2 processing cluster currently consists of three nodes. This architecture takes the benefit of using Spark and Cassandra together, in particular performance is maximized when Spark runs on node where Cassandra has stored data Spark needs for its fast in-memory computations.

Given the sampling and the temporal range of the datasets, the total amount of objects stored inside Cassandra database and used in our analysis is about 20 million items.

3.3. Algorithms

Main computations involve cross-correlation and partial correlation between two paired samples, aiming at finding out

doi: 10.2788/854791

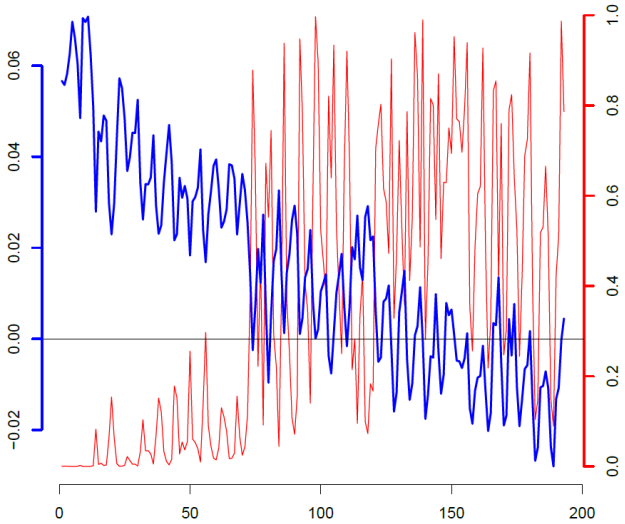

Fig. 3. Partial correlation plot between OMNI proton flux and AIM background on range Oct.10,2014-Apr.18,2015. AIM sources considered here have magnitude between 15 and 16. The left ordinate is the correlation value, while the right ordinate is the significance level (p-value). The abscissa value is the temporal shift, in hours. As you can expect, for time shifts of a few hours the partial correlation is significatively not zero, while increasing the time shift the correlation gets weaker.

if anomalous peaks (i.e., values evidently higher than usual ones) of solar wind cause any effect on Gaia CCDs. Moreover an incrementing time-shift is applied iteratively to one of the series to deduce, e.g., the delay between solar phenomena and their effect on Gaia. This delay depends on the time particles need to reach L2 point, but it can also depend on some additional time it takes for radiation damage to manifest itself. These algorithms pair datasets by entries' times, so they have to rightfully manage the instant times which lack of measures.

Eventually time series can be pre-processed, through advanced moving averages or dimensionality reduction techniques like Principal Component Analysis (PCA), Independent Component Analysis (ICA), and Dynamic Factor Analysis (DFA); or SAX representation ([9]) that may highlight a common behaviour of processed series.

4. RESULTS

First processing has been a stress test, done on a significant time range of six months from October 10th, 2014 to April 18th, 2015. The time shift interval considered is one hour and two hundred shifts have been applied to data. The variables considered in this test are AIM background and flux along with OMNI proton flux, flow pressure, and plasma temperature and speed. The partial correlation results are shown in Fig. 3.

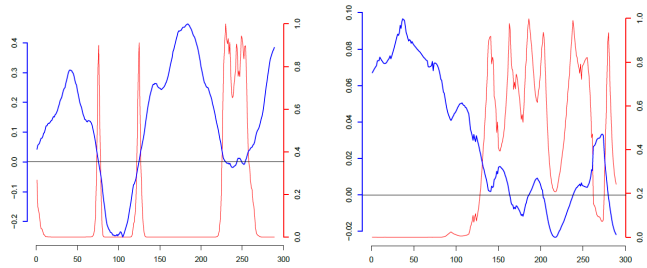

Fig. 4. Partial correlation between OMNI variables and AIM variables on Dec.12-25,2014. OMNI proton density vs. AIM (mag<13) flux; OMNI flow pressure vs. AIM (15<mag<16). background

As Fig. 3 shows, the significant correlation is present inside the first hours of shifting; so there is effectively some correlation between Gaia data and interplanetary environment measurements. In order to refine this analysis, the following results use a shorter time shift interval (5 minutes) than above. Moreover, they will be restricted on days when some peaks of particles fluxes are present, to highlight correlation vs. delay behaviour. These days have been found making a scatterplot of the proton flux variable on the entire OMNI range considered (October 1, 2014 - May 31, 2015).

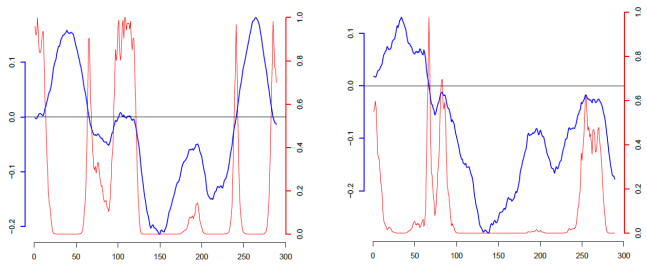

Fig. 5. Partial correlation between OMNI variables and AIM (mag<13) variables on May 11-15,2015. OMNI proton flux vs. AIM flux; OMNI proton flux vs. AIM background.

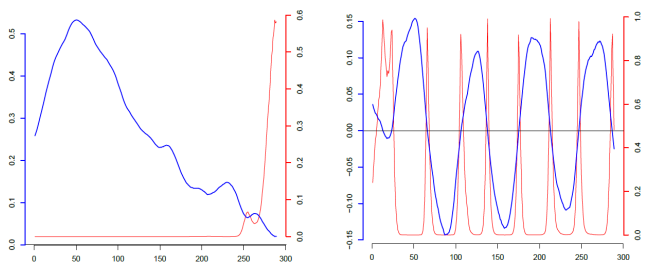

Fig. 6. Partial correlation between OMNI variables and AIM (15<mag<16) variables on May 11-15,2015. OMNI proton flux vs. AIM flux; OMNI proton flux vs. AIM background.

doi: 10.2788/854791

Considering days between December 12 and 25, 2014, when a bigger proton flux has been detected, some common patterns have been seen. In fact between 40 and 50 5-minutes delays a significant peak of correlation is found, as Fig. 4 shows. This happens for both the AIM magnitude ranges specified in Sec. 3.1.

Another interesting period for our analysis is the period between May 11 and 15, 2015. The plots about this period are in Fig. 5, as much as regards AIM data related to magnitude below 13, and Fig. 6 for magnitude between 15 and 16.

These plots highlight a correlation peak at about 45 iterations. This could mean that an higher particles activity shows its effect on Gaia almost four hours later, which could lead to a speed of 200 km/s from L1 to L2; this is compatible at least with the speed of MHD waves [10], [11]. However, further analyses are to be done with some major solar events happened after the time range we considered. Some refinements in the algorithms are also in progress to remove the periodicity from the plots, for example. This periodicity (of 6 hours, as you can see from plots) comes from the time it takes for Gaia to complete a rotation around its axis.

5. FURTHER WORK

Next steps of TECSEL2 project are both infrastructural and scientific. On one hand there is the optimization of the platform to perform data movements and analyses more efficiently, along with an improvement of interface usability.

On the other hand other time intervals containing bigger solar events are to be checked. Moreover AIM data on other CCD rows, and even other Gaia datasets may be used.

This processing can use the same algorithms introduced in Sec. 3.3, eventually with some improvements (e.g., removing time series periodicity) but also convolutive methods ([12]) or methods from topological analysis in order to study and infer the data structure from low dimensional representations.

6. CONCLUSIONS

TECSEL2 is an innovative project for several reasons. It is one of the first studies devoted to the monitoring and characterization of the behaviour of CCD detectors located in the L2 environment and it is therefore a key study for future space missions equipped with CCDs array similar to the one used on Gaia. As a service, TECSEL2 system is a powerful tool for efficient analysis of large and generic time series data, built with big data technologies, the state of art for the treatment of huge amount of data like those coming from the new generation of space and on-ground telescope.

Our results show a correlation between particles fluxes detected at L1 and charge flux and background detected on Gaia CCDs at L2 and processed by AIM. This correlation reaches its maximum with barely 4 hours of delay. Further analyses are in progress to investigate this effect.

7. REFERENCES

[1] J. de Bruijne et al., "Gaia: 1,000 million stars with 100 ccd detectors," *Proc. SPIE 7731, Space Telescopes and Instrumentation 2010: Optical, Infrared, and Millimeter Wave, 77311C*, August 05 2010.

[2] M. A. C. Perryman et al., "Gaia: Composition, formation and evolution of the galaxy," *A&A*, vol. 369, pp. 339–363, 2001.

[3] D. Busonero, "Astrometric instrument modeling in the context of gaia astrometric verification: tasks and activities," *Proc. SPIE 8449, Modeling, Systems Engineering, and Project Management for Astronomy V, 84490F*, September 25 2012.

[4] D. Busonero et al., "Monitoring, diagnostic, and calibration of the gaia astrometric instrument response within the astrometric verification unit," *Proc. SPIE 7731, Space Telescopes and Instrumentation 2010: Optical, Infrared, and Millimeter Wave, 773133*, August 10 2010.

[5] R. Messineo et al., "The italian dpc: infrastructure and operations for the italian contribution to the gaia data processing and analysis consortium," *Proc. SPIE 8451, Software and Cyberinfrastructure for Astronomy II, 84510E*, September 24 2012.

[6] R. Schwenn, "Space weather: the solar perspective," *Living Rev. Solar Phys.*, vol. 3, no. 2, 2006.

[7] T. Prusti, "Gaia: scientific in-orbit performance," *Proc. SPIE 9143, Space Telescopes and Instrumentation 2014: Optical, Infrared, and Millimeter Wave, 914310*, July 14 2014.

[8] I. Giotis et al., "Bridging big sensor data analytics from water management to space," in *Proc. Big Data from Space (BiDS14)*. IEEE, 2014.

[9] J. Lin et al., "A symbolic representation of time series, with implications for streaming algorithms," *DMKD*, pp. 2–11, June 13 2003.

[10] A. Kozlovsky et al., "On the motion of dayside auroras caused by a solar wind pressure pulse," *Annales Geophysicae*, vol. 23, pp. 509–521, 2005.

[11] O. Goncharov, "Interplanetary shock propagation through the magnetosphere to the magnetotail," *WDS'11 Proceedings of Contributed Papers, Part II*, pp. 19–26, 2011.

[12] M. Gandini et al., "Blind separation of manufacturing variability with independent component analysis: A convolutive approach," *Expert Systems with Applications*, vol. 38, pp. 9939–9946, 2011.

OSIRIDE: PROTOTYPE SYSTEM FOR CONTENT BASED IMAGE RETRIEVAL IN SATELLITE IMAGERY

Catalin Cucu-Dumitrescu[1], Florin Serban[1], Adrian Stoica[1], Corina Vaduva[2], Michele Iapaolo[3]

[1]TERRASIGNA Ltd., Bucharest
[2]University ñPolitehnicaò of Bucharest
[3]ESA ESRIN, Frascati

ABSTRACT

A functional prototype software for content based search in satellite images is presented. The system uses pre-ingested information in order to offer fast answers for queries. Search is performed at patch level and is based on image descriptors comparison. All the results presented in the paper are obtained in the context of the ESA GSTP project: Open Source Image Retrieval ï Integration of Developed Tools (OSIRIDE) [1].

***Index Terms*—** EO, CBIR, Image Descriptors.

1. INTRODUCTION

In the current technology era, there are numerous sensors specifically designed to remotely gather data that is further stored into dedicated repositories. In the field of Earth Observation (EO) petabytes of images are daily acquired. Various characteristics of the environment are highlighted in order to provide complementary aspects that are usually hidden to the human sight. Nevertheless, the access is limited due to the variety and complexity of the available information. Data mining and content retrieval techniques become mandatory for the exploration of these archives. Several solutions have been already certified in the field of text and multimedia image analysis. However, the methods are not reliable in the framework of EO data. The particularities of remotely sensed imagery require an algorithmic adjustment that cannot be generalized due to the images diversity. A good data management is fundamental for a dynamic system working with great amounts of data. This involves the development and the integration of indexing and search methods into a dynamic system allowing reconfiguration and user interaction to extract actionable intelligence information to support a variety of applications.

Several operational systems for EO data exploration and knowledge discovery have been developed in the past decade. In [2], a multilayer structure was proposed to merge the embedded features in order to support image ranking and classification. Further human integration into the processing loop offers an advantage towards linking the semantic gap between the numeric representation and the actual meaning of the data. A knowledge-driven information mining approach for remote sensing archives that addressed the EO information issue is presented in [3]. The user-centered learning concept is supporting the feature extraction and the semantic search functions. The proposed technological solution created a bridge between the data and the potential knowledge towards future data exploration [4]. However, searching for relevant knowledge across heterogeneous geospatial archives requires high dimensional database indexing. To this aim, the GeoIRIS system [5] comprise feature extraction, semantic associating from low-level features with visual characteristics and information ranking for fast queries. Recently, search engines have been evolved towards data analytics in need of a proper management and exploration of Big Data. The synergy between heterogeneous data sources such as EO imagery, metadata, semantic descriptors and linked open data has been asserted to develop a system architecture [6] focusing on transforming the spatial products into actionable intelligence.

These are some of the most successful attempts to develop a technique capable to handle EO data. They were proved to be efficient in different application scenarios. Information mining and content based image retrieval (CBIR) were validated on both optical and SAR images. Nevertheless, it is difficult to put them at work in the frame of huge, growing data archives including a variety of information. They must be adapted to deal with thousands full scenes in real time and to be able to efficiently manage the knowledge inside. The existing techniques must be expanded in order to cope with the Big Data concept. Mining from sparse, uncertain, incomplete data, learning the local relationships and knowledge sharing are the main challenges to obtain value-added products [7]. Each applied process will enhance the available information with a new layer. The user receives the lead role in data semantic modeling, imposing a customized extraction of valuable information, given a certain scenario.

To this scope, the paper presents OSIRIDE - a new CBIR system that tries to expand the data mining perspective into Big Data analytics. A modular architecture integrates several methods and algorithms to discover hidden patterns inside the

doi: 10.2788/854791

data matching particular characteristics, to visualize the data given a feature space in order to guide a query, to learn semantic dependencies based on a relevance feedback received from the user and to link the physical parameters (EO data) with semantic labels (value-added products). Even if the access to the system is made via a web platform, the user is not allowed to add new data into the repository. However, he is allowed to perform complex queries and compare his results to a predefined set of reference maps. OSIRIDE was designed to facilitate the exploitation of past, current and future generation of satellite images. The preliminary tests were thus performed on Landsat data.

2. EO IMAGE RETRIEVAL FRAMEWORK

There is a certain general workflow that must be followed in order to discover knowledge from a data collection. EO image analysis in particular combines general expertize from computer vision, image understanding, data mining, machine learning, databases and software. Precise management and a broad concept are compulsory to address the need of making big data volumes accessible and easy to control, boosting an effective use through merging the data processing similarity and the knowledge transferred by the user to the machine.

In line with this aspect, we presents an integrated platform for image retrieval. The proposed framework (Fig. 1) is in fact a use-case diagram that offers an outline on the processing chain and the client interaction with the system. The userôs role is vital, as he is conveying his understanding into semantic data modeling. Based on the human-machine interaction, the workflow is divided in three parts:

- Access web platform and database;
- Content based image retrieval and knowledge discovery;
- Exploratory data analytics and dynamic navigation.

Nevertheless, each step is individually developed, noting that all the intermediate results have been carefully indexed and coordinated. Once the user has connected to the database, a graphical interface (GUI) is guiding him through data searching and selection.

In order to proceed with the information mining specific analysis, the user must set the parameters for data model generation according to his interest. For each of the selected images, a set of features will be computed. This is a challenging task, due to the fact that searching for relevant information will be performed by an analysis of the feature space. The retrieval problem highlights relations between features of defined concepts and it is considered an interactive learning issue used to enhance the capabilities to find adequate similarities. This is leading towards new models generation, with superior semantic meaning, responding to the user's needs. The visualization tools integrated into the GUI will support the dynamic navigation and data analysis.

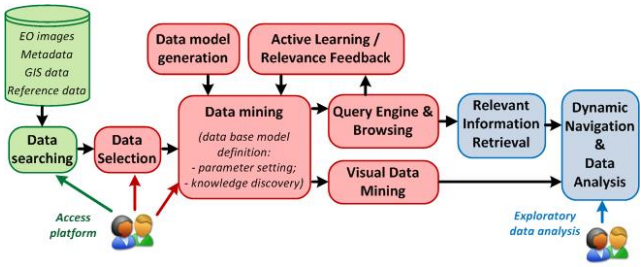

Fig. 1. EO image retrieval framework. In green, the modules enabling the access to the system. In red, the data processing. In blue, the interface guiding visualization and analytics.

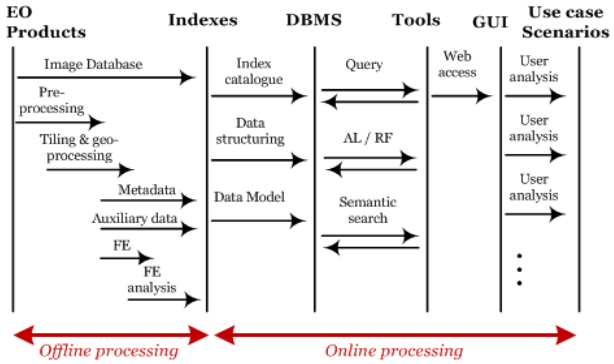

Fig 2. Functional state diagram.

The modularity of the platform sustains open data analysis, providing flexibility for the selection of input data and parameter setting. The functional state diagram matching the proposed framework is presented in Fig 2. The EO products are converted into formats suitable for ingestion and further visualization, then indexed according to the content properties. The database management system (DBSM) will manage the index catalogue and enable semantic queries and learning. The results will be conveyed to the user through web access and displayed for user analysis.

The complexity of EO data analysis will complicate the process. In order to overcome the lack of universal methods and algorithms, we proposed a modular design that can accommodate a wide range of techniques. Information mining will be enhanced, however, the algorithms for ñsmart indexingò tend to overload the system when dealing with huge volumes of data. The ability of dividing, load balancing and parallel computing will enable faster processing. Another drawback is given by the growing speed of Big Data. The results of application scenarios must be delivered as fast as new data is provided in order to fulfill the userôs scope. Thereby, we proposed a 2 step data analysis (Fig 2):

- *Data-driven (offline) processing* ï a time intensive task aiming to prepare a basic content-based indexing, according to the main image characteristics.
- *User-driven (online) processing*ï a real time task centered on the human-machine interaction that focus on learning and searching for the knowledge shared by the user.

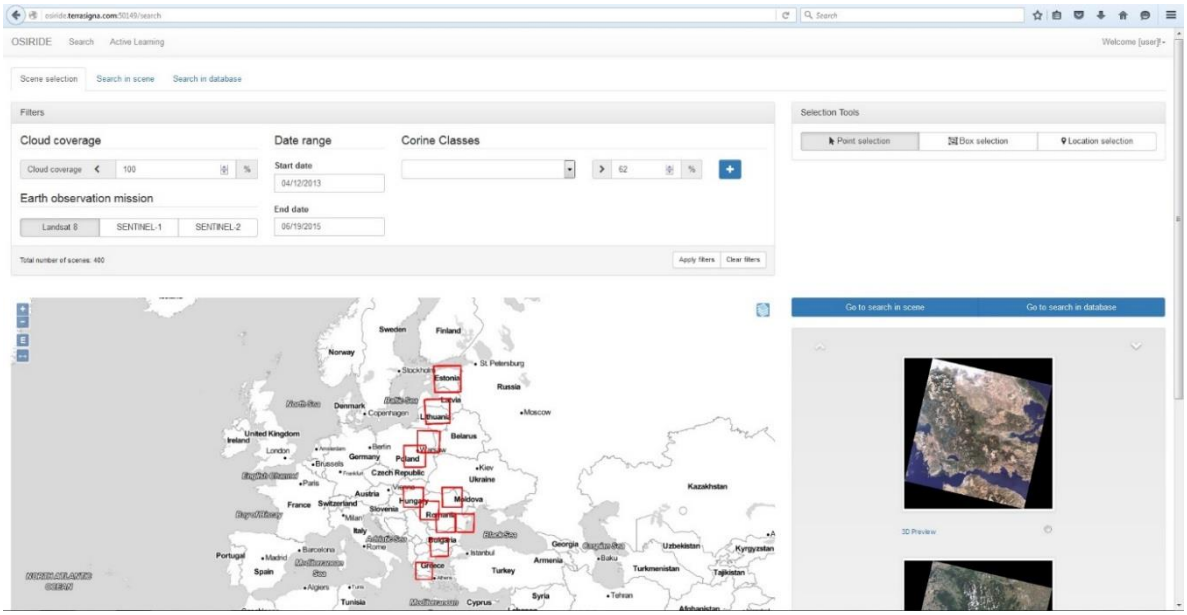

Fig. 3. OSIRIDE Graphical User Interface ï main view.

2.1. Data ingestion

When dividing the process in two parts, storing, management and information access is increased. The first part is focusing on data preparation for further semantic searches, addressing specific preprocessing issues and feature extraction. When ingested into the system, an EO scene is cut into small tiles. These tiles represent the natural semantic units of a scene, in the sense that they contain a simple structure of objects, which can be labeled. The hidden patterns inside each tile is extracted and indexed by means of feature extraction procedures. A series of image descriptors have been adapted for EO imagery and integrated into the platform: color histogram [8], Gabor descriptor [9], Weber Local Descriptor [10] and Local Binary Pattern [11]. They express color or texture, underlining a different aspect of the image. Categories of objects are discriminated using these primitive features and the descriptors act as a vocabulary for the relevant content. Proper indexes will be added to each tile based on the metadata and the descriptors.

2.2. Browsing and data analysis

Once the index catalogue is ready, the user can start exploring the database. The human machine interaction is assisted by a GUI, with very intuitive panels (Fig. 3). The information flow is mainly from the user to the computer, starting with browsing through the data in search of a collection that fits a certain sensor, time and location criteria. A preliminary insight about a scene content is offered to the user through a visualization tool, which displays the content in a 3D manner, based on the computed descriptors. A support Vector Machine (SVM) [12] will be called within a query process, handling in most cases the retrieval problem as an interactive learning issue used to split the feature space into ñrelevant classò and ñnon-relevant classò. The user will supervise the learning process through an interactive loop, with positive and negative examples, such that inferring data models and relations between the features of a certain concept will turn to a high level semantic annotation. In order to minimize the interaction, a maximum ambiguity criteria is added [13]. The system will display the patches more close to the SVM hyperplane, ranked ascending according to the Euclidean distance. This is called a relevance feedback (RF) technique that basically exposes to the user the candidates with the highest degree of uncertainty and make him decide whether that data is relevant to his query or not.

Data analysis and interpretation is mainly based on the visual inspection of the user. However, additional information will be at the userôs disposal to select according to the requirements of the application scenario.

3. PRELIMINARY RESULTS

A prototype system have been developed based on the described framework. The ingestion interface was built in Python and is available in command line with few simple parameters. The data managers can lunch batch processing sessions that extract EO data from archives, generate metadata, previews, and descriptors. The web application was developed entirely in JavaScript using node.js for server side. For accessing the data and dedicated algorithms Web Processing Services standard was used. The WPS have been implemented in order to offer a flexible way to access server

side resources like complex mechanism for searching through all the descriptors or performing active learning sessions in an asynchronous mode.

At the moment, in the database, there are almost 500 images. The data ingestion is work in progress. For the prototype demonstration the numbers of foreseen ingested images are given in the table below.

Sensor	Number of images	Type of data	Period of time
Landsat 8	1858	Optic	April 2013 ï June 2015
Sentinel 1	87	SAR	Oct 2014 ï Sept 2015
Sentinel 2	144 - tiles	Optic	Dec 2015

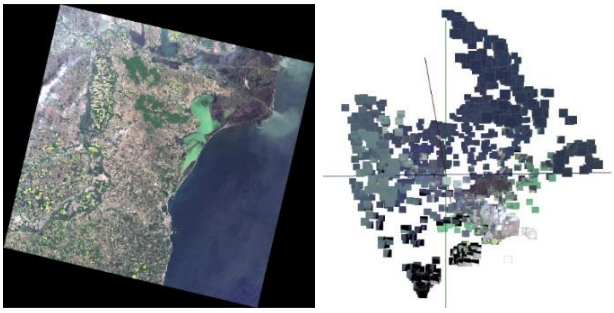

Fig. 4. Landsat 8 image including the query area: natural color preview (left), 3D preview (right).

Fig. 5. The most relevant "coastal area" tiles in the database.

In order to demonstrate the system functionality, we propose an application scenario for coastal monitoring in the past 3 years. The data selection criteria will help extracting only the data relevant for this case. We will refer thus to: Landsat 8 sensor, Cloud coverage less than 5%, Marine waters (Corine Land Class label) at least 5%. There are 31 scene meeting this criteria. We chose the image in Fig. 4 to start the semantic search. We target areas corresponding to tiles of 100x100 pixels picturing the seashore. As we expect to retrieve arrangements of water and agricultural areas/artificial surfaces, we go for a combined feature: the concatenation of the Gabor descriptor and the color histogram. The most relevant 20 tiles in the database are presented in Fig. 5. It is easy to see that they are meeting the imposed requirements.

4. CONCLUSIONS

In this paper we have introduced an integrated platform for image retrieval able to cope with the current emergence of very large image collections. Adapting CBIR techniques for the EO domain and make them available through a web platform is of main importance towards the full exploitation of the next generation of satellite images. The modular architecture facilitates the integration of several methods and algorithms to support information mining. The ability to choose from a list of multiple descriptors, combined with the SVM active learning and the RF criteria provides flexibility when searching through the database.

5. REFERENCES

[1] OSIRIDE Website: http://wiki.services.eoportal.org/tiki-index.php?page=OSIRIDE&highlight=osiride

[2] G. Sheikholeslami, W. Chang, and A. Zhang, "SemQuery: Semantic Clustering and Querying on Heterogeneous Features for Visual Data", *IEEE Tr on Know and Data Eng*, vol. 14, no. 5, 988-1002, 2002.

[3] M. Datcu, and K. Seidel, ñHuman-Centered Concepts for Exploration and Understanding of Earth Observation Imagesò, *IEEE Trans. on Geosc. and Rem. Sen.,* vol. 43, no. 3, pp.601-609, 2005.

[4] M. Datcu, K. Seidel, S. DõElia, and P.G. Marchetti, ñKnowledge driven information mining in remote-sensing image archivesò, *ESA Buletin* 110, pp. 26-33, May 2002.

[5] C. R. Shyu, et all, ñGeoIRIS: Geospatial Information Retrieval and Indexing Systemð Content Mining, Semantics Modelling, and Complex Queriesò, *IEEE Trans. on Geoscience and Remote Sensing*, vol. 45, no. 4, pp. 839-852, April 2007.

[6] D.Espinoza-Molina, et all, ñVery-high resolution SAR Images and linked open data analytics based on ontologiesò, *IEEE J. of Sel. Topics in App. Earth Obs. and Rem. Sens.*, vol. 8, no. 4, pp 1696-1708, April 2015.

[7] X. Wu, X. Zhu, G. Wu, and W. Ding, ñData mining with Big Dataò, *IEEE Tr on Kn and Data Eng*, vol.26, nr.1, pp. 97-107, 2014.

[8] L. Fuhui, H. Zhang, and D.D. Feng, "Fundamentals of content-based image retrieval", *Multimedia Inf. Retrieval and Management*. Springer Berlin Heidelberg, pp. 1-26, 2003.

[9] W. Li, K. Mao, H. Zhang, and T. Chai, ñDesigning Compact Gabor Filter Banks for Efficient Texture Feature Extractionò, *2010 11th Int. Conf. Control, Automation, Robotics and Vision Singapore*.

[10] J. Chen, S. Shan, et all, ñWLD: A robust local image descriptor,ò *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 9, pp. 1705ï 1720, 2010.

[11] M. Musci, et all, ñTexture characterization in remote sensing imagery using binary coding techniquesò, Proceedings *of the 4th GEOBIA*, pp. 437-442 , 7-9 May 2012, Brazil.

[12] M. Babaee et all, ñVisualization-Based Active Learning for the Annotation of SAR Imagesò, *IEEE J. of Sel. Topics in Appl. Earth Obs and Rem. Sen.*, vol.8, no. 10, pp. 4687 ï 4698, 2015.

[13] M. Ferecatu, and N. Boujemaa, ñInteractive Remote-Sensing Image Retrieval Using active relevance feedbackò, IEEE Trans on Geosc and Rem Sen, vol. 45, no.4, pp 818- 826, 2007.

AN ASSESSMENT OF SIMILARITY MEASURES FOR CBIR SYSTEMS

Radu Tănase [1,2], Anamaria Rădoi [1], Corina Văduva [1], Mihai Datcu [3,1]

[1] University "Politehnica" of Bucharest
[2] Military Technical Academy of Bucharest
[3] German Aerospace Center

ABSTRACT

The continuous increase in the amount of acquired satellite images for Earth Observation leads to a growing demand of building content-based image retrieval (CBIR) systems that are able to extract the desired information with a minimum effort by applying similarity measures between a query and objects to be retrieved. This paper presents an assessment and at the same time a benchmarking of some of the most widely used similarity measures that can be used for retrieving objects similar to a given query. Experiments are made on two multispectral images: a 30 meters LandSAT 8 image and a 10 meters Sentinel 2 image, both acquired over the Romanian capital city, Bucharest.

Index Terms— similarity measure, feature extraction, CBIR, WLD, color histogram, KLD, NCD, Levenshtein distance, Hamming distance

1. INTRODUCTION

With the recent advances of technology, nowadays sensors are able to acquire large amounts of high quality Earth Observation data. In this context, an efficient comparison of objects from the image archive is mandatory for any CBIR system. As a consequence, image similarity functions are at the base of this type of systems that can be divided into three main categories: query by example, active learning and relevance feedback [1].

Given this context, the choice of the similarity function used for retrieval has to be done according to the properties of the images for which the CBIR system is applied. In order to achieve even better performances, instead of applying the similarity function directly on image patches, this function can also be applied on feature descriptors. Depending on their quality, one can greatly improve the performances of a CBIR system.

In this paper we employ the Webber Local Descriptors (WLD) [2] and the color histogram (HIST) in order to describe the content of image patches and we query objects from different semantic classes using the following similarity measures: the Euclidian distance, the Manhattan distance, the Hamming distance, the Levenshtein distance, the correlation coefficient, the Kullback-Leibler Divergence and the Normalized Compression Distance.

2. IMAGE DESCRIPTORS

The Webber Local Descriptors [2] are known for their good response to image contrast, therefore they are widely used to characterize image patches with strong texture. The idea is to represent an image patch by a 2D histogram of the excitations and orientations computed for each pixel. These two components are given by:

$$\mathcal{E}(x_c) = arctg \sum_{i=0}^{p-1} \frac{x_i - x_c}{x_c} \tag{1}$$

$$\mathcal{O}(x_c) = arctg\left(\frac{v_s^{11}}{v_s^{10}}\right) \tag{2}$$

where x_c represents a central pixel, x_i represents a surrounding pixel and v_s^{11} and v_s^{10} are the outputs of two local gradient filters.

The color histogram is perhaps one of the most widely used descriptor in image analysis. Simple and efficient, it uses the distribution of the color to describe the patch content, therefore it is sensitive to color changes. In our case, the feature descriptor is composed by concatenating the histograms of all bands of multispectral images.

3. SIMILARITY MEASURES

The choice of a similarity measure has a great influence on the retrieval accuracy of the CBIR systems. In this paper, we evaluate 7 similarity functions applied on the descriptors mentioned in the previous section. Considering two descriptors $\boldsymbol{X}$ and $\boldsymbol{Y}$ of N elements, i.e., $\boldsymbol{X} = [x_1, x_2, \dots, x_N]$ and $\boldsymbol{Y} = [y_1, y_2, \dots, y_N]$, we have implemented a set of different metrics to calculate the distances between two image patches:

a) The Euclidian Distance (EUC):

$$EUC = \sqrt{\sum_{i=1}^{N} (x_i - y_i)^2} \tag{3}$$

b) The Manhattan Distance (MAN):

$$MAN = \sum_{i=1}^{N} |x_i - y_i| \tag{4}$$

where $|a|$ denotes the absolute value of a.

c) The Hamming Distance (HAM) is an information theoretic measure that can be used to compare binary descriptors (i.e., descriptors containing 0's and 1's):

$$HAM = \sum_{i=1}^{N} (x_i \oplus y_i) \tag{5}$$

where $\oplus$ represents the modulo 2 addition [3]. For other types of descriptors, HAM counts the number of elements that differ between the query descriptor and the descriptor of the patch for which the similarity measure is computed.

d) The Levenshtein Distance (LEV) is a text-oriented measure of similarity between two strings that counts the number of deletions, insertions, or substitutions required to transform one string into the other [4]. Mathematically, this distance is defined by :

$$LEV(x,y) = \begin{cases} max(x,y), & if\ min(x,y) = 0 \\ min\begin{pmatrix} LEV(x-1,y)+1 \\ LEV(x,y-1)+1 \\ LEV(x-1,y-1)+1_{\{x_i \neq y_j\}} \end{pmatrix}, & otherwise \end{cases} \tag{6}$$

where $1_{\{a\}}$ is the indicator function (it equals 0 if a is *false*).

e) The correlation coefficient (COR):

$$COR = \frac{1}{N} \frac{\sum_{i=1}^{N} (x_i - m_X)(y_i - m_Y)}{\sigma_X \sigma_Y} \tag{7}$$

where m_X and m_Y are the means of the two descriptors and σ_X and σ_Y are their corresponding standard deviations.

f) The Kullback-Leibler Divergence (KLD) compares two probability densities $p(u)$ and $q(u)$ following [5]:

$$KLD(p,q) = \sum_{u} p(u)\, log\left(\frac{p(u)}{q(u)}\right) \tag{8}$$

where p is the distribution of $\boldsymbol{X}$, and q is the distribution of $\boldsymbol{Y}$. KLD is equal to 0 when $p(u) = q(u)$, namely the patch-based descriptors are characterized by identical distributions.

g) The Normalized Compression Distance (NCD) uses the compressed versions of two objects (both separately and combined) to compute the degree of similarity between them. NCD is given by [6]:

$$NCD(X,Y) = \frac{C(X,Y) - min\{C(X), C(Y)\}}{max\{C(X), C(Y)\}} \tag{9}$$

where $C(X)$ and $C(Y)$ represent the compressed versions of X and Y, respectively, while $C(X,Y)$ represents the compressed version of the objects concatenated. NCD has been already employed in classification, clustering [7] and anomaly detection, and, recently, in a CBIR system. [8]

4. EXPERIMENTAL RESULTS AND DISCUSSIONS

In order to make an assessment of the state-of-the-art similarity measures, a 30 m resolution LandSAT8 multispectral image (8 bands) and a 10 m resolution Sentinel2 multispectral image (4 bands) over Bucharest, Romania, were used. Three main classes have been considered in the scene: urban areas, vegetation areas, and agricultural areas (Fig. 1). These classes were inspired from the 300 m resolution land cover maps provided by the European Space Agency (2008-2012 time period).

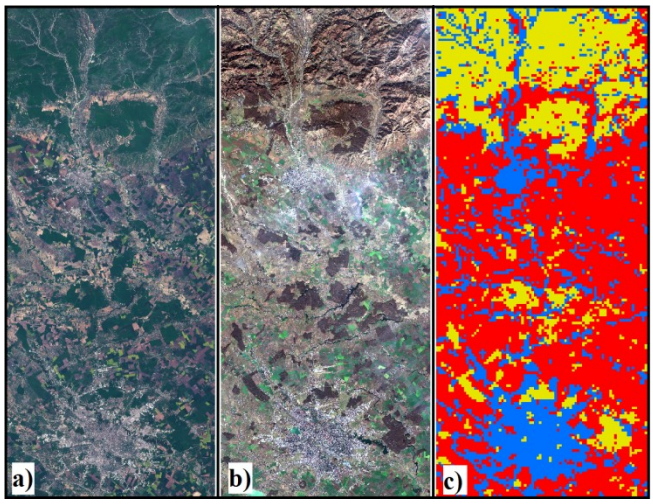

Fig.1. The testing images, together with the ground truth.
a) LandSAT8; b) Sentinel2; c) Ground Truth

All experiments were made on each image separately. Firstly, each image was split into 50x50 pixels patches and the corresponding feature descriptors for each patch were computed. Then, ten representative patches for each class were selected and their centroids were computed. These centroids were used as sliding windows in order to detect similar patches in the images. The detection results (true-positive rates) are shown in Fig. 2 and Fig. 3.

doi: 10.2788/854791

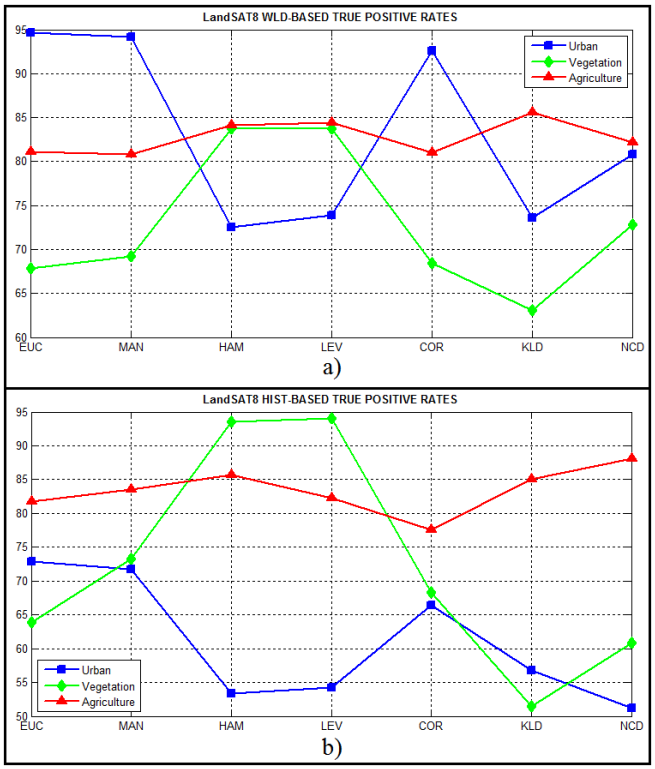

Fig. 2. Detection results on LandSAT8 image.
a) WLD feature descriptors; b) HIST feature descriptors

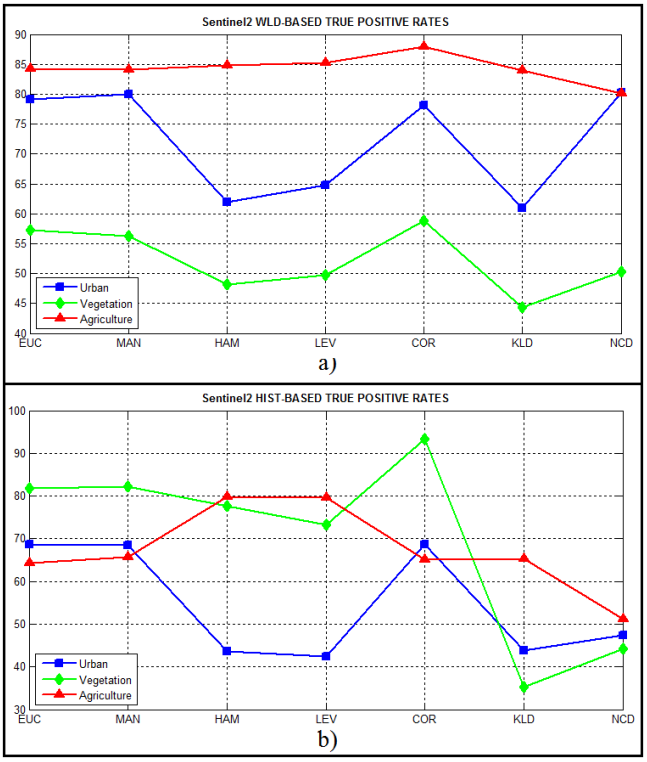

Fig. 3. Detection results on Sentinel2 image.
a) WLD feature descriptors; b) HIST feature descriptors

Each class is best detected when specific combinations of feature descriptors and measures of similarity are employed. At first sight, one can notice that the "*urban*" and "*agriculture*" patches are better detected when using WLD descriptors, while "*vegetation*" patches are better detected when employing HIST descriptors. This happens because WLD descriptors are more sensitive to the heterogeneous scenes (like urban scenes and agricultural parcels) while the HIST descriptors are more sensitive to the relatively homogeneous areas (like forested areas). The true positive rates obtained for "*vegetation*" patches are also influenced by the fact that the Sentinel2 image was acquired during winter. In these conditions, it is more difficult to detect the forests poor canopy.

If we analyze the true positive rates of detection based on the WLD descriptors, we can notice that urban scenes are very efficiently detected when employing one of the EUC, MAN, COR or NCD similarity measures, while the agricultural parcels are very efficiently detected when employing one of the HAM, LEV, COR or KLD distances. However, especially in real-time CBIR systems, the processing time is very important when choosing the similarity measure to employ. This is why it is recommended to avoid choosing the LEV, KLD or NCD, since they are very time-consuming (Fig. 4).

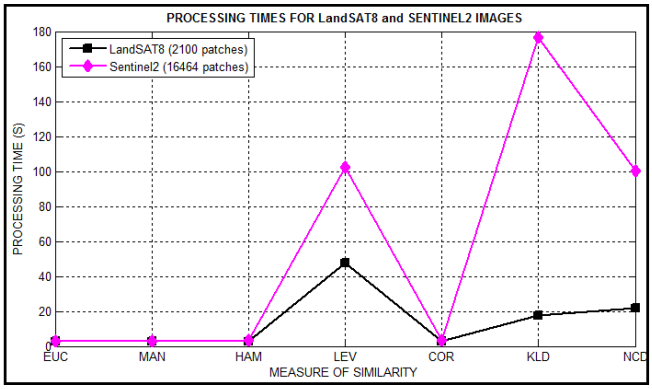

Fig. 4. Processing times for LandSAT8 and Sentinel2 images.

Further, if we analyze the results based on the HIST descriptors, we can notice that "*vegetation*" patches are the most correctly detected when employing HAM, LEV or COR distances. It is however recommended to avoid the use of the LEV distance, since it is time-consuming, and to use the HAM distance instead. The results are almost identically, since the two similarity measures are related (HAM is a particular case of LEV).

Taking into consideration the above statements, examples of good detection results on both images, for each class, are shown in Fig. 5 and Fig. 6, respectively. All detected patches are shown in white.

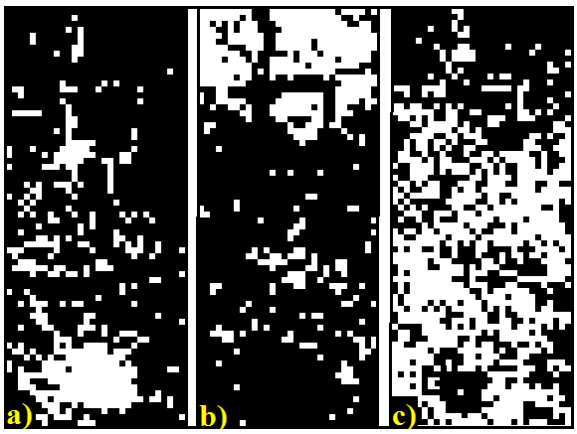

Fig. 5. Detection results on LandSAT8 image. a) Urban-WLD-EUC; b) Vegetation-HIST-HAM; c) Agriculture-WLD-COR.

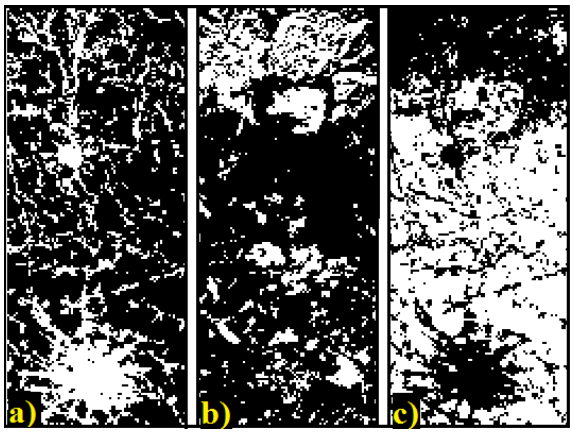

Fig. 6. Detection results on Sentinel2 image. a) Urban-WLD-MAN; b) Vegetation-HIST-COR; c) Agriculture-WLD-COR.

Examples of detected patches on both images, considering the same parameters in Fig. 5 and Fig. 6, are shown in Fig. 7.

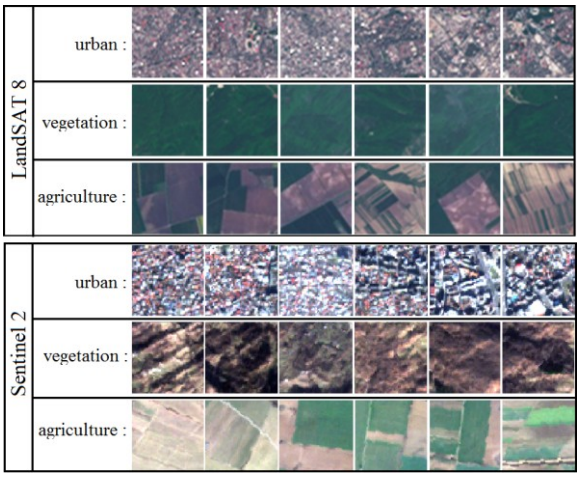

Fig. 7. Examples of detected patches for the three classes, on both LandSAT8 and Sentinel 2 images.

Another observation that can be made if we analyze the above results is that, as a compromise, the COR similarity measure returns fairly good results regardless of the employed feature descriptor or the query;

5. CONCLUSIONS

Regardless of whether image patches or feature descriptors are employed, the choice of a suitable measure of similarity is critical for the performances of a CBIR system. A comparative survey and at the same time a relevant benchmarking of the current state-of-the-art similarity functions were conducted in this paper. It was found that, depending on the type of information that each feature descriptor encapsulates, a considered measure of similarity can lead to different results. Thus, a detailed analysis of the pair query - feature descriptor – similarity measure is mandatory for a CBIR system.

7. REFERENCES

[1] K. Alonso, M. Datcu, "Accelerated Probabilistic Learning Concept for Mining Heterogeneous Earth Observation Images", *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol.8, no.7, pp.3356-3371, July 2015

[2] J. Chen, S. Shan, C. He, G. Zhao, M. Pietikainen, X. Chen, and W. Gao, "WLD: A robust local image descriptor," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 32, no. 9, pp. 1705–1720, Sep. 2010

[3] R.W. Hamming, "Error detecting and error correcting codes", *Bell System Technical Journal*, vol. 29, pp. 147-160, April 1950

[4] V.I. Levenshtein, "Binary codes capable of correcting deletions, insertions, and reversals", *Soviet Physics Doklady*, vol. 10, no. 8, pp. 707-710, 1966

[5] S. Kullback and R.A. Leibler, "On information and sufficiency", *Annals of Mathematical Statistics*, vol. 22, pp.79-86, 1951

[6] M. Li, X. Chen, X. Li, B. Ma, P.M.B. Vitányi, "The similarity metric", *IEEE Transactions on Information Theory*, vol. 50, no. 12, pp. 3250-3264, 2004

[7] R. Cilibrasi and P.M.B. Vitányi, "Clustering by Compression", *IEEE Transactions on Information Theory*, Issue 51, No.4, 2005.

[8] D. Cerra and M. Datcu, "A fast compression-based similarity measure with applications to content-based image retrieval", *Journal of Visual Communication, Image Representation*, Vol. 23, Issue 2, pp. 293-302, 2012

PATCH-BASED MULTISPECTRAL IMAGE CLASSIFICATION ASSESSMENT FOR SENTINEL-2 DATA ANALYSIS

Florin-Andrei Georgescu[1, 2], *Mihai Datcu*[2, 3], *Dan Răducanu*[1]

[1]Military Technical Academy (MTA), 050141 Bucharest, Romania
[2]Politehnica University of Bucharest (PUB), 060042 Bucharest, Romania
[3] Remote Sensing Technology Institute, German Aerospace Center (DLR), 82234 Wessling, Germany

ABSTRACT

In the last decades we can observe a continuous growth of the remote sensing data, acquired with a large variety of sensors, in different acquisition modes. This has lead to large collections of Earth Observation (EO) image data that must be understood and analyzed. Even though content based image retrieval systems designed to handle EO data exists, no general approaches are available. This paper present a straight forward application that uses feature extraction methods such as Gabor (G) filter banks, Weber Local Descriptor (WLD), Spectral Histogram (H) and Bag of Words (BoW) in the context of patch based EO image classification. The goal is to provide an assessment for Sentinel-2 data analysis. Another important objective of our system is to determine the best patch size for classification and the best number of classes that can be observed inside the analyzed remote sensing scene.

Index Terms— EO image classification, Gabor filter banks, Histogram, Bag of Words, Weber Local Descriptor, Sentinel-2

1. INTRODUCTION

The need of searching through large amounts of EO image data in the minimum amount of time, have led to the development of EO content based image retrieval and indexing systems which can increase the productivity and can be of vital importance in the case of critical infrastructures monitoring.

For a better understanding and efficient retrieval of EO image data from large database collections have been developed powerful tools such as GeoIRIS [1], a content-based multimodal Geospatial Information Retrieval and Indexing System, which includes automatic feature extraction, visual content mining from large-scale image databases, and high-dimensional database indexing for fast retrieval [1], KIM - knowledge-driven information mining, which is based on human-centered concepts (HCCs) and implements new features and functions allowing improved feature extraction, search on a semantic level, the availability of collected knowledge and interactive knowledge discovery [2], SemQuery which is a semantics-based clustering and indexing approach, used to support visual queries on heterogeneous features of images [3] and others make use of Interactive learning [4] and even of Visual Grammar [5].

Our goal is to compute an assessment of patch based multispectral classification methods that can be integrated into a system designed for EO data analysis and information mining. This could have a great impact in land cover applications such as agriculture, vegetation, urbanism, hydrology or hazard monitoring.

The application we propose involves a step of feature extraction and another one of feature classification. In the first step, feature descriptors such as Gabor [6], WLD [7], Spectral Histogram and BOW [8] are computed and in the second step, the obtained features are given as inputs to the classification algorithms such as k-Nearest Neighbors (k-NN) or Support Vector Machines (SVM). The goal of this classification is to accurately obtain coverage maps of the input scene. Also, in our assessment we analyze the classified data accuracy and precision, determine the best patch size and the optimum number of classes that can be identified in the analyzed image.

2. PATCH-BASED CLASSIFICATION OF EARTH OBSERVATION IMAGES

As a consequence of the continuous growth of EO data collections, the development of remote sensing data analysis tools and methods is required. In most of the cases, when dealing with large amounts of EO image data, the search for a specific image or phenomenon is a well known problem.

The authors are using multispectral image data in order to establish the conditions in which the proposed feature extraction methods perform at their best in a patch-based classification scenario. In Figure 1, we present the main workflow used for image classification.

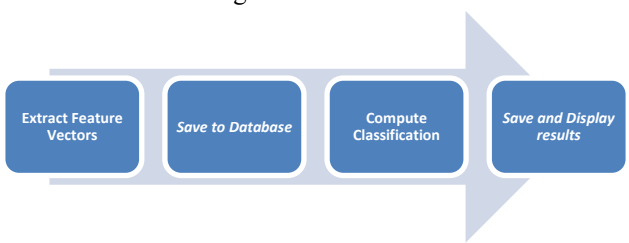

Figure 1. Workflow for Feature Extraction and Earth Observation image classification

2.1. Feature Extraction

It is known that the best results are obtained with specific algorithms for specific types of data. Also, depending on the porpoise of the classification, we should choose between feature extraction methods that are more sensitive to texture, color or both texture and color.

Standard image descriptors such as Gabor, WLD, Spectral Histogram and BoW are computed at patch level. For Gabor features [6] extraction method, which is known to be well suited for texture detection and classification, we used $\theta = 6$ orientations and $\varphi = 4$ scales. According to [7], WLD computation is done using $C = 18$ differential excitation histogram bins and $T = 8$ bins for orientation histogram. In the case of spectral histogram features, we computed $H = 64$ bins for each spectral band [9] and in the BoW [8] model approach the codebook is generated from the spectral values of each pixel, using in our computation 64 generated spectral words.

2.2. Earth Observation Image Classification

In the frame of EO image classification step, we used supervised methods like Support Vector Machines (SVM) and k-Nearest Neighbors (k-NN). Regarding the parameter setup, for the SVM classification we used a radial basis function (RBF) kernel that have the gamma coefficient $\gamma = 3.0518 \times 10^{-4}$ and the regression parameter $C = 5$, and for the k-NN classification case we used $k = 10$ neighbors.

3. PERFORMANCE ANALYSIS

Even though different patch sizes can be used, some of patch-based feature extraction methods are imposing a minimum patch size that can be used. Also, when choosing the patch size is important to take into account the spatial resolution of the analyzed image in order to cover the objects of interest. Therefore, for observation of Sentinel-2 images using the 10m spatial resolution spectral bands, a patch of 25x25 pixels will cover relevant objects of different thematic classes such as agriculture, urban areas, water bodies, etc.

To establish the optimum number of classes that can be obtained using unsupervised classification, we used rate distortion (RD) theory, as it is presented in [10]. In order to compute the RD algorithm we have taken into account the k-Means clustering of the entire scene (Figure 3) using the extracted BoW features. The reason behind choosing these features is lying in the property of the BoW model of preserving both the texture and color distribution in the analyzed patches. As it can be seen in Figure 2, the optimum number of classes to be analyzed for a patch size of 25x25 pixels is between 5 and 30 thematic classes.

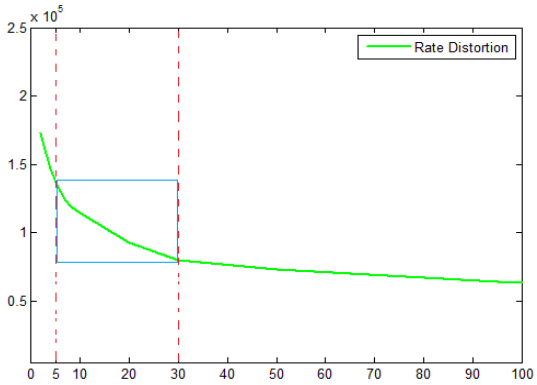

Figure 2. Rate Distortion. X axis represents the number of analyzed classes and Y axis represents the associated MSE

The results presented in this paper are computed for 5 generic classes such as agriculture fields with abundant vegetation (Agriculture-HV – C1), agriculture fields with less vegetation (Agriculture-LV – C2), forests (C3), urban areas (C4) and water bodies (C5), as presented in Figure 4.

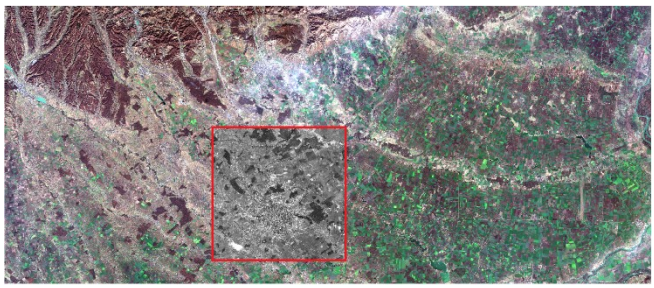

Figure 3. Entire Sentinel-2 analyzed scene.
The red rectangle represent the region used for quality assessment.

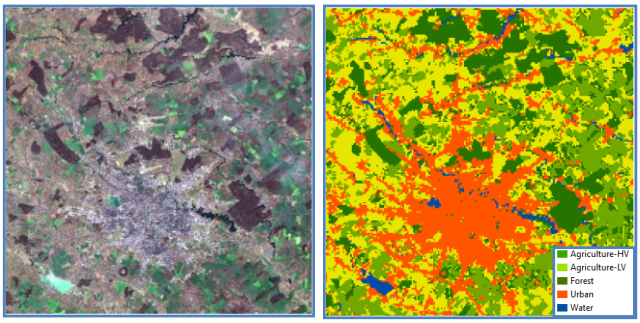

Figure 4. Sentinel-2 scene used for quality assessment (Left).
Manual Annotation into 5 thematic classes (Right)

For quantitative and qualitative assessment we make use of a small region, representing almost 10% of the original scene, which contains all the analyzed classes of the initial scene. As it is stated in [11], our computation assumes a sample image for which we have ground truth knowledge and a set of well known classifiers like SVM and k-NN.

Using a sample image covering a surface of 2500 km^2, illustrated in Figure 4, we realized the manual annotation of

doi: 10.2788/854791

40000 patches with a size of 25x25 pixels into 5 thematic classes, taking into account all the spectral bands in order to assign to each patch the best suited class. This annotation is used further for automatic generation of the training sets and also to compute the confusion matrixes needed in the qualitative and quantitative evaluation of the classifications.

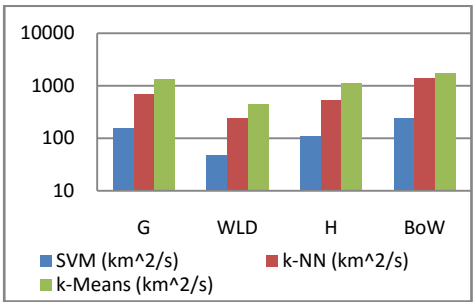

Figure 5. Classification speed expressed in km^2/s

To demonstrate the classification performances, we computed the classification speed expressed in km^2/s for the full analyzed scene of 25725 km^2. As it can be seen in Figure 5, the top speed of classification is achieved in the case of k-Means clustering into 5 thematic classes, followed by k-NN and SVM classification. The hardware used for this benchmark has 16 GB of RAM memory and a 2.4 GHz 8 core CPU.

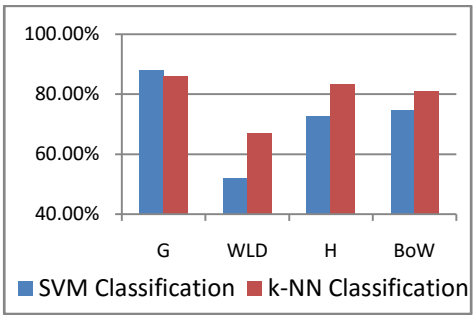

Figure 6. Mean Precision of SVM and k-NN

The mean accuracy of the evaluated supervised classification algorithms is presented in Figure 6. These values are computed using 10 random training sets obtained from the manual annotation. The training sets are 20% of the sample image and 1.9% of the full scene being analyzed. As it can be observed, the notable results are obtained for Gabor, BoW and for Spectral Histogram feature extraction methods, while for WLD the accuracies obtained are between 50% and 65%.

4. RESULTS

Even though the assessment of the methods is done mostly for a small region, the goal of this paper is to show the performances of the classification of feature extraction methods on a large Sentinel-2 image. The analyzed image represents a Sentinel-2 scene covering an area of 245x105km^2 over Romania. For our analysis, we used four spectral bands of 10m spatial resolution (Blue, Green, Red and Infrared) available in the Sentinel-2 image to classify 416600 patches of 25 x 25 pixels into 5 thematic classes.

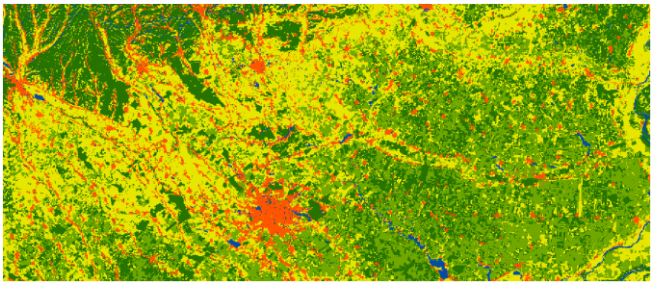

Figure 7. SVM Classification – Gabor Features

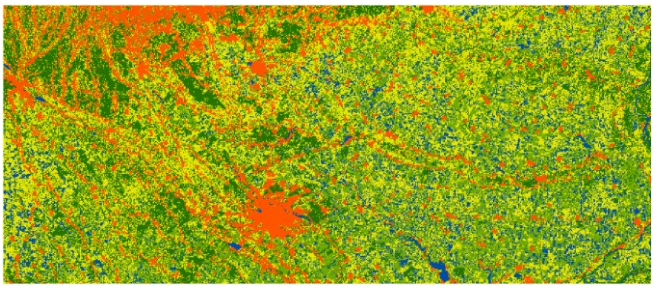

Figure 8. SVM Classification – WLD Features

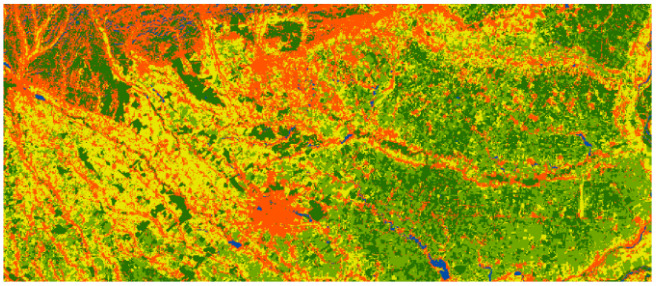

Figure 9. k-NN Classification – Spectral Histogram Features

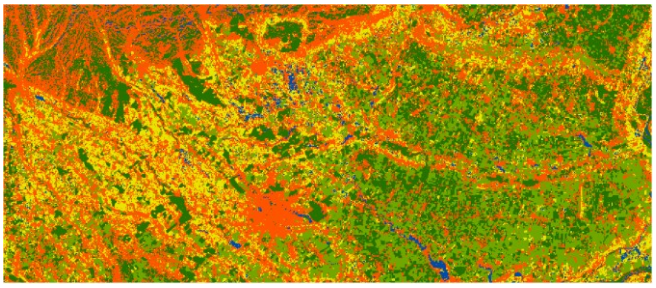

Figure 10. k-NN Classification – BoW Features

As presented in Section 3, we computed our results for a very small area in order to generalize the results for the full scene as is suggested in [11].

The qualitative and quantitative assessment of the entire scene can be approximated from the Precision - Recall tables computed for supervised SVM (Table 1) and k-NN (Table 2) classification methods. As it can be observed from

the Tables 1-2, and with respect to Figure 6, Gabor features performances are the best, followed by Spectral Histogram features, BoW and WLD for both classification methods.

		C1	C2	C3	C4	C5
G	P	**92.5%**	**83.0%**	**89.8%**	**97.0%**	**77.6%**
	R	87.3%	92.0%	81.9%	96.6%	84.3%
WLD	P	**47.1%**	**36.3%**	**43.1%**	**90.4%**	**42.3%**
	R	49.0%	58.4%	26.9%	90.8%	62.1%
H	P	**88.6%**	**60.9%**	**59.4%**	**76.9%**	**77.7%**
	R	77.8%	64.3%	62.3%	77.1%	89.4%
BoW	P	**90.2%**	**67.0%**	**56.6%**	**80.5%**	**77.9%**
	R	82.0%	68.3%	77.0%	77.5%	92.1%

Table 1. SVM Precision-Recall rates for a small image region (P=precision, R= recall)

		C1	C2	C3	C4	C5
G	P	**91.0%**	**94.9%**	**83.9%**	**92.4%**	**66.5%**
	R	94.9%	85.7%	89.0%	97.7%	95.8%
WLD	P	**63.6%**	**67.5%**	**70.7%**	**90.4%**	**42.3%**
	R	79.8%	71.5%	55.2%	94.6%	55.1%
H	P	**88.4%**	**84.1%**	**80.0%**	**88.5%**	**76.2%**
	R	93.8%	82.9%	84.9%	82.1%	91.4%
BoW	P	**82.9%**	**74.8%**	**76.7%**	**93.1%**	**77.2%**
	R	95.2%	86.9%	88.0%	71.3%	91.9%

Table 2. k-NN Precision-Recall rates for a small image region (P=precision, R= recall)

In Figure 7 we present the best classification result obtained with SVM Classification for Gabor features. Figure 8, 9 and 10 are illustrating other classifications on the full scene.

5. CONCLUSIONS

This paper demonstrates the high potential of using Sentinel-2 data for generation of land cover maps with possible applicability in agriculture, vegetation, urbanism, hydrology and hazard monitoring. In order to do so, we computed feature extraction methods in a patch-based approach for the 10m spatial resolution spectral bands of Sentinel-2, and then we applied supervised classification algorithms to generate the thematic maps. Also unsupervised k-Means clustering is used in order to determine the suitable number of classes that can be extracted from the analyzed scene.

Even though the computation speed is dependent on the hardware performances, the classification time is relevant for an approximate evaluation of the time consumption in similar hardware architectures.

The results we obtained show a good separation into different classes of Sentinel-2 recently released data into a few generic land cover classes such as agriculture, forest, urban and water.

6. REFERENCES

[1] Shyu, C., M. Klaric, G. Scott, A. Barb, C. Davis and K. Palaniappan - *GeoIRIS: Geospatial Information Retrieval and Indexing System - Content Mining, Semantics Modeling, and Complex Queries*, IEEE TGRS, Vol. 45, No. 4, 2007

[2] Datcu, M. and K. Seidel - *Human-Centered Concepts for Exploration and Understanding of Earth Observation Images*, IEEE TGRS, Vol. 43, No. 3, 2005

[3] Sheikholeslami, G., W. Chang and A. Zhang, *SemQuery: Semantic Clustering and Querying on Heterogeneous Features for Visual Data*, IEEE TKDE, Vol. 14, No. 5, 2002

[4] Schröder M., H. Rehrauer, K. Seidel and M. Datcu, *Interactive Learning and Probabilistic Retrieval in Remote Sensing Image Archives*, IEEE TGRS, Vol. 38, No. 5, 2000

[5] Aksoy S., K. Koperski, C. Tusk, G. Marchisio and J. Tilton, *Learning Bayesian Classifiers for Scene Classification with a Visual Grammar*, IEEE TGRS, Vol. 43, 2005

[6] Georgescu F.-A., C. Văduva, M. Datcu and D. Răducanu – *A framework for benchmarking of feature extraction methods in Earth Observation image analysis,* Image Information Mining 9th Edition, Bucharest, 2014.

[7] Chen, J., S. Shan, C. He, G. Zhao, M. Pietikäinen, X. Chen and W. Gao, *WLD: A Robust Local Image Descriptor*, IEEE TBC, Vol. 32 , No. 9, 2010

[8] Cui S., G. Schwarz, M. Datcu, *Remote Sensing Image Classification: No Features, No Clustering*, IEEE Journal, Nov 2015

[9] Tuncel E., P. Koulgi, K. Rose, *Rate-Distortion Approach to Databases: Storage and Content-Based Retrieval,* IEEE Transactions on Information Theory, Vol. 50, No. 6, June 2004

[10] Faur, D., I. Gavat, and M. Datcu, *Salient Remote Sensing Image Segmentation Based on Rate-Distortion Measure*, IEEE TGRS, Vol. 45, No. 4, 2007

[11] Baraldi A, L. Bruzzone, P. Blonda - *Quality Assessment of Classification and Cluster Maps Without Ground Truth Knowledge*, IEEE Transactions on Geoscience and Remote Sensing, Vol. 43, No. 4, April 2005

DATA STREAM PROCESSING FOR SITUATIONAL AWARENESS

Roger King, Osama Abuomar, and Nicolas Younan

Department of Electrical and Computer Engineering
Center for Advanced Vehicular Systems
Mississippi State University

ABSTRACT

Data stream mining refers to the extraction of different structures, patterns and models from continuous data streams from, for example, an Earth orbiting satellite. There are different storage, computational, mining, and querying challenges when analyzing data streams. Mining time-series data is considered one of the key milestones in data stream mining where models and patterns can be extracted from continuous and rapid time-series data. This paper outlines some techniques that can be used to mine big datasets from space including two phase techniques, Hoeffding bound-based techniques, SAX based techniques, granularity-based techniques, and data stream classification. Therefore, these techniques can be utilized by the Earth remote sensing community for mining time series data in real time.

Index Terms— data stream mining, situational awareness, multitemporal

1. INTRODUCTION

The Earth is facing unprecedented climatic and environmental changes, which require global scale monitoring and a host of new Earth Observation (EO) satellite missions. Addressing the need for geospatial intelligence for the environment and security requires data from EO satellites, operating as single missions or constellations, which will enable a higher revisit frequency. Sensors on board these satellites will represent a host of new measurable phenomena that will increase the variety, amount, and resolution (spatial, spectral, and temporal) of data. This increase in EO missions is expected to result in unprecedented amounts of data to be processed.

At the same time, the need for timely delivery of focused information for decision making is increasing. Therefore, EO data, which is not immediately usable, requires chained transformations before becoming the needed "information products" (easily understandable and ready-to-use without further manipulations) offered as on-demand or systematic services. Different entities may perform these transformations using their own processes, which require specific knowledge, experience, and possibly data or information from domains other than EO. Today, these information products are primarily developed in a semi-automated fashion by experts or specialized companies operating in specific application domains.

By using some approaches to design systems that can learn and/or apply knowledge, there should be increased efficiencies from the reduction of the information extraction time through the real-time automation of such processes. One such approach is data stream mining. Traditional data mining techniques are suitable for simple and structured datasets like relational databases, transactional databases and data warehouses. Data Stream mining refers to informational structure extraction as models and patterns from continuous data streams.

Data streams differ from the traditional data mining methods on several ways [1]:

- The data elements in the stream arrive online.
- The system has no control over the order in which data elements arrive to be processed, either within a data stream or across data streams.
- Data streams are potentially unbounded in size.
- Once an element from a data stream has been processed, it is discarded or archived. That is, it cannot be retrieved easily unless it is explicitly stored in memory, which typically is small relative to the size of the data streams.

2. BACKGROUND

There are three main constraints of the knowledge discovery systems: memory, time and sample size. Sample size has been the main limitation in traditional machine learning applications. That is, despite the availability of the computational software and hardware resources required for a massive search, over-fitting or "data dredging" problems will be likely expected if such a search is conducted over small sample size (i.e. less than 10,000 samples). However, in most of the data mining applications being analyzed nowadays, the limitations are memory and time instead of sample size [2].

Structured and simple datasets can be analyzed using traditional data mining techniques. Examples of such datasets are transactional databases, relational databases, and data warehouses. As a result of continuous and fast development of data collection techniques and advanced database systems, data has rapidly grown in various complex forms such as spatial and temporal datasets, nonstructured and semistructured datasets, and hyper-text and multimedia datasets. Therefore, it is now becoming necessary to mine and analyze such complex datasets [2, 3]. There has been a trend recently to design data algorithms in order to analyze continuous data streams that arrive at very high rates (i.e. algorithms that respond to user queries by scanning the relevant data samples one time only and in specific order controlled by the data streams rate) [4]. Some of data streams applications are; sensor networks (i.e. satellites), economic applications, web applications, telecommunication data management applications, network monitoring, and security applications.

Data stream mining refers to the extraction of different structures, patterns and models from continuous data streams from an Earth orbiting satellite. There are different storage, computational, mining, and querying challenges when analyzing data streams. Due to its continuous and rapid pace, it has become very important to design new techniques and algorithms that deal with different data streams. Traditional data mining and machine learning techniques require the whole data to be stored first and then analyzed after several passes over the data (i.e., from an archive).

Based on the above, when designing a data stream mining algorithm, there are two main challenges; designing rapid data streams mining methods and dealing with concept drift due to the data streams' highly dynamic nature [5,6].

3. DATA STREAM MINING ALGORITHMS

In this section, four classes of data stream mining techniques are briefly discussed. The first three classes provide techniques to develop different algorithms, but the last class is a generic controller that could be applied along with any data stream mining algorithm.

3.1. Two Phase Techniques

The two phase techniques are introduced in [7]. These techniques primarily utilize microclusters to maintain an online summary of data. Over the last 10 years many algorithms for clustering data streams have been proposed. Most data stream clustering algorithms deal with the problems of unbounded stream size, and the requirements for real-time processing in a single pass by using a two-stage online/offline approach introduced in [8].

3.2. Hoeffding Bound-Based Techniques

A very fast machine learning (VFML) was proposed as a strategy to scale up different learning algorithms [2]. Based on this strategy, the accuracy in every step of the algorithm has been determined using the number of data samples (data records). A Hoeffding bound is used as an accuracy bound for the VFML techniques [9]. According to the Hoeffding bound criteria, the true mean value, r, is at least $\left(\left(\hat{r}\right)-\varepsilon\right)$ with probability 1 - δ, where $\hat{r}$ is the mean value. That is,

$$\varepsilon = \sqrt{\frac{R^2 \ln(1/\delta)}{2n}}, \qquad (1)$$

where R is the approximate range and n is the number of data samples. This idea is utilized in some algorithms and techniques. This includes VFKM as a traditional extension to k-means clustering algorithm, very fast decision trees (VFDT), and classification techniques used in decision trees. The concept drift problem in data streams has been addressed as an extension of VFDT [10] called CVFDT. The most recent classifier is produced by running VFDT using fixed-length sliding windows.

3.3. SAX Based Techniques

In addition to its usage in data mining algorithms and techniques like classification, indexing, and clustering, a Symbolic ApproXimation (SAX) as a new representation for time series is proposed in [11].. In SAX, there are three main steps for converting a numerical time series to a symbolic time series. Piecewise Aggregate Approximation (PAA) is performed as the first step in which a size n time series is converted to an arbitrary size w time series using Equation (2):

$$C_i = \frac{w}{n} \sum_{j=\frac{n}{w}(i-1)+1}^{\frac{n}{w}i} C_j, \qquad (2)$$

where C_i is the approximated time series ith time point.

Symbolic discretization is the second step where equal areas under the Gaussian distribution's curve are produced and then respective breakpoints are set. The third step involves computing the accumulated distance measure between time series subsequences through computing this measure between each two characters in the lookup table.

3.4. Granularity-Based Techniques

The Granularity-based approach is introduced in [12]. The main motivation behind this approach is the fact that sensor nodes, smart phones, and other resource constrained devices are not compatible with data stream techniques. In essence, the granularity-based approach balances between the resources available and the resource consumption patterns of the data mining techniques.

The computational resources monitoring is the main requirement of algorithm granularity. This is performed over a predetermined time, frames/intervals, typically referred to as *TF*. The resulting overhead of the adaptation will be lower as *TF* increases. However, the consumption of resources might be high in case of the long time frame. Some data stream mining techniques and algorithms have been proposed after the Granularity-based approach. A recent tutorial that lists some of these algorithms and techniques was introduced in [13].

3.5. Data Stream Classification

Maintaining decision trees in streaming environments was considered the main motivation to develop the VFDT system [14] which is, once again, a learning system that uses Hoeffding trees. Using the data mining results of the implemented data model, general matrix multiply (GEMM) and FOCUS are two algorithms that have been developed in order to maintain the model as well as to detect any changes between every two datasets [15]. The frequent itemset model and decision tree models are some examples where these algorithms have been applied. However, other than using them for one-pass data stream mining, decision trees cannot be used to develop an on demand classifier that can be used with data streams.

4. DATA STREAM MINING APPLICATIONS

The techniques and the algorithms of data stream mining have been applied in various applications. In this section, sensor data streams mining [16], mining of mobile data streams [16], and data stream mining of electric power systems [17] are among these applications and are discussed in this section.

4.1. Sensor Data Stream Mining

Research in data stream mining has been fueled by the unprecedented growth of computational power in wireless sensor nodes. These wireless sensor networks present some challenges [16].

- Data duplication in wireless sensor networks that are densely deployed.
- If local models are generated by individual wireless sensors, then multilevel data stream mining will be necessary.
- Since the streaming data of wireless sensors are mostly noisy, then data cleaning in real time becomes important.
- Because of the limitation in resources, every wireless sensor is expected to adapt to this environment. The granularity-based approach (see Section 3.4) tends to be operated successfully in wireless sensor networks in the streaming environment.

4.2. Power Systems Data Stream Mining

Power systems is considered an excellent use of data stream mining for situational awareness in a real world application. Phasor Measurement Units (PMUs) use satellite-based GPS signals to synchronize measurements across the electric grid. This has facilitated power systems' wide area and real time monitoring. The typical data streaming rate of PMUs parameters is 30 samples per second whereas it is 1 sample per 2-5 seconds in a conventional Supervisory Control And Data Acquisition (SCADA) system. There has been a notion to increase the reliability of electrical grids due to the enhanced availability of time stamped power systems' data. Moreover, extracting useful information and patterns from this massive amount of data has become a challenge nowadays. Several data mining algorithms are now being used to extract information from synchrophasor data for improving situational awareness of power systems [17]. The extracted information can be used for event detection, for reducing the dimension of data without losing information, and also to use it as a heuristic to process future measurements [18].

Data stream mining in power systems can be classified into two parts: stream mining and dimension reduction. Stream mining algorithms utilize some state-of-the-art mining algorithms such as Hoeffding Trees (HT). HT algorithm scans the incoming data stream only once and then builds a decision tree by. The tree itself holds sufficient statistics in its leaves in order to make classification decisions of incoming data streams [17].

5. CONCLUSIONS

Data stream mining is the process of extracting models and patterns from continuous and rapid data records. This field has been a focal area of research for over a decade and is driven by the massive amount of data that is continuously generated at high rates by different emerging real time

doi: 10.2788/854791

applications, from power systems (where stream mining algorithms have been implemented for synchrophasor data to meet quick decision making requirement of future situational awareness applications in power systems and smart grids) to its potential for handling time-series data from space.

It is important for the Earth remote sensing community to explore new analysis techniques for efficiently handling time series data in real time. Some of the data stream mining approaches discussed in this paper are excellent research topics for the community.

6. REFERENCES

1. B. Babcock, S. Babu, M. Datar, R. Motwani, and J. Widom, "Models and issues in data stream systems," In *Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems (PODS).* Madison, Wisconsin, pp.1-16,2002.
2. P. Domingosand G. Hulten,"Mining high-speed data streams," In *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '00)*. ACM, New York, NY, USA, 71-80. DOI=10.1145/347090.347107, 2000.
3. M. Kholghiand M. Keyvanpour, "An analytical framework for data stream mining techniques based on challenges and requirements," *Intl J Eng Sci and Technol*. arXiv preprint arXiv:1105.1950, 2011.
4. M. Garofalakis, J. Gehrke, and R. Rastogi,"Querying and mining data streams: you only get one look a tutorial," *SIGMOD Conference*. 635,2002.
5. Aggarwal,C. *Data streams: models and algorithms*, Springer, New York,2007.
6. Chu,F., *Mining Techniques for Data Streams and Sequences*,Ph.D. Dissertation, University of California, 2005.
7. C. Aggarwal, J. Han, J. Wang, and P. Yu, "A framework for clustering evolving data streams," In *Proceedings of the 29th VLDB Conference*,Berlin, pp. 81–92,2003.
8. C. Aggarwal, J. Han, J.Wang, and P. Yu, "A framework for projected clustering of high dimensional data streams", In *Proceedings of 2004 International Conference on Very Large Data Bases*,Toronto, Canada, 2004.
9. W. Hoeffding, "Probability inequalities for sums of bounded random variables," *J Am Stat Assoc* 58, pp. 13–30, 1963.
10. G. Hulten, L. Spencer, and P. Domingos,"Mining time-changing data streams," In *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*,ACM Press. New York, pp. 97–106, 2001.
11. J. Lin, E. Keogh, S. Lonardi, and B. Chiu, "A symbolic representation of time series, with implications for streaming algorithms,"In *Proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*,San Diego, California, pp. 2–11, 2003.
12. M. Gaber, A. Zaslavsky, and S. Krishnaswamy, "Mining data streams: a review,"SIGMOD Rec. 34, 2, pp. 18-26,DOI=10.1145/1083784.1083789, June 2005.
13. J. Gama, M. Gaber, and S. Krishnaswamy, "Data stream mining: from theory to applications and from stationary to mobile," *In Twenty-Fifth Symposium on Applied Computing*,Sierre, Switzerland. Accessed October 20, 2011.
14. P. Domingosand G. Hulten, "Mining high-speed data streams," In *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining* (KDD '00),ACM, New York, NY, USA, pp. 71-80,DOI=10.1145/347090.347107, 2000.
15. V. Ganti, J. Gehrke, and R. Ramakrishnan, "Mining data streams under block evolution," *SIGKDD Explor 3*, 2, 2002.
16. M. Gaber M, "Advances in data stream mining," *Wiley Interdisciplinary Rev: Data Min knowl Discov* 2, 1, pp. 79-85,2012.
17. N. Dahal, O. Abuomar, R. King, and V. Madani, "Event stream processing for improved situational awareness in the smart grid," *J. Expert Syst Appl*,DOI: 10.1016/j.eswa.2015.05.003, 2015.
18. N. Dahal, R. King, and V. Madani,"Online dimension reduction of synchrophasor data," *IEEE Transmission and Distribution Conference and Exposition (T&D)*,IEEE PES,2012.

LINKING BIG DATA FROM SPACE TO APPS ON EARTH

Uwe Voges[1], Peter Becker[2], Udo Einspanier[1]

[1]con terra GmbH (Münster, Germany)
eMail: u.(voges|einspanier)@conterra.de

[2]Esri Inc (Redlands, CA, USA)
eMail: pbecker@esri.com

ABSTRACT

For an easy and reliable access to Earth Observation (EO) data it should be served through cloud based web-servers, accessible by GEO-applications via simple REST-interfaces. This is offered by Esri´s ArcGIS Online (AGOL) combined with ArcGIS Image Extension for Server. For easier finding and discovery, metadata about the EO data published as RDF-based Linked Open Data (LOD) and being linked to the EO- and Non-EO world could be a promising approach.

Index Terms— AGOL, ArcGIS, Cloud, Earth Observation, GeoDCAT-AP, Linked Open Data, LOD, RDF

1. INTRODUCTION

For typical developers of GEO-enabled applications (GEO-apps) it is often hard (even impossible) to find and get access to EO data and products. The reasons are manifold:
1. EO data is not directly accessible as it is mostly stored in vendor specific archives. If generally accessible, a user has to register first, then try to find, order and download the data.
2. EO data is delivered as huge datasets with specific formats requiring preprocessing (e.g. format conversions or re-projections) before it can be used.
3. Finding EO data is complicated as often being described with metadata based on EO/GEO specific metadata models which are stored in specific metadata catalogs providing EO/GEO specific search interfaces. The specific metadata models do mostly not rely on common vocabularies and ontologies and non-geospatial objects are often re-modeled instead of linking to it. This prevents the data from being part of a wider network (which would make the finding easier just by browsing to it via linked data descriptions). Further the data objects itself are often not accessible by simple RESTful URLs. An example is ESA HMA (Heterogeneous Missions Access) [http://earth.esa.int/web/gscb/hma-standards]. The metadata models, discovery and access mechanism standardized here are well suited for B2B-scenarios. But for easy access as required in an end-user-market (B2C) these metadata models and catalog interfaces are too complex and not widely known.

2. REQUIREMENTS / SOLUTION

2.1. Reliable and performant web accessible EO services

The most important requirement for a GEO-app developer is an easy, reliable and performant web based access to EO data and products. EO data and products should be served by EO services through simple REST-interfaces. Ideally, the web-servers should reside in a cloud environment, gaining the advantages for cloud-based access: scalability, performance, availability, security and administration.
Esri´s ArcGIS Online (AGOL) [arcgis.com] combined with ArcGIS Image Extension for Server [esri.com/software/arcgis/arcgisserver/extensions] provide a scalable cloud-based content management system (control who can or cannot get access), which allows sharing geospatial data including EO imagery as services. Everyone can gain access to this and can develop a range of GEO-apps. Such services include ArcGIS Image Services, which are able to serve out very large collections of EO data. These services reference low-resolution browse imagery or the full resolution source datasets along with all associated metadata. The dynamic Image Services can efficiently transform the data from the source pixels into a wide range of information products by applying on the fly processing functions (e.g. orthorectification and projection, pixel based processing, rendering, download) as the data is accessed. A single service is therefor able to return not only the original data value, but also products such as vegetation indices from optical imagery or hill shades from elevation data. Such access enable the authoritative data source to be stored once, while providing a great range of products without the need for duplicate storage. The creation of these services can easily be automated such that the services are continually updated with new imagery as soon as it becomes available.
The client applications can access such services as if all the different image products were stored locally. Users are then able to define additional processing to be performed by the server with only the resulting data being returned. Access to the EO data for GEO-app developers is provided by the Image Service REST API, but also by OGC conformant

doi: 10.2788/854791

(temporal enabled) WMS- and WCS-services. A WMS [1] provides pictures of geospatial data to be portrayed as static maps. Unlike this a WCS [2] provides access to coverage data in forms that are useful for client-side rendering, as input into scientific models, and for other clients. Use of WMS is very restrictive in terms of its ability to define processing to be applied. In contrast, GEO-apps using the Image Service REST API can be much more interactive and dynamic by directly getting access to all metadata and processing through the same service. GEO-apps consuming Image Services can define aspects such as the time period of the imagery, projection and re-sampling methods to be used. They can define processing to be applied on the server and the compression for transmission, enabling use of the Image Services over low-bandwidth Internet connections.
It is easy to develop webapp builder apps or integrate the service into mobile-based GEO-apps, or use in storymaps etc.

2.2. Easier finding and linking to web accessible EO data

If a client already knows the EO services´ URL he can directly use it within a GEO-app. If he knows that the desired service is available in a specific service repository (like AGOL) he can try to search or browse for it there. The more likely case is that the client needs to search for a desired service with some kind of search engine. To be indexed by a search engine there is a need for metadata about the services. In the GEO/EO-domain usually specific ISO19115/ ISO19119 [3][4] based metadata models (in Europe additionally following the INSPIRE [5] rules) are used. The corresponding metadata instances are searched with specific catalogue services providing specific filter languages.
To more easily find, discover and link to the EO Services Linked Open Data (LOD) is a promising approach. LOD is usually based on the common Resource Description Framework (RDF) [6] where statements about arbitrary things (resources) are described by Subject, Predicate and Object Triples. Any kind of these elements are identified by Unique Resource Identifiers (URI´s). When those URI´s are HTTP URLs someone can directly look up the elements. The most important thing is to include links to other URIs by using concepts of common vocabularies and ontologies providing the semantic of the concepts (e.g. FOAF [7], GeoDCAT-AP [8] or GeoSPARQL [9]). The links shall also comprise links to the EO Service URL(s) using the concept "accessURL" of the GeoDCAT-AP vocabulary.
The linking based on these concepts to other URIs generate graphs which allow the discovery of more things and which minimizes replications (e.g. linking to and from the LOD cloud (http://lod-cloud.net), e.g. DBpedia, GeoNames).
RDF-based EO metadata can be stored in a triple store and searched with the domain independent standardized query language SPARQL which is just extended by a few Geo-domain specific capabilities [9]. Suitable triple stores (mostly providing a (Geo)SPARQL interface) are those of well-known Open Data Portals. Here the single RDF information should also be accessible via their URL. With adequate (Geo)SPARQL processors semantic assisted searches with reasoning capabilities (using the power of the domain specific ontologies) could better find the desired information.

3. EXAMPLE

In the following a limited example will be presented where a few kinds of EO data (products) are served as ArcGIS services providing Image Service and OGC WMS/WCS interfaces. Different generic GEO-apps consume these services. Apps consuming Image Services can define dynamic access e.g. dynamic definition of compression for transmission, download and access to all metadata. The services are described with RDF metadata which is being harvested by a Data Portal´s triple store where the services can be searched and discovered.

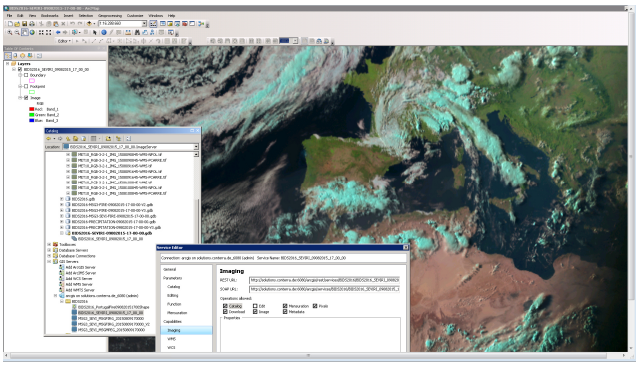

Figure 1: Publishing SEVIRI Image Data as EO service with ArcMap

3.1. Serving EO data as ArcGIS services

The example is based on three EUMETSAT products which were ordered online via EUMETSAT´s EO Portal (eoportal.eumetsat.int) for one point in time (9.8.2015 17:00):

- Active Fire Monitoring (GRIB) - MSG (fire detection product indicating the presence of fire within a pixel)
- High Rate SEVIRI Level 1.5 Image Data - MSG
- Multi-Sensor Precipitation Estimate (GRIB) - MSG

Additionally a dataset (shape) is included with precalculated buffers around fires (just Spain/Portugal) for better visibility. After reception from EUMETSAT the datasets were preprocessed (e.g. translated with GDAL [12] into GeoTIFF). The SEVIRI data were already translated with a special tool at EUMETSAT into GeoTIFF. After the preprocessing, serveral products were loaded into so called mosaic datasets (using Esri´s ArcGIS Image Extension) and then the mosaic

datasets and raster files were published from within Esri´s ArcMap as ArcGIS services providing an Image Service REST interface and/or OGC WMS/WCS interfaces. For the example it was sufficient to publish the services into a local web-accessible ArcGIS installation. For a real scenario the better choice for sharing would be a publication into AGOL (as being more reliable, scalable and performant).

3.2. Publishing EO services with 5* LOD

In our sample the services were described manually by RDF based on the common controlled vocabularies and ontologies GeoDCAT-AP, EuroVoc [10], INSPIRE [11] and FOAF [7] and include RESTful links to the ArcGIS services. An example is shown in Fig. 2.

```
<dcat:dataset>
  <dcat:Dataset rdf:about="de.conterra...FireUseCase.SEVIRI.09082015_17_00_00">
    <dct:title xml:lang="en">High Rate SEVIRI Level 1.5 Image Data - MSG...</dct:title>
    <dct:description xml:lang="en">Rectified (level 1.5) Meteosat SEVIRI...</dct:description>
    <dct:spatial>
      <dct:Location><locn:geometry rdf:datatype="http://.../ont/geosparql#gmlLiteral">
        <![CDATA[<gml:Envelope srsName="http://www.opengis.net/def/EPSG/0/4326">
        <gml:lowerCorner>-81.00 -79.00 </gml:lowerCorner><gml:upperCorner>
        81.00 79.00</gml:upperCorner></gml:Envelope>]]></locn:geometry>
      </dct:Location>
    </dct:spatial>
    <dct:temporal><dct:PeriodOfTime>
        <schema:startDate rdf:datatype="http://...#date">2015-08-09</schema:startDate>
      </dct:PeriodOfTime>
    </dct:temporal>
    <dct:publisher>
      <foaf:Organization>
        <foaf:name>EUMETSAT</foaf:name>
        <foaf:mbox rdf:resource="mailto:ops@eumetsat.int"/>
      </foaf:Organization>
    </dct:publisher>
    <dct:created rdf:datatype="http://.../XMLSchema#date">2015-12-23</dct:created>
    <dct:type
    rdf:resource="http://inspire.ec.europa.eu/metadata-codelist/SpatialDataServiceType/view"/>
    <!-- Atmospheric Conditions -->
    <dcat:theme rdf:resource="http://inspire.ec.europa.eu/theme/ac-mf"/>
    <dcat:theme rdf:resource="http://eurovoc.europa.eu/89"/>
    <dcat:keyword xml:lang="en">Cloud</dcat:keyword>
    <dct:provenance><dct:ProvenanceStatement>
        <rdfs:label xml:lang="en">Sat:MSG,Instrument:SEVIRI(Type:OPTICAL)</rdfs:label>
      </dct:ProvenanceStatement>
    </dct:provenance>
    <dcat:distribution><dcat:Distribution>
      <dct:title lang="en">view</dct:title>
      <dcat:accessURL rdf:resource="http://solutions.conterra.de/arcgis/services/BIDS2016/
        BIDS2016_SEVIRI_09082015_17_00_00/ImageServer/WMSServer?
        service=WMS&amp;request=GetCapabilities"/>
      <dct:format rdf:parseType="Resource"><rdfs:label>WMS</rdfs:label>
      </dct:format>
      </dcat:Distribution>
    </dcat:distribution>
  </dcat:Dataset>
</dcat:dataset>
```

Figure 2: GeoDCAT-AP Example (condensed)

During RDF ingestion every RDF description gets assigned a final URL based on the identifier. Via this URL the description can be accessed in different formats, e.g. HTML or RDF. The format returned depends on the requested mime-type. HTML encodings could be used to let mainstream search engines index the EO services, e.g. by providing the URL´s in a sitemap registered with the search engine. But the problem here is that targeted searches (e.g. using domain specific vocabularies/ontologies or spatial/ temporal filters) cannot be processed. Consequential (in a real world scenario) the client would often get lots of very similar search results from which it is hard to detect the required service.

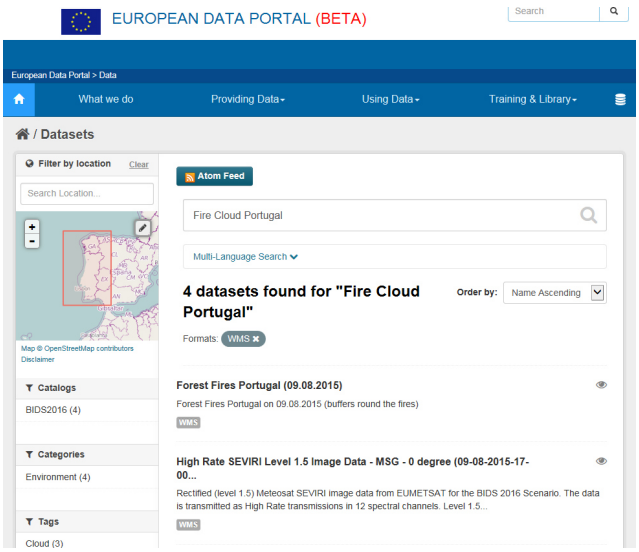

Figure 3: Search in the EU Data Portal

It is more promising to publish or harvest (as done here) the RDF descriptions into a well-known data portal (running an RDF store) which is more tailored to finding and using GEO/EO data and services. In our example we used a (protected) test version of the new EU Data Portal (EDP) [(test.)europeandataportal.eu/] (Fig 3+4).

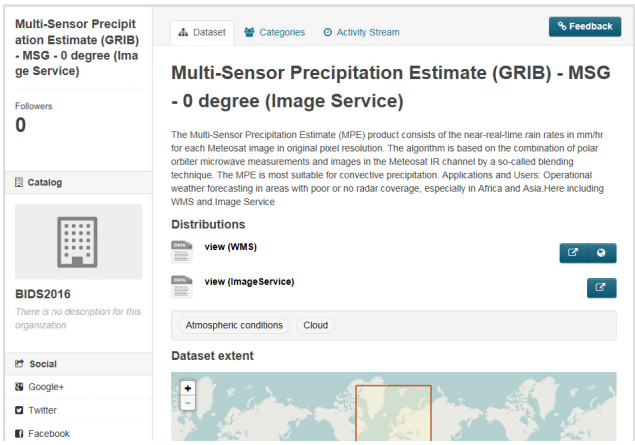

Figure 4: Search results (details) in the EU Data Portal

For broader visibility (in a real world scenario) it would be important to link the RDF data with the LOD cloud (lod-cloud.net).

doi: 10.2788/854791

3.3. Finding and using EO Services with GEO-apps

With the RDF descriptions ingested into the EDP it is now possible to proceed searches e.g. based on categories, tags, keywords, boundingBox, format, etc. (see Fig. 3). With the SPARQL endpoint of the EDP it is further possible to proceed much more targeted searches (see Fig. 5).

```
SELECT ?AccessURL ?DT
WHERE {
        ?Dataset a dcat:Dataset .
        ?Dataset dcat:keyword "Cloud" .
        ?Dataset dcat:distribution ?Distribution .
        ?Distribution dcat:accessURL ?AccessURL .
        ?Dataset dct:temporal ?Temp .
        ?Temp schema:startDate ?DT .
        FILTER (?DT > "2015-08-05"^^xsd:date && ?DT < "2015-08-15"^^xsd:date) }
```

Figure 5: SPARQL query on EU Data Portals´ RDF-store

From the detail information of a search results´ entry (providing a WMS interface - see Fig. 4, globe icon) it is possible to open the EDP map client which loads the WMS (see Fig 6). This map.apps (conterra.de/mapapps) based map client is able to interpret the GeoDCAT-AP metadata and to identify the WMS access point (URL).

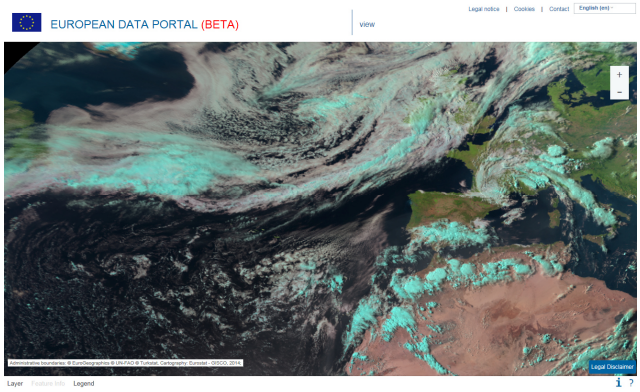

Figure 6: EDP WMS client - SEVIRI Data - MSG (9.8.2015 17:00)

The WMS endpoints (URLs) found can also be used by an arbitrary WMS client (see Fig 7).

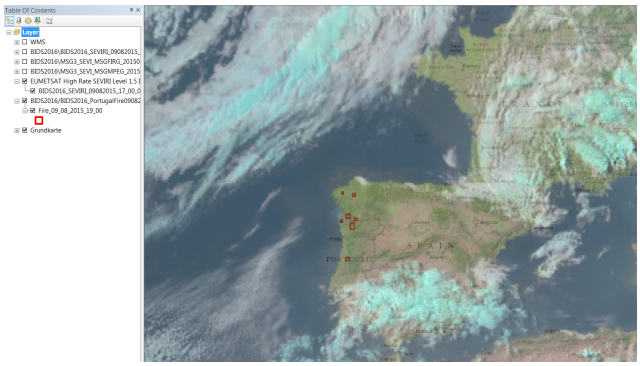

Figure 7: ArcMap using SEVIRI + Forest Fires WMS

When search results entries provide dynamic ArcGIS Image Service endpoints those can be used by a proper client (as ArcMap, Fig 8), e.g. for applying on the fly processing or visualization and downloading parts of the image (Fig. 8.)

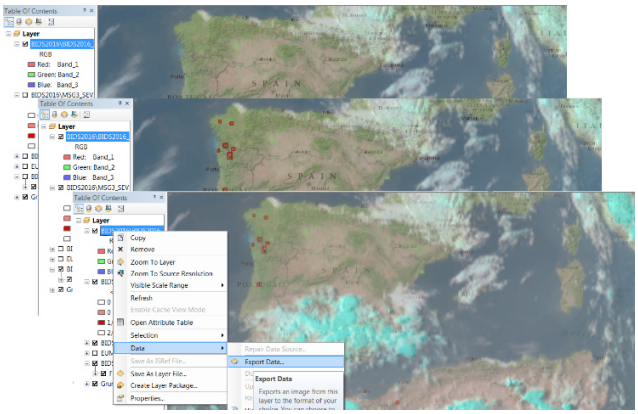

Figure 8: Image Services (9.8.2015): a: SEVIRI, b: a+Fire, c: a+b+ Precipitation estimation (incl. download / metadata access options)

4. CONCLUSION

The publication of EO data and products as OGC WMS/WCS and Esri´s ArcGIS Image Services (ideally deployed in a cloud-based environment as ArcGIS Online), accessible via REST-interfaces and their description with Linked Open Data and publication (for search and discovery) into a Data Portal (e.g. European Data Portal) is a promising approach to bring big data from space to geo-applications on earth.

5. REFERENCES

[1] OpenGIS Web Map Server Implementation Specification. Open Geospatial Consortium Inc. OGC 06-042. Version 1.3.0.

[2] OGC WCS 2.0.1. Open Geospatial Consortium. OGC 09-110r4

[3] ISO 19115:2003, Geographic information – Metadata

[4] ISO 19119, Geographic Information – Services, 2006.

[5] INSPIRE Metadata Implementing Rules: Technical Guidelines based on EN ISO 19115 and EN ISO 19119, EC JRC. 2008.

[6] Resource Description Framework (RDF). W3C. 2014. www.w3.org/RDF

[7] FOAF Vocabulary Specification. http://xmlns.com/foaf/spec/

[8] GeoDCAT-AP 1.0. EC. joinup.ec.europa.eu/node/139283

[9] GeoSPARQL - A Geographic Query Language for RDF Data. Open Geospatial Consortium Inc. OGC 11-052r4. V 1.0. opengeospatial.org/standards/geosparql

[10] EuroVoc, EU multilingual thesaurus. eurovoc.europa.eu

[11] INSPIRE metadata code list register. EU. http://inspire.ec.europa.eu/metadata-codelist

[12] GDAL - Geospatial Data Abstraction Library. www.gdal.org/

RETURN ON EXPERIENCE AND PERSPECTIVES ABOUT HOW TO TURN BIG DATA FROM SPACE INTO BIG VALUE ON EARTH

Olivier Flebus, Lydwine Gross-Colzy, Julien Joyeux

Capgemini Aerospace & Defense, Toulouse, France

ABSTRACT

In this paper we present our own vision for building spatial applications that make use of "big data from space" and create digital downstream services generating revenues and jobs for organizations and people on Earth ☺. This vision is based on our 5-year experience developing Earth Observation scientific algorithms in a way that follows the rules of the Digital world and state-of-the-art Big Data approaches and solutions[1]. We've also started to implement and deliver this vision into our own Big Data Platform: Tech4Earth [1]. We also provide our point of view about what future Big Data approaches and solutions like deep learning could mean for spatial applications.

Index Terms— Tech4Earth, Earth Observation, Digital Services, Platform, New Processing Capabilities, Hadoop, Spark, Machine Learning, Deep Learning.

1. SPACE SCIENTISTS HAVE STARTED TO FACE NEW CHALLENGES

Multiple factors have contributed to space scientists getting aware that new approaches are required for the future.

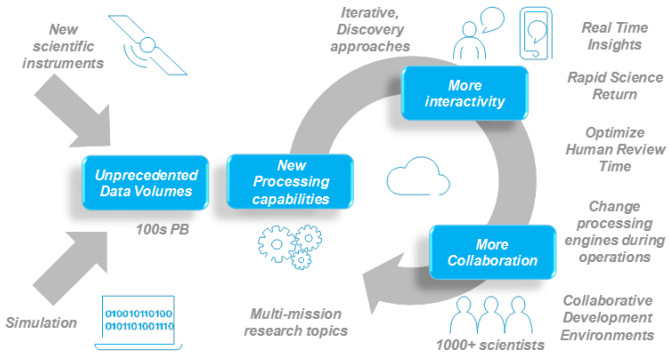

Figure 1 - Space Scientists have started to face new challenges

Unprecedented data volumes

It's a fact: data volumes acquired by space instruments and generated by spatial missions are growing exponentially, and thus curiosity toward "Big Data" technical solutions coming from Internet Giants has naturally risen. Nevertheless, we believe that these growing volumes are not the sole reason that justifies new "Big Data" approaches for processing chains.

More interactivity with data

The traditional – quite rigid – way of designing processing chains sometimes does not guarantee to make scientific discoveries. While working with astrophysicists mapping the dark universe (Science Ground Segment of the Euclid Mission [2]) and the ones building the first gravitational waves observatory in space (Data Processing Center of the eLISA mission [3][4]), we have been convinced that finding a more agile way to conduct experiments with processing chains is critical for these open, discovery-oriented contexts.

New kinds of processing chains

Conducting experiments based on measurements datasets typically requires human analyses to be conclusive. When simulation and/or a large number of experimentations are involved, new kinds of analysis and modeling methods are required because the human working time is far more difficult to increase than the computing resources. That's when statistical approaches and machine learning techniques become particularly relevant and useful.

More collaboration

Collaboration is more and more relevant for spatial missions. They combine multiple science areas and sometimes involve more than one thousand people. Collaboration is also a key success factor for creating new services that disrupt the value that can be extracted from imagery. Moreover, collaboration is key at every stage (design, implementation, testing and operations) of the processing chains lifecycle.

More value extracted from existing data

Getting new insights does not require new data, existing data can be leveraged; for instance, historical data are valuable to train machine learning algorithms. Combining multiple kinds of data (e.g. spatial data and data from in-situ sensors) is another way to create value, beyond traditional calibration/validation processes. Even if they design new instruments that will provide much more data, space

[1] We define what we call Big Data approaches and solutions in part 3 of this paper.

doi: 10.2788/854791

scientists also have access to huge archives that are not processed for most of them.

Focusing on (business or scientific) value, setting up a "flow" of value by delivering more frequently, relying on people, collaboration, feedback and empiricism… All these characteristics are the ones of the agile movement that started with software development [5] and is now evolving towards bringing agility to systems engineering (including hardware) and product or services management, covering both development and operations (DevOps).

2. DIGITAL SERVICES

Digital Services are the main reason why disruptive big data solutions (e.g: NoSQL storage and distributed frameworks like hadoop) had been invented by Internet Giants[2]. Digital Services have to be running at any time, accessed from anywhere in the world from any type of device. They also have to deal in an invisible way with extremely volatile workloads, hardware and software failures, always providing results to more than one billion active end-users in less than a second. Digital Services have redefined the state-of-the-art for the most flexible, versatile and elastic components ever engineered.

3. HOW TO BEST LEVERAGE BIG DATA SOLUTIONS FOR SPATIAL APPLICATIONS

What do we call Big Data Solutions?

We consider that Big Data approaches and solutions are defined by the following statements/principles:

- They are derived from approaches and solutions primarily developed by Internet organizations that were facing contexts **beyond the frontier of what was achievable** with established approaches and solutions. This encompasses the common Big Data definition with the 3-to-5 V's: Volume, Velocity, Variety, Veracity and Value. You enter the Big Data world when one of the first three characteristics goes beyond "normal". Thus Fast Data (with "normal" volumes) can also be considered as Big Data, as soon as it still provides true value.
- On the one hand, Big Data approaches and solutions are **disruptive** in the sense that they rely on **radically different design principles or patterns**: *design for failure* as a way to build very resilient (distributed) systems by including (or even generating) failures in the real life of the system, make extreme scalability easy to use, move the processing where the data is stored, etc.
- On the other hand, Big Data approaches and solutions **have offered a new life to some legacy approaches**, especially in the field of Artificial Intelligence. Traditional, "legacy" machine learning algorithms have become far more valuable by being trained with extremely large datasets.
- Big Data approaches and solutions **provide agility**. They are designed to help organizations discover new insights from data in motion as well as data at rest; decrease the time-to-market for digital services by combining development and operations; scale to unanticipated use without compromising the user experience, especially without any outage, etc. The agility of Big Data platforms also comes from underlying Cloud services.

Big Data Solutions for Scientific Processing Chains?

Spatial processing chains have traditionally been considered as sub-system of the whole satellite mission. They have been designed deriving mission requirements, implemented and fully tested before mission launch, and it has always been difficult to change them in the operation phase.

Putting together the traditional approach for processing chains and the definition of Big Data approaches and solutions clearly highlights the fact that using Big Data for Science cannot be considered as obvious and direct. Especially it's really not only a matter of data volumes.

Our lessons learned from 5 years of experimentations with Big Data for Science

We started implementing scientific processing chains with Big Data Solutions (Hadoop Map/Reduce in that case) in 2011 with the ESA FAAPS project (Fully Automated Aqua Processing Service) [6][7]. The project objectives were to develop, validate, demonstrate and assess the benefits for flood management authorities of an operational service delivering NRT flood extent maps generated by processing ESA satellite data.

[2] GAFA (Google, Apple, Facebook, Amazon) are the most popular Internet Giants, which have grown among the largest market capitalizations in the world and started disrupting other not-only-digital markets, such as Space launchers (Blue Origin) or Earth Observation (Skybox Imaging).

doi: 10.2788/854791

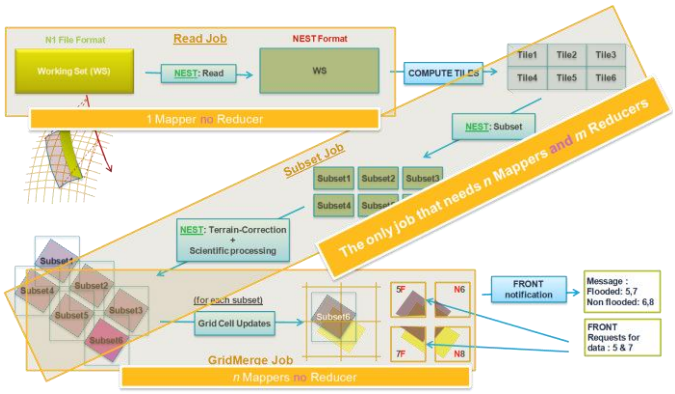

Figure 2 – Map/Reduce applied to FAAPS processing chain

A part of the whole FAAPS processing chain was implemented in a "pure" Map/Reduce way. The main achievement has been to demonstrate:

- that it's possible to process imagery in near real-time (NRT) without having to care about (and code) any orchestration issue,
- that it's possible to handle huge reprocessing tasks with the same code and software environment than the one used for NRT processing, just by relying on the native scalability of Map/Reduce (and scaling the Cloud infrastructures).

Our FAAPS experience also highlighted specific difficulties in trying to "hadoopize" existing scientific processing chains:

- How sensor data (especially imagery) should be explicitly (we are now cutting down sentinel-1 images into 81 pieces) or implicitly (HDFS chunks) distributed.
- How the processing chain and most of its components should be reengineered according to foreseen workload and gains.

Following the FAAPS achievement we kept experimenting with three objectives:

- Implement scientific processing chains using not only Map/Reduce but also additional "engines" that have started being integrated in Hadoop (Spark, R).
- Imagine new algorithms and approaches (such as neural networks and machine learning) to model and solve Earth Observation problems.
- Putting together the previous activities to design and build a Platform leveraging Big Data Solutions for Spatial Applications [1].

Our experience has inspired the following vision and principles for implementing Scientific Processing Chains with Big Data solutions:

- A **single platform** and unified approach **for the complete value chain**, from scientific processing to digital services. It concretely means leveraging approaches and solution from the Digital world for implementing scientific processing chains, not the opposite. It also means embracing a DevOps approach combining the stakes, objectives and constraints of both Development and Operations.
- **Keep things simple**, fast & easily understandable even if it's less efficient from a resource point of view. We're convinced that it's better to be able to scale and update an algorithm that is not 100% efficient than building a 100% efficient algorithm that is understood by a few people and runs only on a fixed infrastructure.
- Be ready to **completely reconsider the workload and lifecycle of your processing chain**. When using machine learning, design and development phases are replaced by training and prediction phases. The training phase requires massive infrastructures and human work. The prediction phase is more lightweight: it could ultimately be executed with a laptop or even be embedded into sensors!

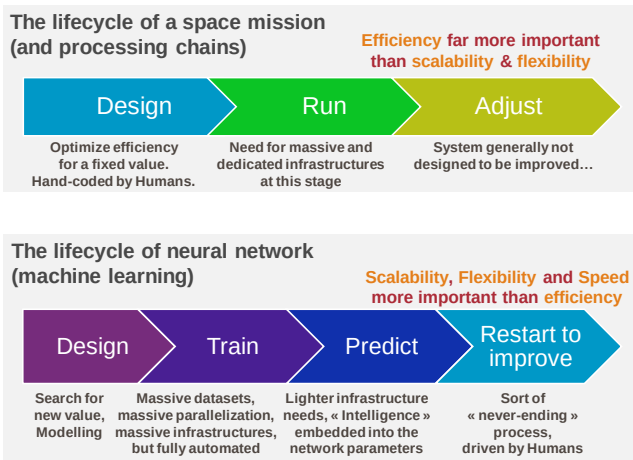

Figure 3 - The lifecycle of a traditional processing chain and the one of a neural network

Thus, it leads to radically different lifecycles and workloads distribution. The training phase involves massive datasets and high processing requirements. In most cases (as most applications still use supervised learning), it also requires these massive datasets to have previously been labeled, most certainly "manually" by humans [8][9].

4. PERSPECTIVES: WHAT CAN WE EXPECT FROM FUTURE BIG DATA APPROACHES AND SOLUTIONS

One of the main characteristics of the world of Big Data Solutions is that they all evolve at a very fast pace. This is particularly true for NoSQL databases and software frameworks like Hadoop, Spark, Kafka or Storm, which constitutes the current building blocks of modern Big Data platforms.

Whereas Map/Reduce is relying on disk only, Spark has brought an efficient to use memory – which is incomparably faster – while keeping it very simple to program distributed tasks.

A next wave is to be able to use the best-in-class solution for massive distributed processing that has been used for decades by scientists: hardware accelerators (GPU). Combining Hadoop-like solutions for scalability and GPUs for performance could bring the better of two worlds to future processing chains.

This trend has already started for deep learning: Big Data players (Google, Microsoft, Facebook, Baidu) as well as Hardware vendors (NVIDIA) are working on solutions that exploit clusters of GPUs. Thus it's quite possible to imagine platforms that will combine:

- hadoop-style high-level programming frameworks hiding the complexity of distribution, automatically enabling elasticity and flexibility for analytics platforms,
- Unified memory between GPUs, CPUs and the previous high-level frameworks bringing top-class performance for every kind of processing,
- Improved hardware accelerators (and clusters) dedicated to deep learning.

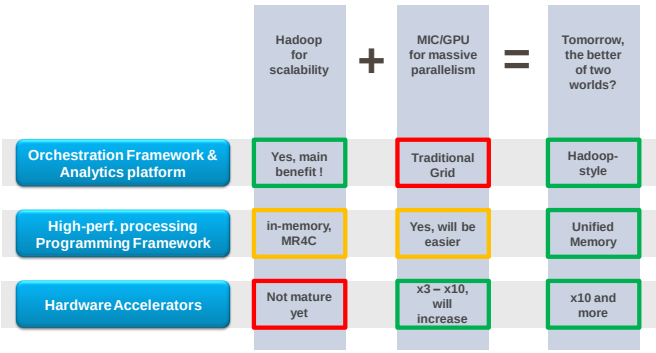

Figure 4 - Better of two worlds: Hadoop + GPU

6. REFERENCES

[1] O. Flebus, L. Gross-Colzy, J. Joyeux "Tech4Earth: Capgemini's platform to foster the digital transformation of the Earth Observation ecosystem", 2015, http://bit.ly/tech4earth

[2] F. Pasian & al, "The Science Ground Segment for the ESA Euclid mission", 2012, SPIE http://adsabs.harvard.edu/abs/2012SPIE.8451E..04P

[3] K. Danzmann, et al., The Gravitational Universe, 2013. Whitepaper on eLISA for the ESA. https://www.elisascience.org/whitepaper/

[4] Pau Amaro-Seoane, et al., "Low-frequency gravitational-wave science with eLISA/NGO", 2012, http://dx.doi.org/10.1088/0264-9381/29/12/124016

[5] Kent Beck, Mike Beedle, Arie van Bennekum, Alistair Cockburn, Ward Cunningham, Martin Fowler, James Grenning, Jim Highsmith, Andrew Hunt, Ron Jeffries, Jon Kern, Brian Marick, Robert C. Martin, Steve Mellor, Ken Schwaber, Jeff Sutherland, Dave Thomas, "Manifesto for Agile Software Development", 2001, http://www.agilemanifesto.org/

[6] FAAPS - Fully Automated Aqua Processing Service, https://artes-apps.esa.int/projects/faaps

[7] C. Cormont, "Generic architecture for managing huge volume of data", PV2013, http://bit.ly/1moZJ8H

[8] Fei-Fei Li, et al., "ImageNet - crowdsourcing, benchmarking & other cool things", 2010, http://bit.ly/1SFwYmW

[9] Fei-Fei Li, "How we're teaching computers to understand pictures", 2015, http://bit.ly/20SLTuJ

Author Index

Keyword Index

Page numbers below always refer to the *first* page of the paper(s) indexed by the corresponding keyword.

Abstract

Big Data from Space refers to Earth and Space observation data collected by space-borne and ground-based sensors. Whether for Earth or Space observation, they qualify being called 'big data' given the sheer volume of sensed data (archived data reaching the exabyte scale), their high velocity (new data is acquired almost on a continuous basis and with an increasing rate), their variety (data is delivered by sensors acting over various frequencies of the electromagnetic spectrum in passive and active modes), as well as their veracity (sensed data is associated with uncertainty and accuracy measurements). Last but not least, the value of big data from space depends on our capacity to extract information and meaning from them.

The goal of the Big Data from Space conference is to bring together researchers, engineers, developers, and users in the area of Big Data from Space. It is co-organised by ESA, the Joint Research Centre (JRC) of the European Commission, and the European Union Satellite Centre (SatCen) and was held at the auditorio de Tenerife (Santa Cruz de Tenerife, Spain) from the 15th to the 17th of March 2016.

These proceedings consist of a collection of 108 short papers corresponding to the oral and poster presentations presented at the conference. They are organised in sections matching the order of the conference sessions followed by the contributions that were presented during the poster session, also organised by topics. They provide a snapshot of the current research activities, developments, and initiatives in Big Data from Space.

LB-NA-27775-EN-N

JRC Mission

As the Commission's in-house science service, the Joint Research Centre's mission is to provide EU policies with independent, evidence-based scientific and technical support throughout the whole policy cycle.

Working in close cooperation with policy Directorates-General, the JRC addresses key societal challenges while stimulating innovation through developing new methods, tools and standards, and sharing its know-how with the Member States, the scientific community and international partners.

Serving society
Stimulating innovation
Supporting legislation

Publications Office

doi:10.2788/854791
ISBN 978-92-79-56980-7