JRC TECHNICAL REPORTS

Maritime Awareness Systems Performance in the Western Indian Ocean 2014-2015

Results from the PMAR-MASE project

Harm Greidanus, Marlene Alvarez,
Vincenzo Gammieri, Carlos Santamaria,
Alfredo Alessandrini, Pietro Argentieri

2015

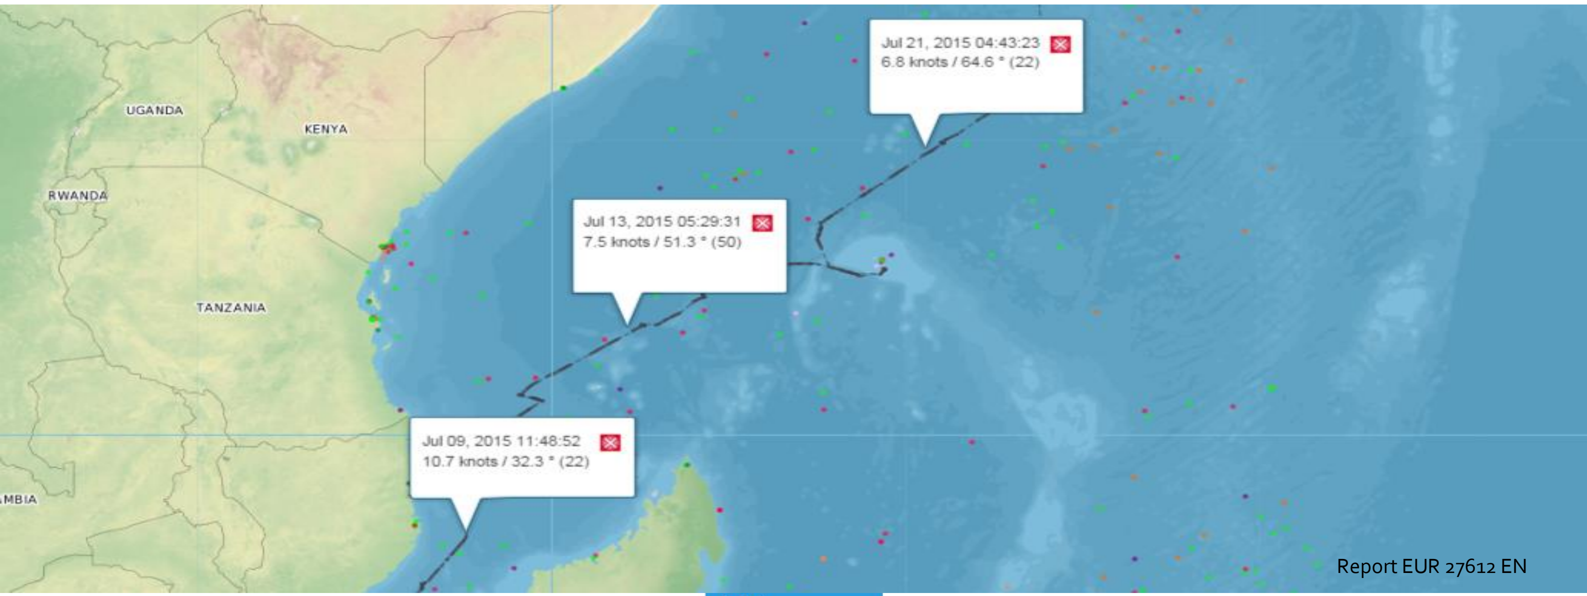

Report EUR 27612 EN

Joint Research Centre

European Commission
Joint Research Centre
Institute for the Security and Protection of the Citizen

Contact information
Harm Greidanus
Address: Joint Research Centre, Via Enrico Fermi 2749, TP 670, 21027 Ispra (VA), Italy
E-mail: harm.greidanus@jrc.ec.europa.eu
Tel.: +39 0332 78 9739
Fax: +39 0332 78 9156

https://ec.europa.eu/jrc/en/institutes/ipsc
https://ec.europa.eu/jrc/en/research-topic/piracy-maritime-awareness-and-risks

JRC97935

EUR 27612 EN

ISBN 978-92-79-54033-2

ISSN 1831-9424

doi:10.2788/420868

Luxembourg: Publications Office of the European Union, 2015

Cover illustration: Ship positions coloured by ship type and the track of one ship as seen in the PMAR Viewer.

Executive summary

The PMAR-MASE project aimed to provide practical experience to authorities in the Eastern-Southern Africa / Indian Ocean region on wide-area maritime awareness. Data from automatic ship reporting systems (AIS and LRIT) have been received during one year over a large Area of Interest (AOI) in the Western Indian Ocean. The data were used to construct the Maritime Situational Picture (MSP) over the AOI, i.e. the map with the current positions of all reporting ships. The MSP has been served continuously in real time to two operational users in the region, between 15 Sep 2014 and 15 Sep 2015. The data have also been used to make monthly ship traffic density maps. Finally, a small number of satellite images have been collected to probe for the presence of ships that are not seen on the reporting systems. This report assesses the quality (completeness and timeliness) of the ship reporting data and the resulting MSP; it presents the derived monthly ship density maps, split between slow moving and fast moving traffic in order to understand ship activity; and it present results on the fraction of non-reporting ships based on the analysis of the satellite images.

The PMAR-MASE AOI stretched from the East African coast nearly out to the Maldives in the east, and between Oman in the north to Durban in the south. The data that were used are AIS data received by up to 16 satellites of four different operators, AIS data from various coastal networks in the region, LRIT data from the Flags that use the EU LRIT Data Centre (“EU Flags”), and space-based radar images from two different satellites.

The performance of the ship reporting systems was analysed in detail for one representative month, June 2015. During that month, on average 567k ship position messages were received every day (all from within the AOI). More than 80 % of these are satellite AIS, most of the remainder terrestrial AIS. LRIT only represents 0.1 % of the total number of position reports. Less than 2% are messages from Class B AIS ships (smaller ships). The number of different ships seen each day in the AOI is on average just over 1,500. About 8 % are Class B ships. The number of different ships seen during the entire month is 5,350 – much higher than the daily number, because ships move in and out of the AOI every day.

These numbers are the result of combining the data from all sources together (in June, that is 11 AIS satellites, several coastal networks and the EU LRIT). If only a single AIS satellite would be available, it would see only 82 % of the number of ships seen by all sensors together on a daily basis. This is because only a fraction of all AIS messages that are broadcast are received by any AIS receiver. The more sensors are combined, the more complete the ship traffic picture becomes. The combination of two satellites sees 89 % of what all sensors see; the combination of three satellites, 92 %; etc.; up to 96.5 % from 11 AIS satellites, the remainder being seen only by the terrestrial AIS and LRIT. LRIT on its own sees daily on average 178 ships, a low number because it is only the ships of the “EU Flags”. If the LRIT is combined with a single AIS satellite, which by itself sees daily on average 1,252 ships, then still the LRIT sees on average 27 ships that the AIS does not see. To the constellation of 11 satellites, that reveal the presence of 1,476 ships, LRIT still adds 11 ships missed by the AIS on an average daily basis.

Now looking at the number of ships in the MSP, meaning the number of ships present in the AOI at any one time, which is less than the daily number seen, it is found that one AIS satellite will put 1,188 ships in the MSP; two AIS satellites yield 1,226 ships; thee give 1,262; etc., up to 11 AIS satellites that with diminishing increases result in 1,345 ships in the MSP on average. LRIT alone yields 165 ships in the MSP on average; but adding LRIT to 11 AIS satellites still adds another 10 ships (0.7 %) to the AIS total. The MSPs are constructed from the data by a tracking process, which recognises and removes data errors (that would have a serious impact if left in). The tracking has a number of parameters that influence the results in absolute numbers, but less so in relative numbers (as in comparing N satellites to N+1 satellites).

The ship positions in the MSP are interpolated between position reports that have been received prior to and after the MSP time, or extrapolated from one of these if the other is not available. The further the MSP is away in time from the prior or later position report, the less accurate the interpolation result is. The time difference between the MSP time and the nearest position report is therefore a measure of the accuracy of the MSP positions. If one AIS satellite is used, this time difference is 2.6 hours on average. But the distribution of these time differences is very broad, so many individual time differences are much smaller or much larger. If 7 AIS satellites are used, this mean time difference is reduced to 1.0 hours. Adding more AIS satellites continues to reduce the mean time difference, but not by so much anymore. Using only LRIT, the time difference is on average 1.8 hours. This compares to the time difference when using 4 AIS satellites. Of course, the LRIT only tracks 165 ships in the MSP whereas 4 AIS satellites have 1,276 ships. The LRIT is more accurate despite its much lower number of messages because the messages are distributed more evenly over the day, leaving fewer large gaps.

Concerning the monthly ship density maps, they were constructed from satellite AIS and LRIT data, excluding the terrestrial AIS – which is only concentrated in a few coastal locations – to prevent a spatial bias. The maps were made with a spatial resolution of 0.15 degrees which is approximately 9 nautical miles or 16.7 km. The density counts the average number of ships present at any one time per unit of surface area. The maps were split in fast moving traffic (> 6.5 knots) and slow moving traffic (< 6.5 knots). The fast moving traffic is dominated by the main traffic routes, and the pattern does not change much from month to month. The slow moving traffic indicates activities at sea other than transiting, such as fishing or exploration, besides concentrations of ships at anchor or moored in or near ports. The patterns of these activities change much from month to month.

To sample the non-reporting ship traffic, 46 satellite images were analysed in detail. Of those, 29 were purchased, 12 were publicly available, and 5 were used via scientific collaboration. The images were located in front of the Somali coast, between Kenya-Comoros-Mozambique, over La Réunion and Mauritius, in the Gulf of Aden and in the middle of the Indian Ocean. In total, 272 targets were detected at sea in the images. Of these, 34 % could not be correlated with reporting ships. This percentage however varied much from area to area, the highest fraction of non-reporting ships being found off Somalia. In addition to these 46 satellite images that were analysed with human verification, more than 1,000 Sentinel-1 images that were collected in the area and during the project period under the EU’s Copernicus program, were analysed in fully automatic mode. Here, a fraction of 40 % of the targets could not be correlated with reporting ships. But this number is less accurate due to the lack of verification and is probably an overestimate.

Contents

1 Introduction

This report analyses and discusses the data that have been collected during the PMAR-MASE project for maritime surveillance of the Eastern-Southern Africa / Indian Ocean region. It is an Annex to the PMAR-MASE Final Report.

PMAR stands for Piracy, Maritime Awareness and Risks. The PMAR-MASE project was executed by the European Commission's Joint Research Centre (JRC) for the Indian Ocean Commission under the MASE program. The MASE program aims to support 'MAritime SEcurity' in the Eastern-Southern Africa / Indian Ocean region and is funded by the 10th European Development Fund (EDF). The purpose of PMAR-MASE was to provide practical experience to users and stakeholders of maritime surveillance in the region. Under PMAR-MASE, one year of ship reporting data has been collected over the Western Indian Ocean, during 15 Sep 2014 – 15 Sep 2015. The Area of Interest (AOI) was between 31 and 68 degrees East and between 30 degrees South and 19 degrees North (see Figure 1-1). These data have been used to provide a continuous, real-time Maritime Situational Picture (MSP) to two operational centres for one year – the Indian Ocean Commission's Anti-Piracy Unit in the Seychelles, and the Regional Maritime Rescue Coordination Centre (RMRCC) of the Kenya Maritime Authority in Mombasa.

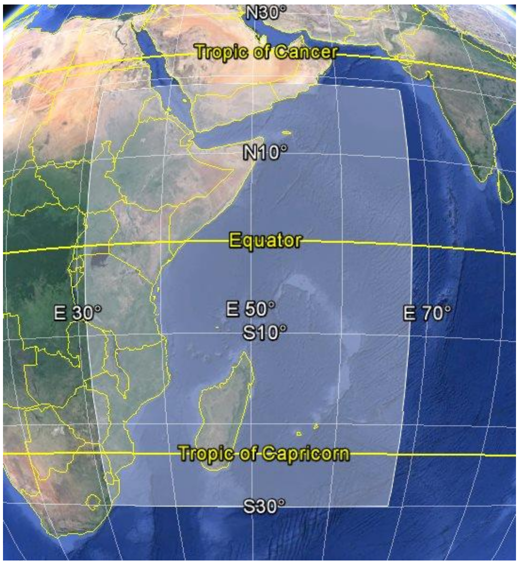

Figure 1-1. The AOI of the PMAR-MASE project, the area within which the maritime surveillance data were collected and the MSP was constructed. (Screenshot from Google Earth.)

Two main types of data have been collected. The foremost one is automatic ship reporting data from the AIS and LRIT systems. AIS (Automatic Identification System) is a system that has transponder equipment installed on ships, automatically broadcasting short VHF radio messages with the ship's identification, position and some other information. These messages are received by the AIS transponders on other ships in the vicinity, so that each ship is aware of the local traffic situation. The short messages are broadcasted frequently, at a variable rate which is of the order of once a minute. By using receivers on the coast or mounted on satellites, and feeding the data into a network, the AIS messages can be centrally collected to provide the monitoring of a certain AOI. LRIT (Long Range Identification and Tracking) is a similar system with transponders on ships, however in this case the short messages are only sent once every 6 hours and only to the Maritime Authority of the ship's Flag State. The messages that contain the ship's geographic position (latitude, longitude) are referred to as position messages or position reports. In the AIS message, the ship is identified by its so-called MMSI number. In the LRIT message the identification is by the ship's IMO number, which can be linked to an MMSI number in order to merge the two types of messages of the same ship into one stream. Only the larger ships carry the AIS and/or LRIT equipment, following global regulations from the UN's International Maritime Organisation (IMO). The precise carriage requirements are rather detailed, but roughly it concerns the ships larger than 300 tonnes. Some smaller ships may carry voluntary "Class B" AIS (the mandatory one for the larger ships being Class A).

The secondary type of maritime surveillance data that was used are ship detections from satellite images. This is a completely different type of data, where the ships are not identified, but only the presence and position of some object can be established, not even with the certainty that the object is a ship. These data are used to assess how many ships may still be present that are not seen by the reporting systems. Ship detections from satellite images are sometimes referred to as VDS (for Vessel Detection System). The satellite images have only a very limited coverage in space and time, and therefore the non-reporting ship traffic cannot be monitored as was done for the reporting traffic, but only very sparsely sampled.

In this report, three types of analysis have been performed on the data:

- Assessment of the quality of the ship reporting data, in terms of completeness and accuracy, and the contributions of the various individual data sources to the final result;
- Construction of ship traffic density maps over the AOI on a monthly basis, to show and quantify the spatial distribution of the ship traffic;
- Assessment of the amount of non-reporting ships, the ships that are not seen in the AIS or LRIT, through the use of VDS / satellite images.

As mentioned, all the data have been used in the first place to compute the Maritime Situational Picture (MSP), the map of the positions of all the identified ship traffic in the AOI. The MSP was being displayed in real-time to the two operators in the region that were mentioned, with an update rate of once per 15 minutes. An example is shown in Figure 1-2 and Figure 1-3. The contents and functionalities of the PMAR Viewer are explained in detail in “PMAR Viewer User Manual, Issue Feb 2015”, JRC94785[1].

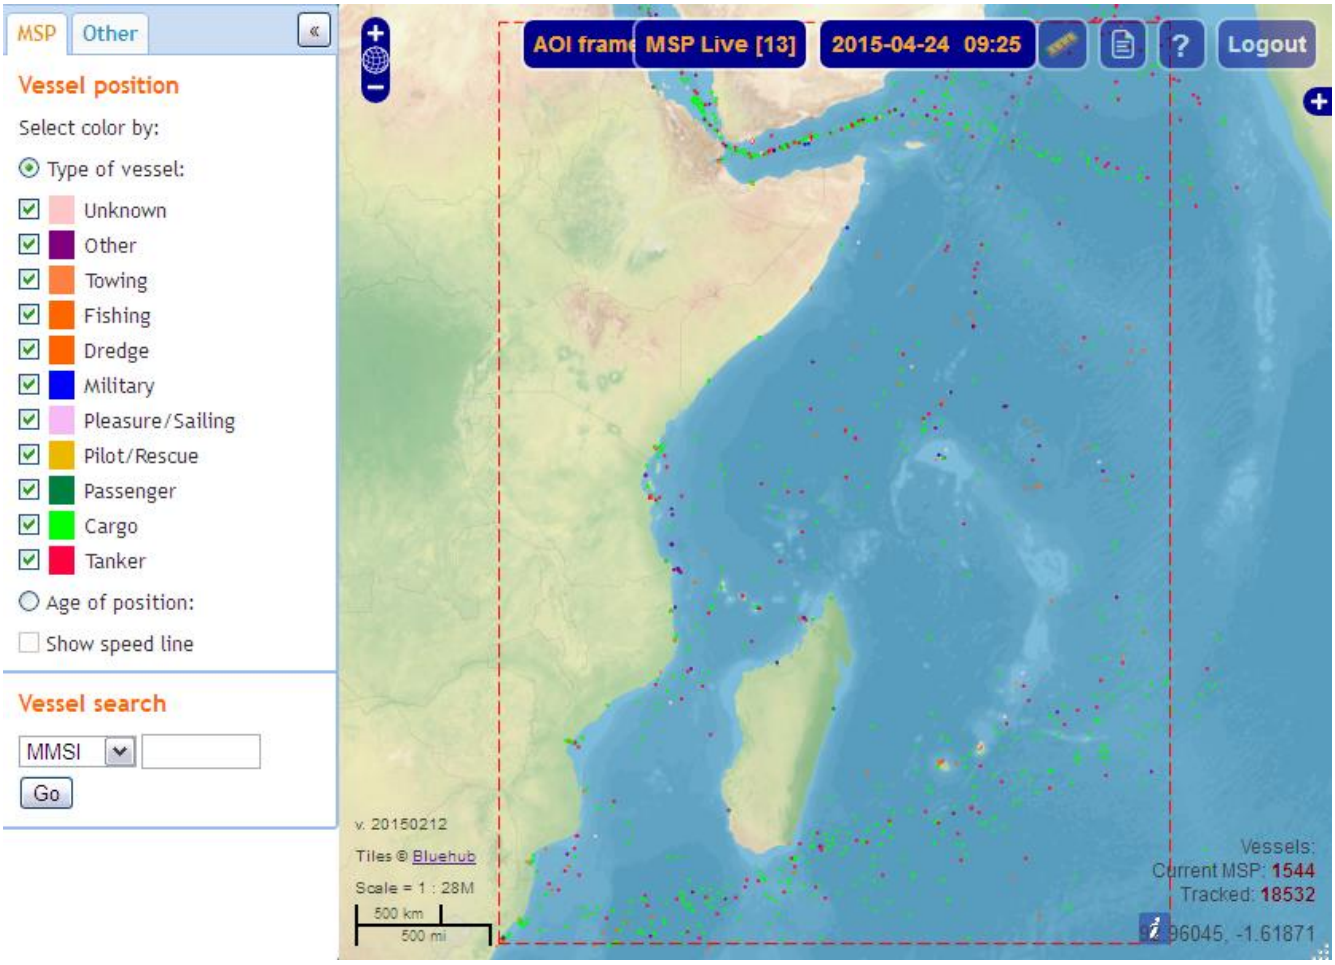

Figure 1-2. The PMAR-MASE Maritime Situational Picture, captured at 24 April 2015 09:25 UTC. Each reporting ship in the AOI (red box) is indicated by a dot, colour coded by ship type. Map background from Natural Earth.

[1] http://publications.jrc.ec.europa.eu/repository/bitstream/JRC94785/lbna27117enn.pdf

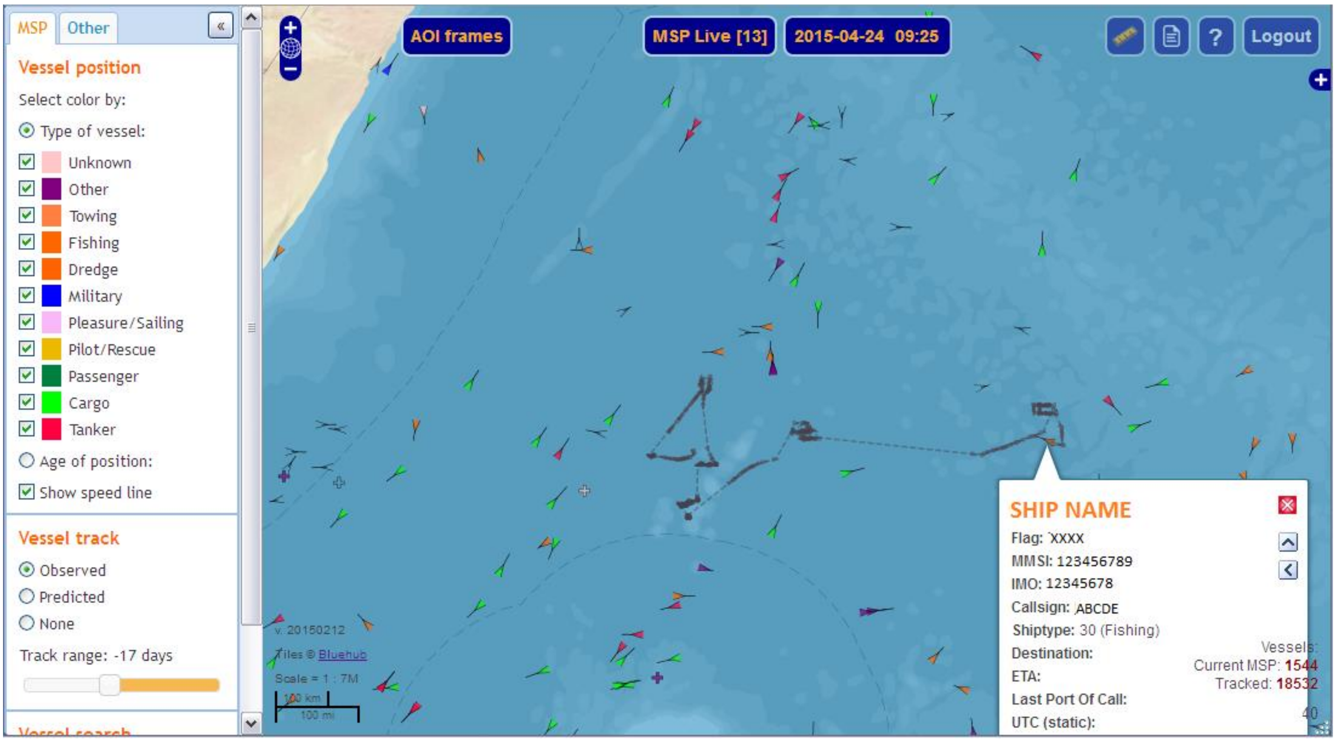

Figure 1-3. The PMAR-MASE Maritime Situational Picture zoomed in to show more detail. At this zoom level, the ship symbols are triangular arrows, with the short protruding line indicating the ship's expected displacement in the next hour. Any ship symbol can be selected to display its information and past track, as shown here. The erratic track of the selected ship confirms its type as 'Fishing'. On account of the public distribution of this report, the ship in question was anonymised. Map background from Natural Earth.

The set-up of the PMAR-MASE system (to collect and process the data) was discussed in the main PMAR-MASE Final Report, and is not repeated here in detail. Figure 1-4 shows the high-level system design. The "Data sources" on the left side of the figure are primarily the AIS and LRIT ship reports, but also piracy incident data, wind and wave predictions, and static map background layers.

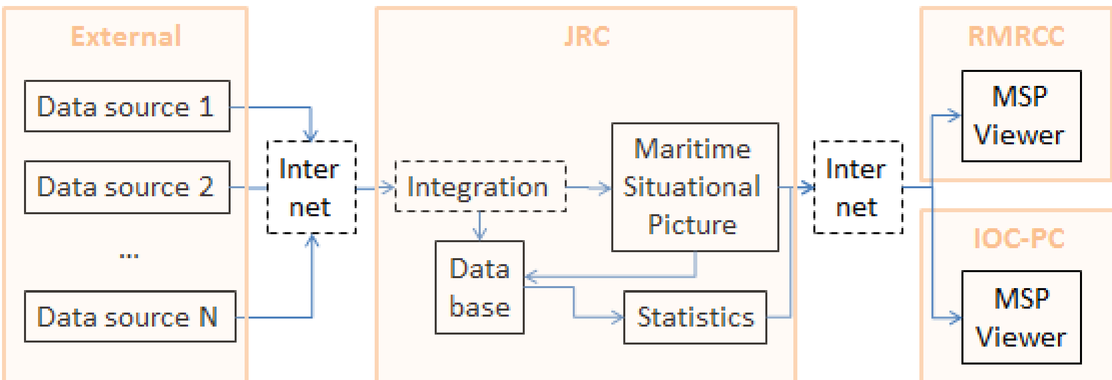

Figure 1-4. The PMAR-MASE high level system design. On the right are the two operational users, who could access the information as exemplified in the previous two figures via a web viewer.

2 Completeness and timeliness of AIS and LRIT

2.1 Introduction

This chapter analyses the completeness and accuracy of the information content in the AIS and LRIT data.

AIS and LRIT data that can be gathered from a certain AOI are fundamentally incomplete. This is in the first place because such information will only come from ships that carry those systems. But accepting that limitation, still only a fraction of the AIS and LRIT reports that are transmitted by the ships in the AOI will reach the user. For AIS, this is because terrestrial receivers cannot reach beyond the VHF horizon, so have only a local coverage; and satellite-based AIS receivers are only overhead of the AOI for a short amount of time, and even then miss many reports because the signal strength is extremely weak at the distance of the satellite, and too many messages may reach the receiver simultaneously which leads to saturation. In the case of LRIT, only a few reports do not reach their destination in the Flag administration, but there the issue is more to get the permission to access the data. Furthermore, LRIT under-samples the ship movements by design on account of its low update rate.

This chapter aims to quantify the information content of the maritime awareness that is based on AIS and LRIT data by two aspects. The first aspect is completeness. This is essentially quantified by counting how many different ships are seen in the data. The second aspect is timeliness. This relates to how often the ship positions are updated, and it is essentially quantified by counting how many position reports are received from the ships per unit of time, and how well they are equally distributed in time. Different measures with various implementations will be used to assess these two aspects.

Three essentially different data categories are distinguished in the ship reporting data: Satellite AIS, Terrestrial AIS and LRIT. The data sources that have been used for each category are:

- Satellite AIS:
  - Satellites operated by the Norwegian Coastal Administration
  - Satellites operated by the company exactEarth
  - Satellites operated by the companies Orbcomm/LuxSpace
  - Satellites operated by the company SpaceQuest
- Terrestrial AIS:
  - Coastal AIS network of MSSIS
  - Coastal AIS from the company exactEarth
  - Coastal AIS from the companies Orbcomm/LuxSpace
- LRIT:
  - LRIT from the Flags that use the EU LRIT Data Centre via EMSA

The satellite AIS data used in the project represents *all* data of that category coming from (semi-) operational systems that were in existence and known to us. No other AIS satellites exist, except for some experimental systems that only produce limited amounts of data.

Data from these sources have been received in the PMAR-MASE project for the period of one year, 15 Sep 2014 to 15 Sep 2015. For the analysis in the present chapter, the data from the month of June 2015 have been used. This month is representative for the rest of the year.

2.2 Types of position messages and their numbers

The AIS position message types are:

1, 2, 3: Class A position report

4: Base station report

11: Response to base station request (indistinguishable Class A or B)

18: Class B position report

19: Extended Class B position report

AIS message 27 (long range report) has not been used.

The position reports from the LRIT are only of one type and are identified in this study by the number 313.

The distribution of the received positions reports over these types is plotted in Figure 2-1. AIS messages 1, 2 and 3 have the same content. From Figure 2-1 it can be seen that by far most messages are AIS Class A position reports (1, 2, 3). The Class B position reports (18, 19) are much fewer, and the LRIT (313) even less again. Message 4 is not relevant for plotting ship positions, but it can be used to check the time stamping of a receiver.

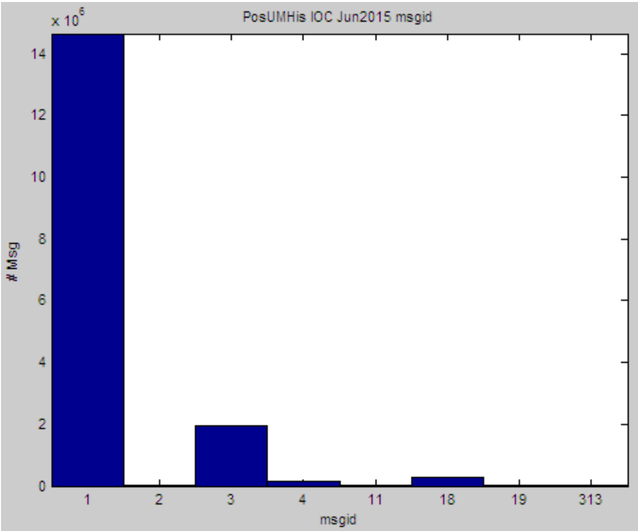

Figure 2-1. The number of messages received for each position message type. The numbers are for the whole month of June 2015 and the vertical scale is in units of one million messages.

2.3 Analysis per data category

Here we look at the number of position reports received and the different number of ships contained in them, split over satellite AIS, terrestrial AIS and LRIT.

Daily number of position messages and distribution over Class A/B

The number of messages received varies per day, because one day there may be more satellite overpasses over the AOI than another, and because the number of ships inside the AOI varies over time. The figure below (left side) depicts the variation, as well as the difference in message numbers received from LRIT, terrestrial AIS, and satellite AIS. The total number of messages received (on average 566,522 per day) is the sum of those three. It can be seen that satellite AIS provides the bulk of the position reports. On the right, the figure shows the percentage of AIS Class B messages. For the terrestrial AIS, this varies between 4 and 11 %; for the satellite AIS, it is much lower at around 0.5 %. There are two reasons for this. First, the Class B message is weaker, so it is less easily received by satellite. Second, Class B is carried by the smaller ships, which are found closer to the coast, within the coverage of the terrestrial AIS.

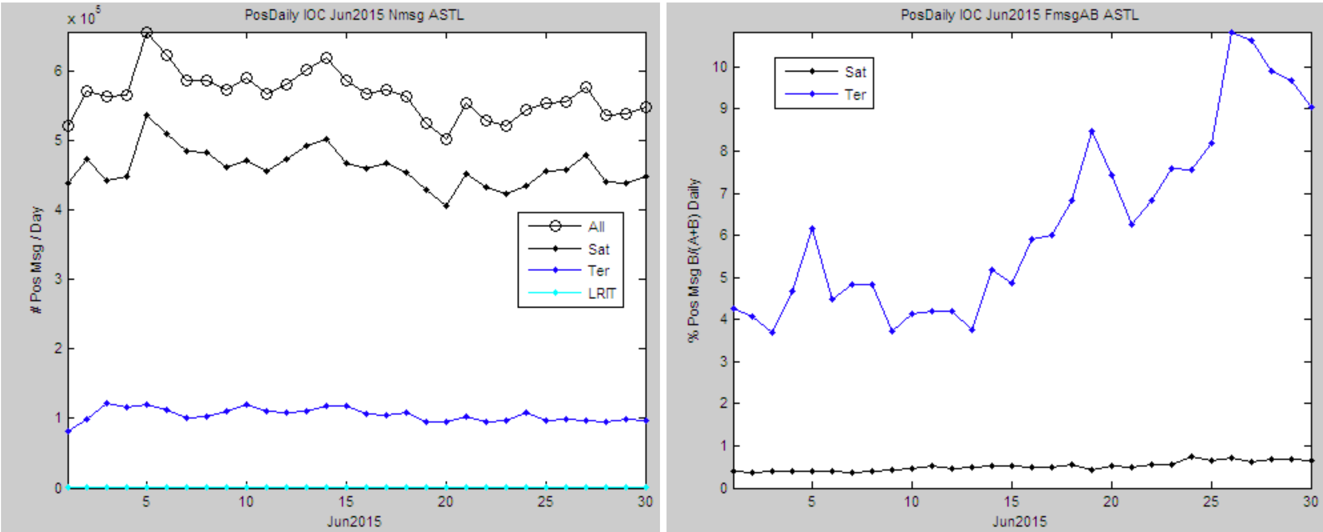

Figure 2-2. Left: The number of position messages received on each day of June 2015 from: LRIT (light blue, at the bottom); all terrestrial AIS sources combined (dark blue, around 1); all AIS satellites combined (black, around 4.5); all sources combined (black circles, top line). The vertical scale is in units of 100,000 messages. Right: The percentage of Class B position reports of the total Class A + Class B position reports received each day in June from: all AIS satellites combined (black, lower); all terrestrial AIS sources combined (blue, higher).

Daily number of different MMSI numbers seen in position reports

The number of unique MMSI numbers is close to the number of different ships in the AOI. In principle, a ship is uniquely defined by its MMSI number in the AIS message, and in the LRIT message by its IMO number which is linked to an MMSI number. Therefore, counting the different MMSI numbers that occur should count the different ships. However, when extracting the tracks of the ships, it can be recognized that some MMSI numbers are used by more than one ship. This happens in particular for ships that use wrong MMSI numbers, such as 0, 1, 12345, 111111111, and the like. The tracking shows that several tens of ships use duplicate MMSI numbers. On account of this effect, the number of different MMSI numbers would underestimate the number of different ships by that amount. However, another effect leads to an overestimation. A small number of AIS messages contain errors, that are introduced either in the transmission or in the reception of the VHF radio broadcast. Some of the errors occur in the MMSI value. This leads to MMSI values that are seen only one single time ever, or at least only once during a long period. If during one month an MMSI number is seen in only a single message, that message is most probably an error. (Even if it is not an error and there is a real ship of which only a single message is received, then that is not useable for the tracking of that ship.) Errors may also occur in the geographic position value; this is especially noticeable when ship positions are indicated on land. As the number of AIS messages per month is so high (in June, 566,522 x 30 = 17 million), a low error rate can still lead to a significant number of erroneous messages. It is found that many hundreds of messages in a month occur with an MMSI number that is only seen once in that month – in June, the number was 868. That corresponds to 29 per day on average. That effect overestimates the number of ships when counting the number of different MMSI. So both effects together, sharing of MMSI numbers and spurious MMSI numbers, in order of magnitude cancel each other. Therefore, the number of different MMSI numbers seen is used as a good indicator of the actual number of different reporting ships present.

The total number of different MMSI numbers seen each day in position reports is shown in the left figure below. It varies daily for the same reason as the number of messages discussed earlier. On average, the total number of MMSI numbers seen each day from all data sources together is 1,563. The figure shows that most of those ships are seen by satellite AIS. About half are seen by terrestrial AIS. (Note that, unlike for the number of messages, the number of MMSI numbers seen by the various sources does not add up to the number of MMSI numbers seen by all sources together.) Terrestrial AIS sees fewer ships because it has a much smaller spatial coverage. Only 12 % of the MMSI numbers are seen by LRIT. This is because only the ships that carry “EU Flags” are included in the LRIT, which is a small fraction of all ships in the AOI.

Figure 2-3 right side shows the percentage of Class B MMSI with respect to Class A + Class B. It is around 7 %, and not much different as seen from satellite or by terrestrial AIS. (LRIT does not cover Class B ships.) This should be compared with Figure 2-2 right side, where it was found that the number of Class B messages from terrestrial AIS was more than 10 times higher than from satellite. This reflects that fact that even though the satellite sees a similar number of Class B AIS ships as the terrestrial AIS, the Class B ships that are seen from the coast are much better tracked than those seen by satellite.

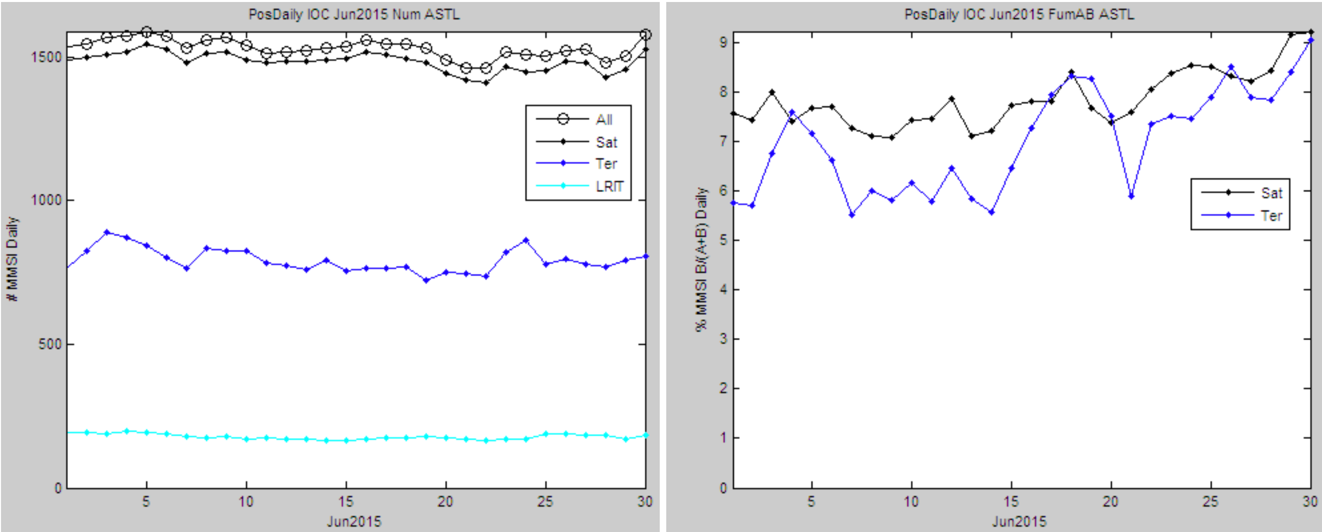

Figure 2-3. Left: The number of different ships seen each day in June 2015 by LRIT (light blue, bottom line); all terrestrial AIS (dark blue, second from bottom); all satellite AIS (black, second from top); and by all systems together (top, circles). Right: The fraction of Class B ships of the total number of ships (A+B) seen each day in June by all terrestrial AIS (blue, mostly the lower line); and by all satellite AIS (black).

The average daily number of 1,563 different MMSI numbers (corresponding to Figure 2-3) includes all that were seen in June (from all position message types, all data sources). When MMSI numbers that occur only once in the month of June, and position reports on land, i.e. likely erroneous ones, are excluded, the remaining average daily number of MMSI numbers is 1,532. This still includes MMSI numbers from AIS message type 4 (base station message), so MMSI numbers from base stations which are not relevant for the MSP. Excluding those, the remaining average daily number is 1,522.

Over the entire month of June 2015, the total number of different MMSI numbers seen is 6,217 (from all position message types, all data sources). Excluding MMSI numbers that are seen only once in that entire month, which are likely errors, the number reduces to 6,217 – 868 = 5,349. This is much higher than the average daily number, because ships move in and out of the AOI constantly, so new ships are seen every day. The average daily renewal rate is (5,349 – 1,532) / (30 – 1) / 1,532 = 8.5 %.

Geographical coverage

The AIS satellites and the LRIT cover the whole AOI uniformly. The terrestrial AIS on the other hand has only the coverage of the coastal ranges where the participating receiving stations are. Figure 2-4 plots the maps of the density of the received position messages, for all data sources together, and for satellite AIS, terrestrial AIS and LRIT separately. The limited geographic coverage of the terrestrial AIS is obvious, as is the much lower density of the LRIT. The density map of all sources together is dominated by the satellite AIS.

The density maps have a very high dynamic range – in some places, e.g. in or near ports, the number of messages per area can be very high, while in others, e.g. on the open ocean outside the main traffic routes, it can be very low. For a meaningful visualisation, the density colour scale is therefore logarithmic. The scale in the top left map ranges from 0 to 4.1, which means from 10^0 = 1 message per unit surface area to 10^4.1 = 13,000 messages per unit surface area. The greenish parts in the middle of the scale have 10^2 = 100 messages per unit surface area.

See the next chapter for all monthly density maps and further details on their construction and contents.

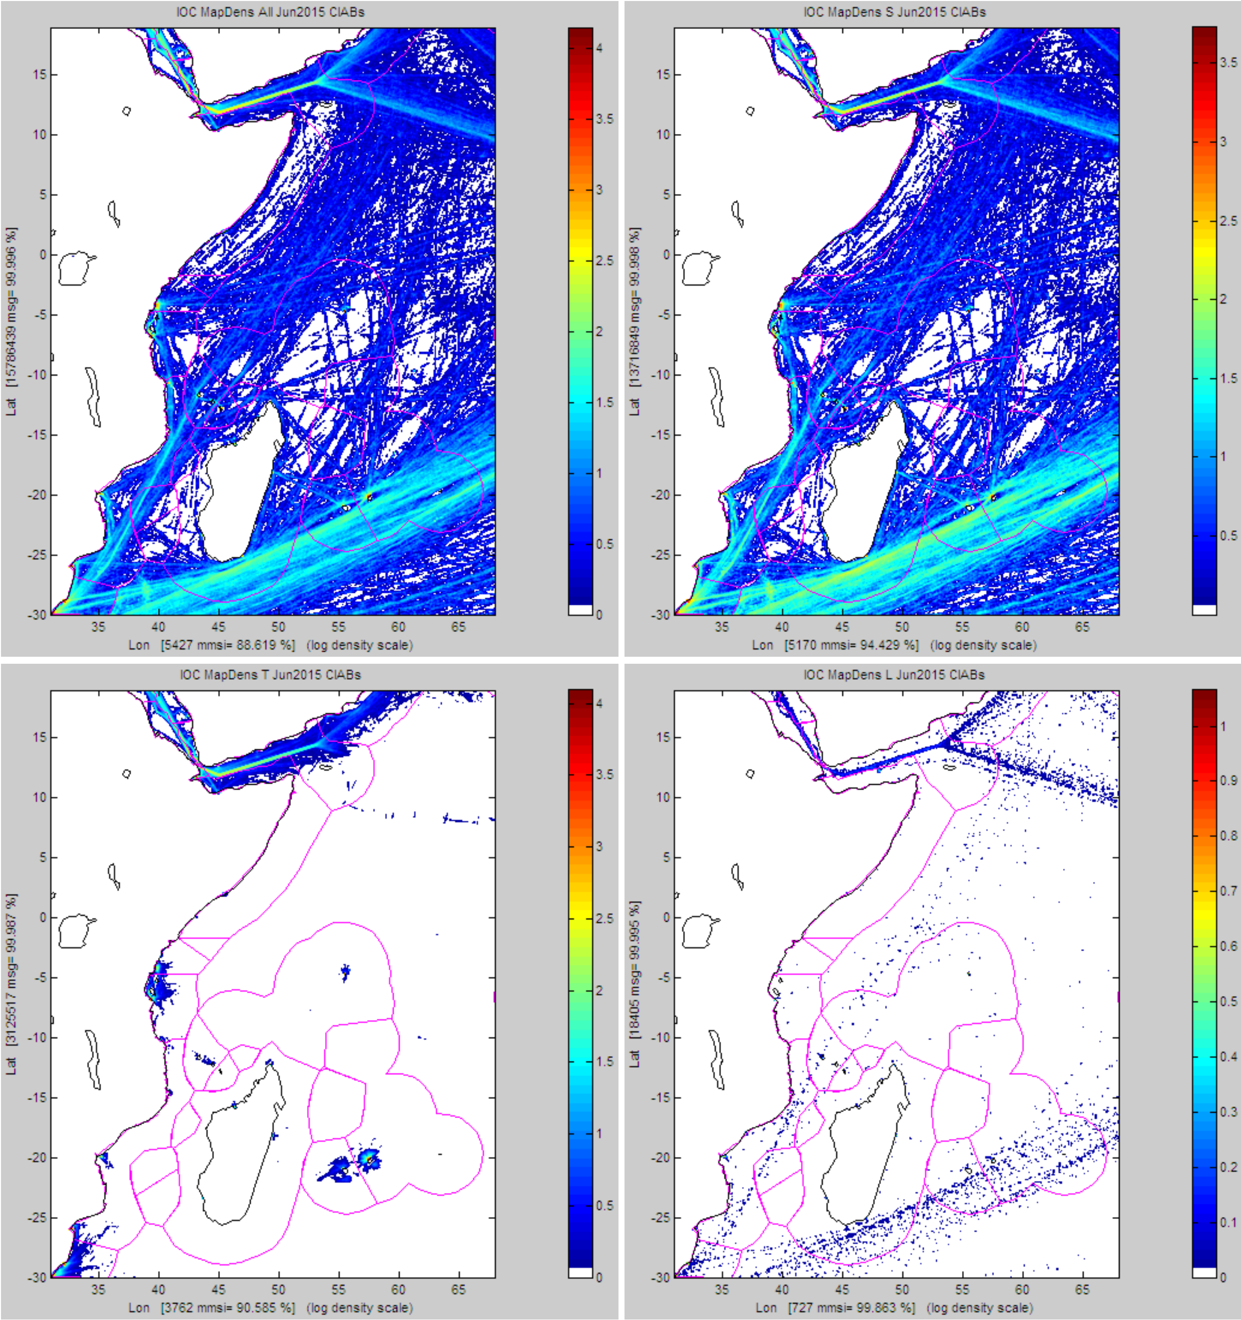

Figure 2-4. Density maps of all ship positions received in June 2015. Top left: from all data sources combined. Top right: from all AIS satellites. Bottom left: from all terrestrial AIS. Bottom right: from LRIT. The colour scaling is logarithmic.

2.4 Daily number of position reports and MMSI numbers

This section documents the number of position reports that have been received, and the number of distinct MMSI numbers that are present in the data. It does so per data type: satellite AIS, terrestrial AIS, LRIT, and all combined. It gives a number for each day between 1 Sep 2014 and 30 Sep 2015, in the form of monthly graphs. MMSI numbers that are seen only once ever in a certain month are not counted, because they are most likely due to data errors. Also messages and MMSI numbers of AIS message type 4 (base station) are not counted. Therefore, the numbers should accurately reflect the ship-borne MMSIs.

The following pages show the graphs for each month, the number of messages on the left and the number of MMSI numbers on the right. All have the same scale, up to 700,000 messages and 1,700 MMSI numbers respectively.

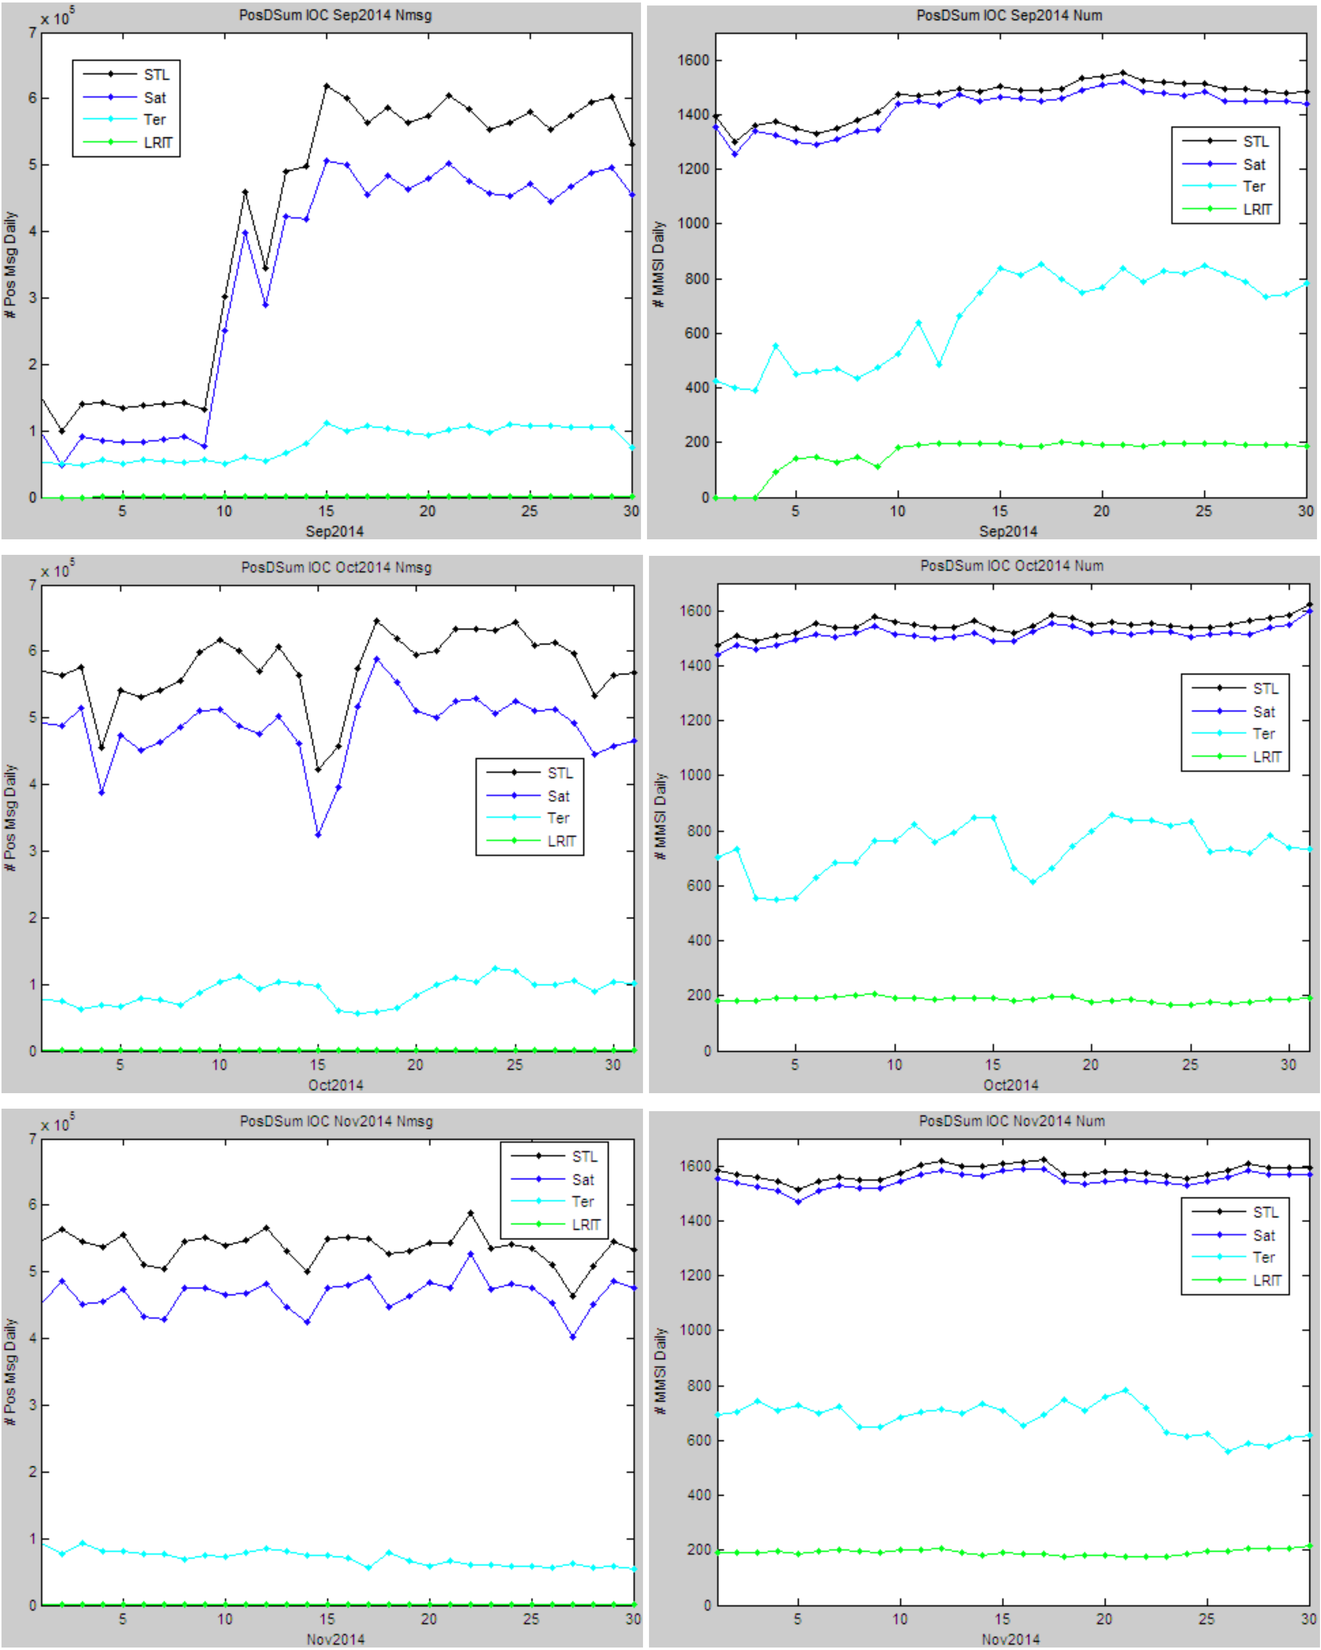

Figure 2-5. Daily number of position reports received (left) and unique MMSI numbers seen (right). Separate curves for LRIT (green, lowest), terrestrial AIS (light blue, second from below), satellite AIS (dark blue, second from top) and total over those three (black, topmost curve). The months of September, October and November 2014.

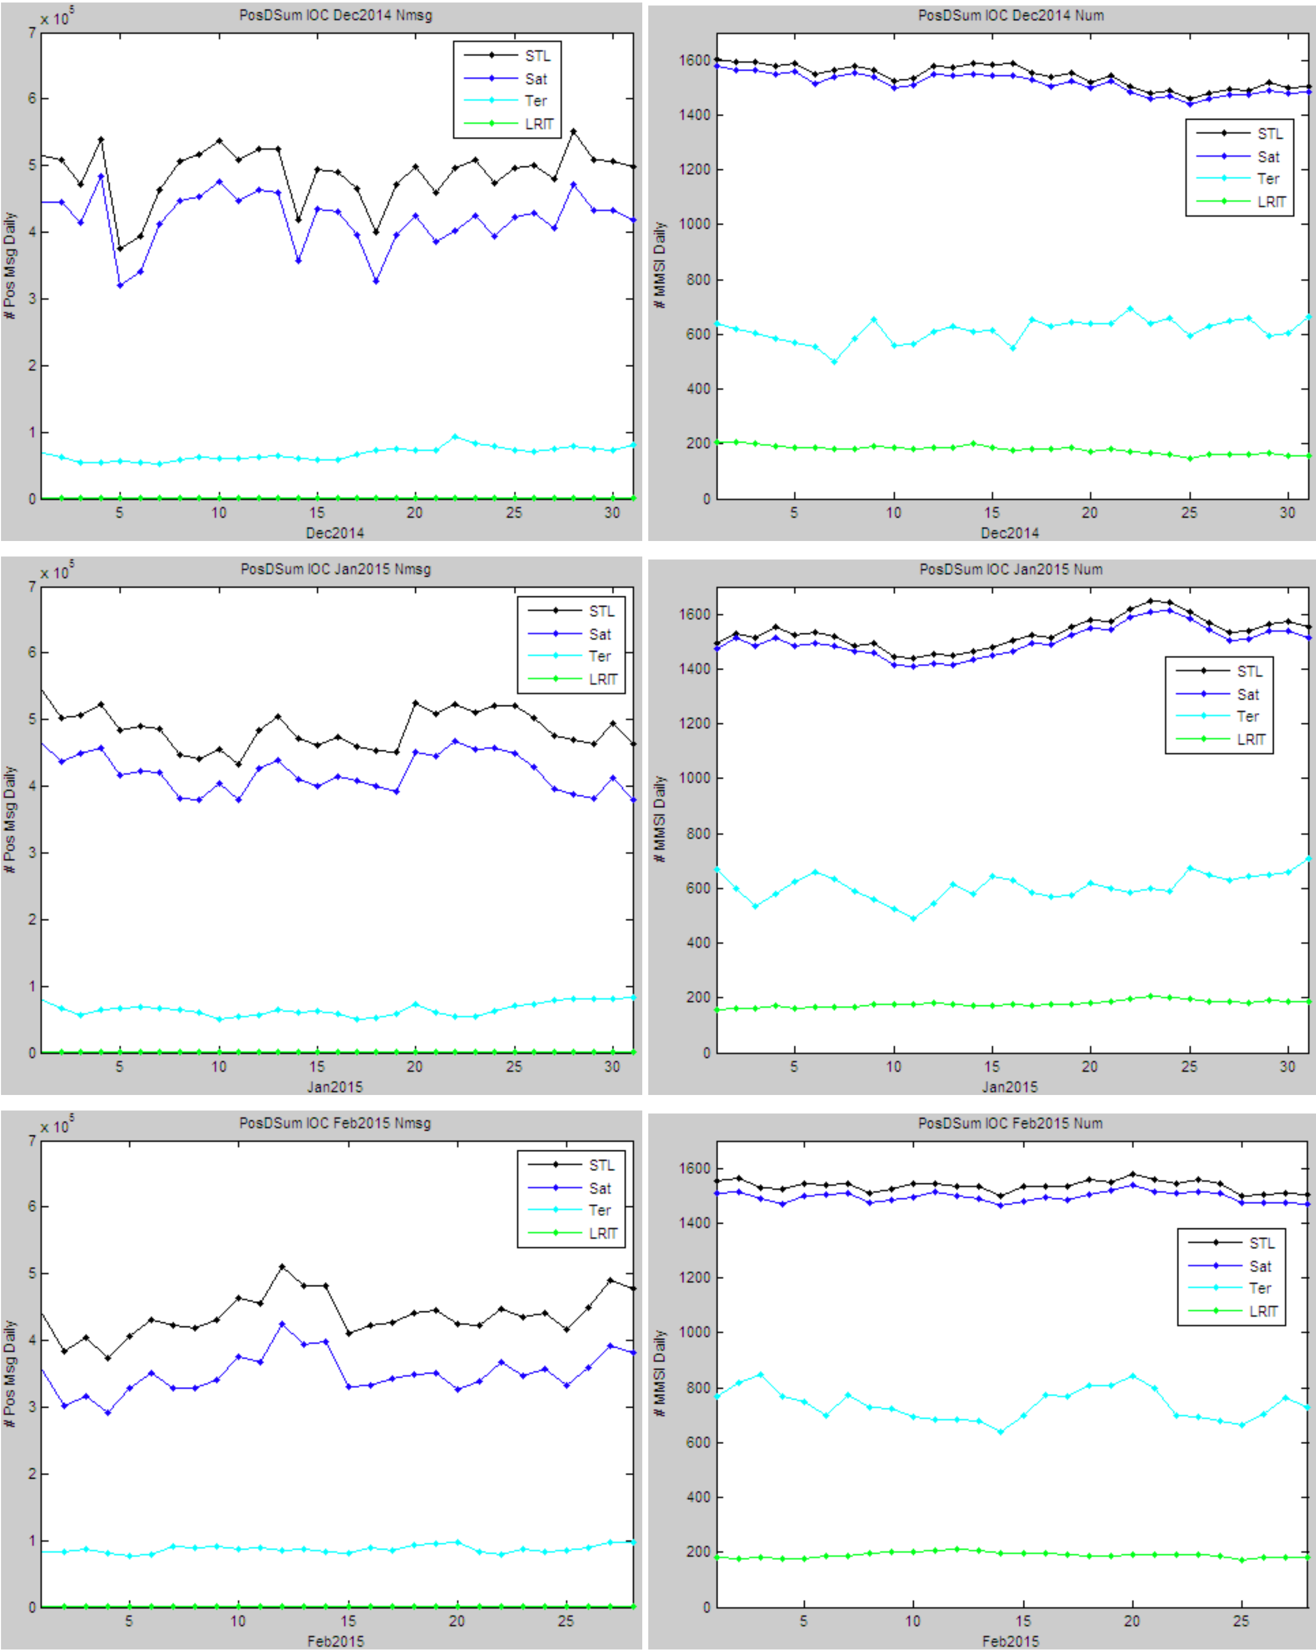

Figure 2-6. Daily number of position reports received (left) and unique MMSI numbers seen (right). The months of December 2014, January 2015 and February 2015.

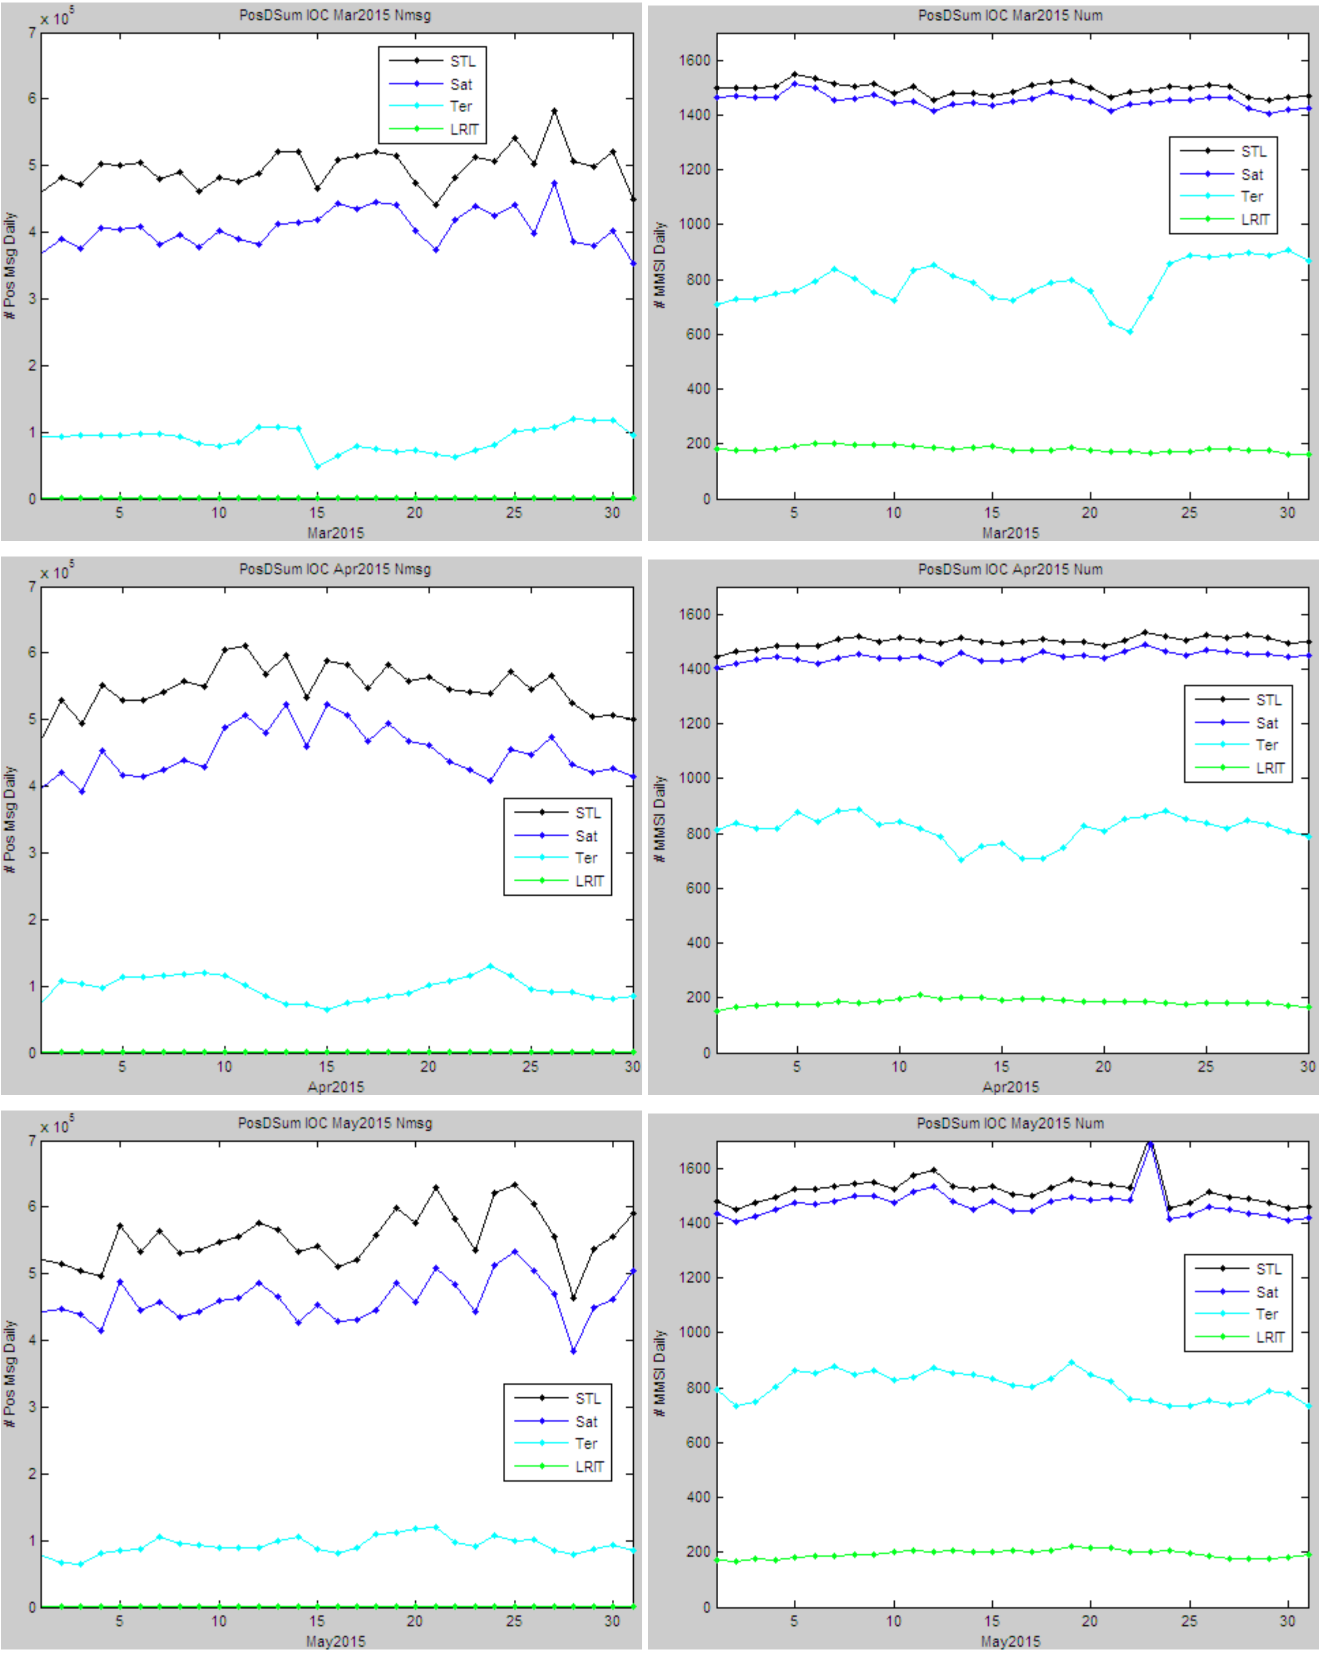

Figure 2-7. Daily number of position reports received (left) and unique MMSI numbers seen (right). The months of March, April and May 2015.

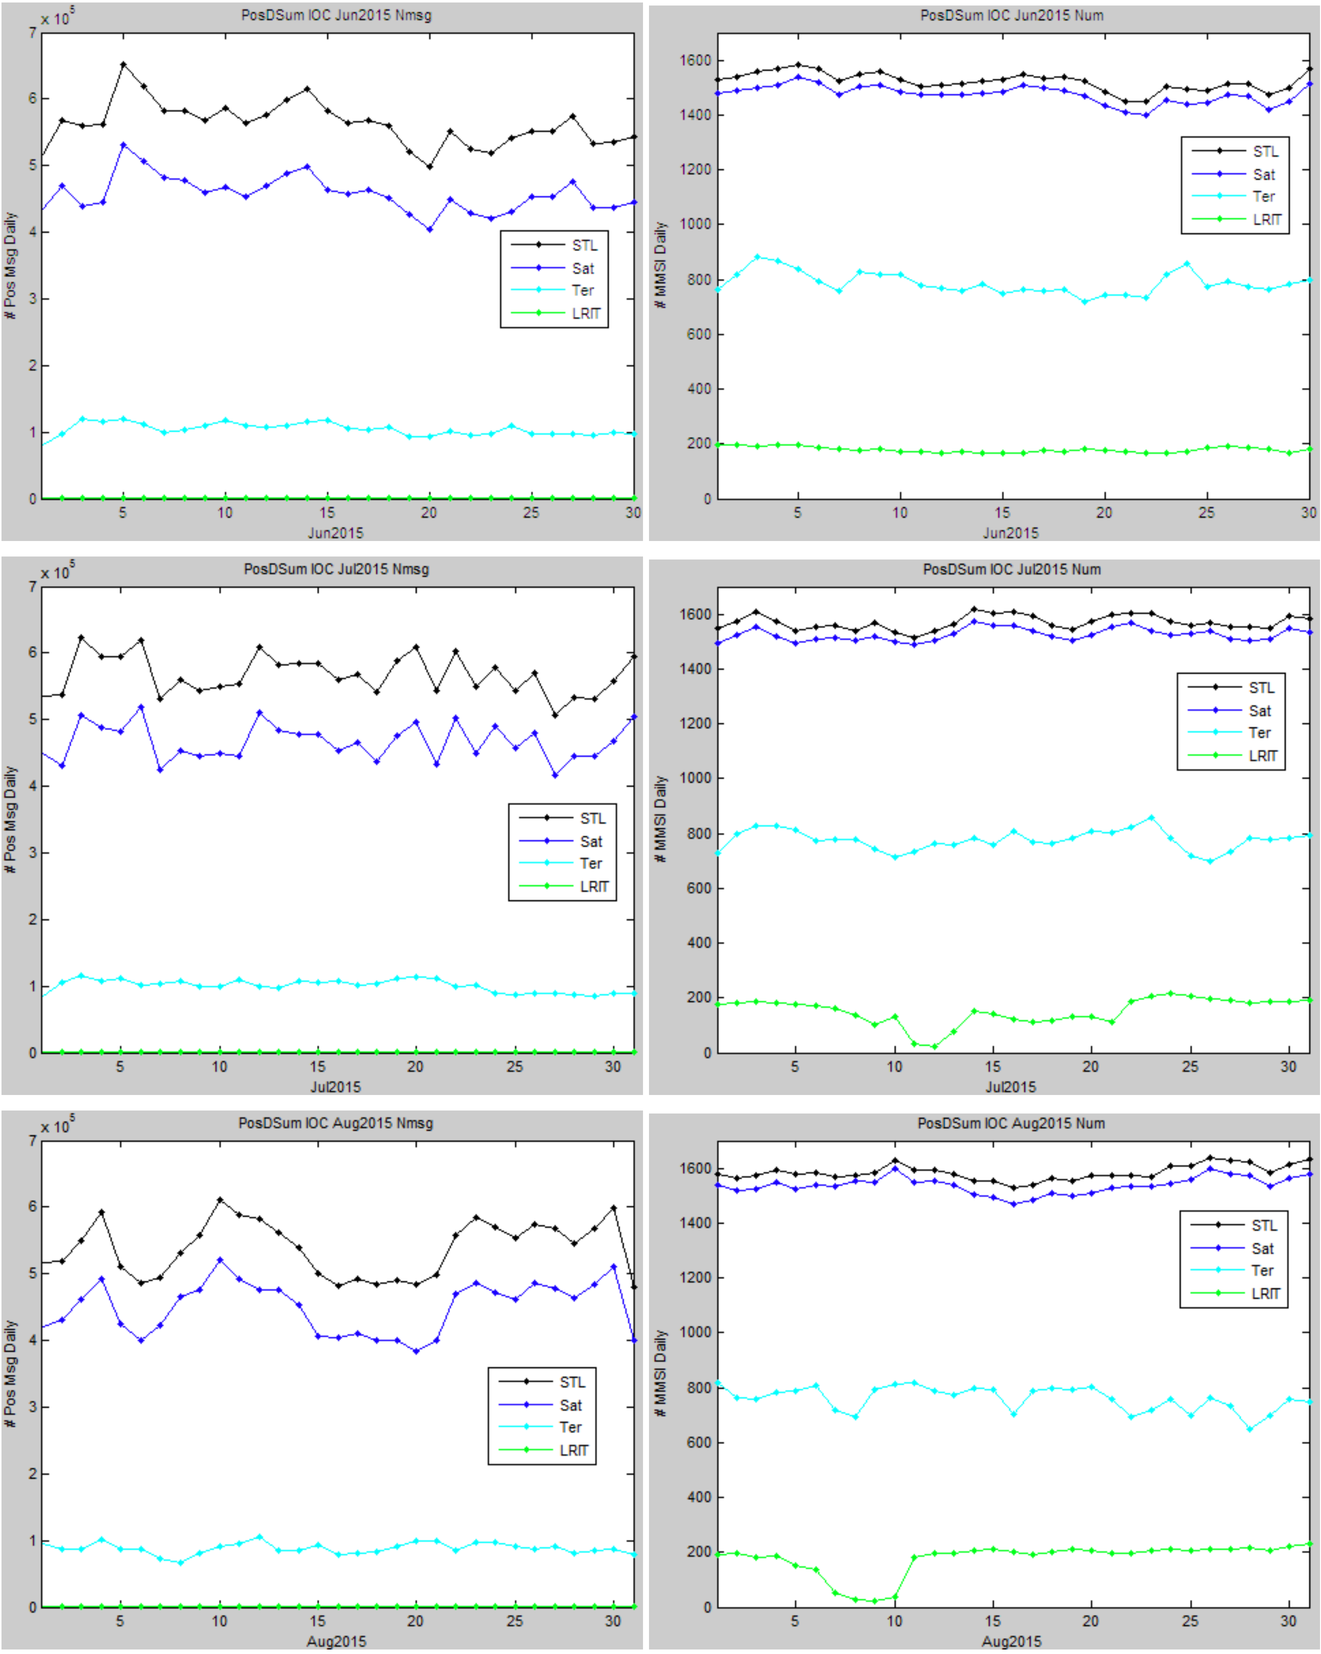

Figure 2-8. Daily number of position reports received (left) and unique MMSI numbers seen (right). The months of June, July and August 2015.

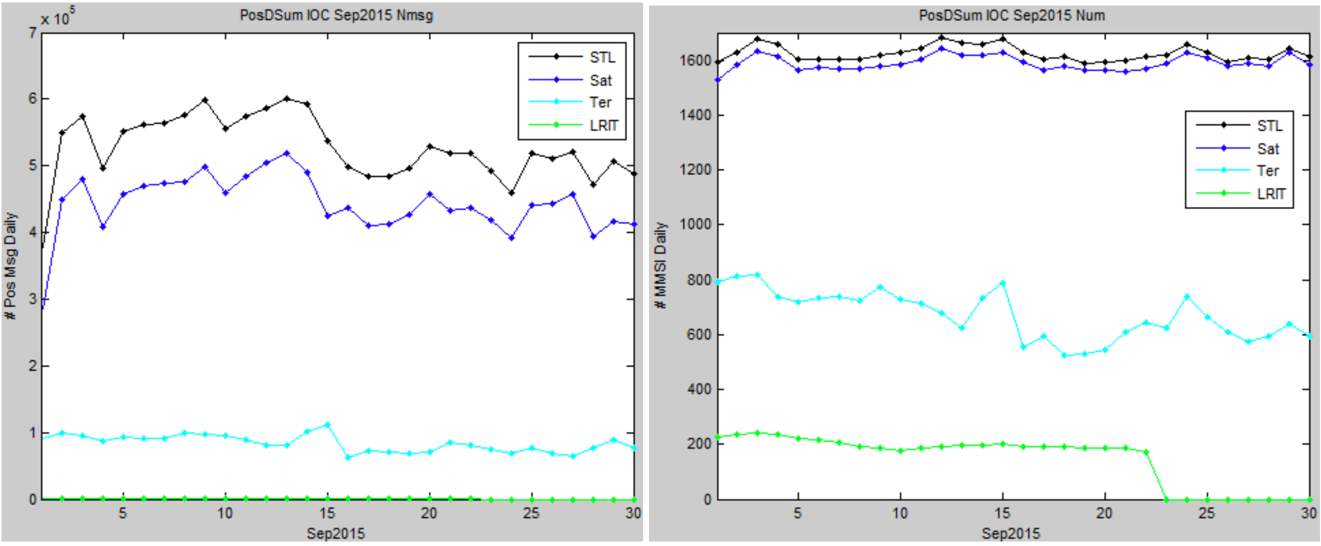

Figure 2-9. Daily number of position reports received (left) and unique MMSI numbers seen (right). The month of September 2015.

2.5 The contributions per data source

This section will look at the contributions of individual data sources, separately and in combinations.

Separate contributions of individual platforms

In the month June 2015, there were 11 distinguishable AIS satellite data sources available. (In other months the number differed.) One of these 11 is Orbcomm Second Generation (OG2) which consists of 6 separate satellites. So in total there is data from 16 satellites included. However, in the data stream, the six satellites of OG2 are not distinguished. Therefore, in this analysis, they count as one.

Figure 2-10 plots the daily number of MMSI numbers seen by each of the 11 satellites separately (drawn lines) and by each terrestrial AIS network and LRIT separately (dotted lines). All satellites have a broadly similar performance in the range of 1,150 to 1,350 MMSI numbers, but still systematic differences can be seen. Sometimes individual satellites suffer from outages, such as shown by the dips on 13 and 17 June. As long as there are several other satellites available, that does not lead to big problems in the surveillance. The differences between the coastal networks are bigger; the larger coastal networks have a similar performance of around 500 MMSI numbers, but also some networks are included with a much smaller geographic distribution, leading to the dotted lines near the bottom. The LRIT can be recognized in the purple line (around 180 MMSI numbers), the same line that was light blue in the previous graph of Figure 2-3 left.

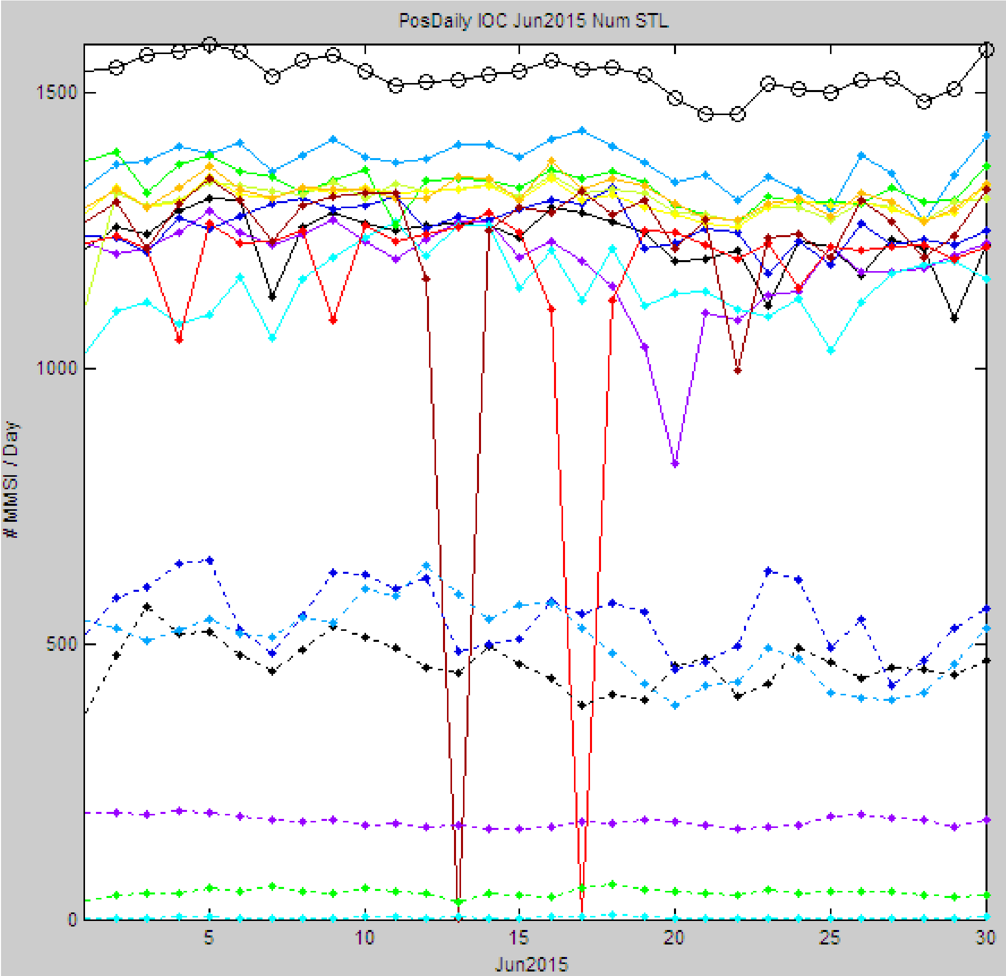

Figure 2-10. The daily number of MMSI numbers seen for each data source separately. The drawn (higher) lines are the AIS satellites. The dotted (lower) lines are the terrestrial AIS networks and LRIT. The top line (circles) is all data sources combined.

Figure 2-11 left shows again the number of MMSI numbers seen each day by the individual AIS satellites, but now as a percentage of the total number of MMSI numbers seen each day by all AIS satellites combined (so excluding the terrestrial AIS and LRIT). So it shows the performance of each AIS satellite with respect to the others. The right side of that figure shows the percentage of Class B MMSI numbers seen by each satellite. This varies significantly from around 2 % to around 6 %. As the satellites have similar geographic coverage, and the Class B signals being weaker than the Class A ones, the difference is ascribed to the difference in sensitivity of the satellite receivers.

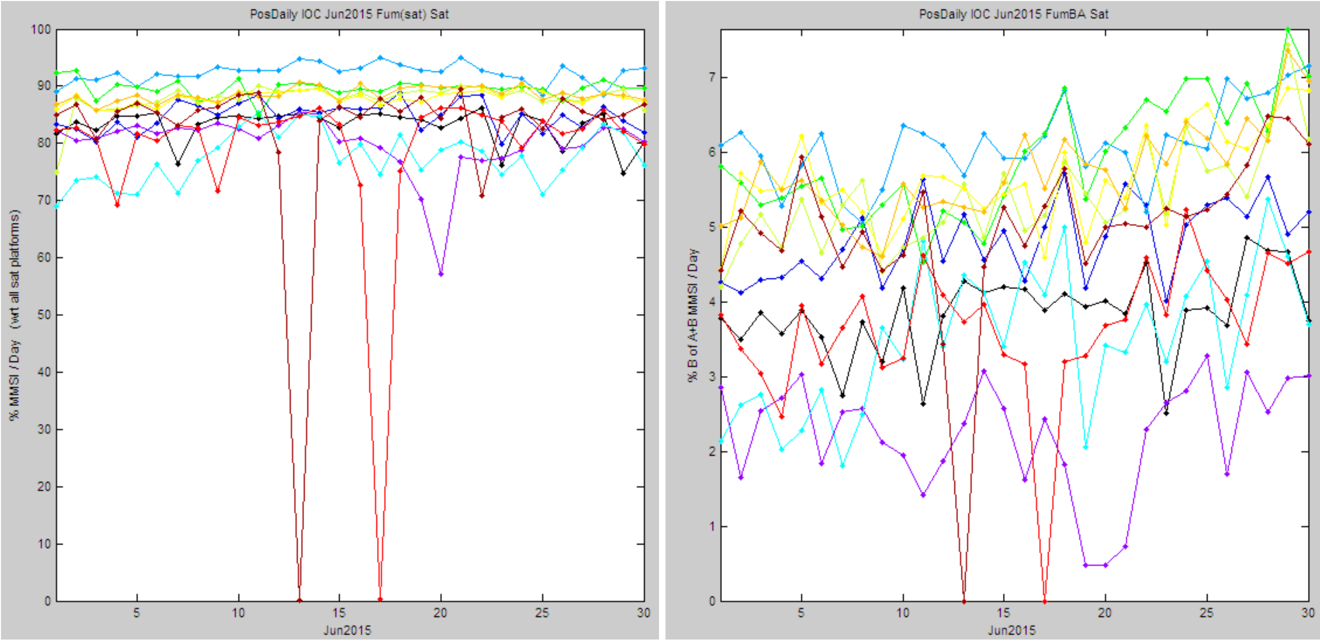

Figure 2-11. Left: The number of MMSI numbers seen by the individual AIS satellites on each day in June 2015, as a percentage of the number of MMSI numbers seen by all AIS satellites together on that day. Right: The percentage of Class B MMSI numbers seen each day for each satellite separately.

Cumulative contributions of AIS satellites

With 11 satellite platforms, it is possible to quantify the performance of one platform on average, by making the average of the performance of each of the 11 platforms. (Performance as before in terms of the number of different MMSI numbers seen per day.) It is also possible to measure the performance of a constellation of two satellite platforms, by calculating the average over all possible combinations of two platforms out of 11. So we calculate the performance of the combination of platforms 1 and 2, of the combination 1 and 3, 1 and 4, etc., up to 1 and 11; then the combination 2 and 3, 2 and 4; etc. There is a total of 11 x 10 / 2 = 55 possible combinations of two platforms out of 11. The performances of all these 55 combinations are calculated each day, and then averaged to obtain the average performance of a constellation of two satellites on each day. This is also done for a constellation of 3 satellites out of the 11, 4 out of 11, 5, etc.; up to 11 out of 11 which is of course only one possible combination. In general, the number of possible combinations of k out of n is n! / ((n-k)! k!). This becomes large for the mid ranges of k, like for n=11 and k=6 it is 462. The total number of combinations to calculate, from k = 1, 2, 3, … to 11 satellites, would be 2047. To avoid excessive computing time, the number of computed combinations for each k is limited to a random selection of maximally 55 different combinations out of the possible total. This leads to an overall total of 463 combinations to calculate.

The result is displayed in Figure 2-12, which shows the number of different MMSI numbers seen each day in June by the average combination of k AIS satellites out of the total available 11, for k = 1, 2, 3, …, 11. The number is expressed as a percentage of the total number of MMSI numbers seen that day from all data sources including terrestrial and LRIT. In these calculations, MMSI numbers that are only seen once ever in the whole month, position reports on land, and AIS message type 4 (base stations) have been excluded. The bottom (black) curve for k = 1 corresponds to the average of the 11 individual satellite curves of Figure 2-11 left (although at a different scale). The top curve (brown) is the result of all 11 satellites combined, so corresponds to the 'Sat' curve in Figure 2-3 left (although again at a different scale). It can be seen that the benefit of adding satellites diminishes with each satellite added. It can also be seen that already a combination of 2 satellites is less sensitive to individual satellite outages: the curve of k = 2 is less variable than the curve of k = 1, which shows dips at 13 and 17 June corresponding to the individual satellite outages on those days.

The diminishing return behaviour is summarized in Figure 2-13 which shows how the performance of a constellation of satellites grows with the number of satellites. The figure plots the mean values of the 11 curves of Figure 2-12 averaged over the whole month. The graph grows quickly in the beginning but flattens off after 3-4 satellites.

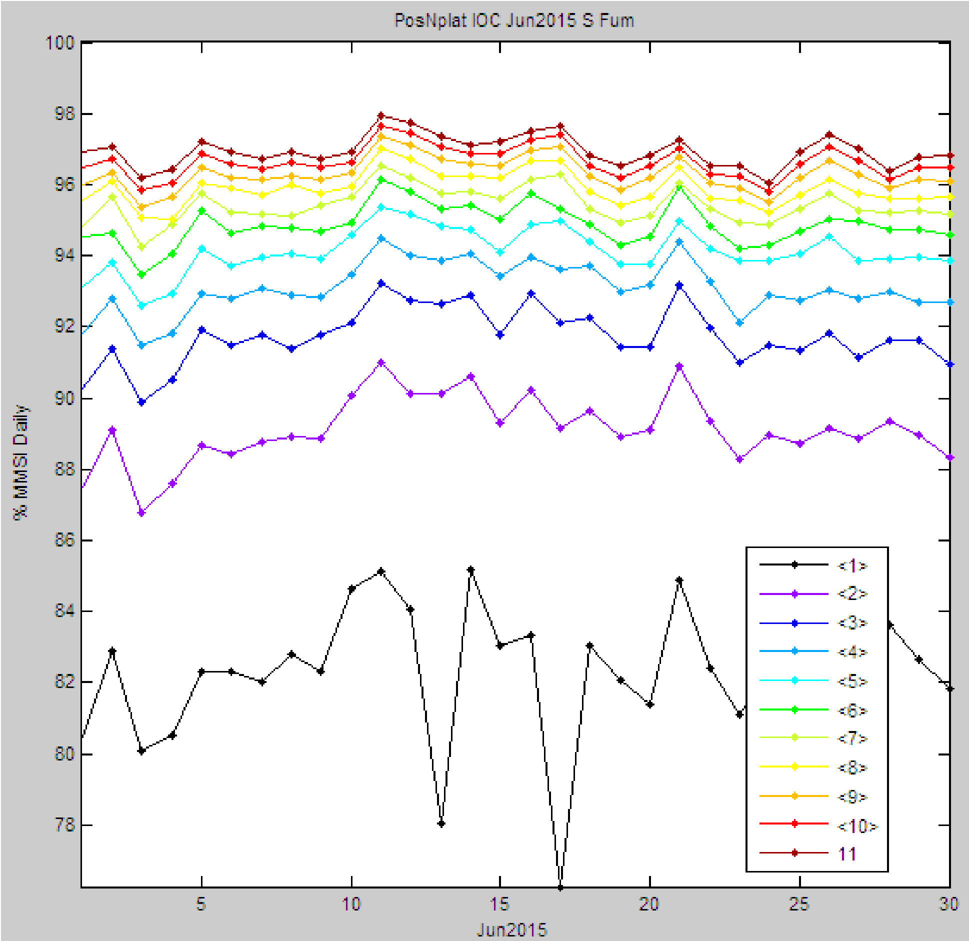

Figure 2-12. The number of different MMSI numbers seen each day in June on average by a combination of k AIS satellites out of the total available 11, for k = 1, 2, 3, ..., 11. The number is expressed as a percentage of the total number of MMSI numbers seen that day from all data sources including terrestrial and LRIT.

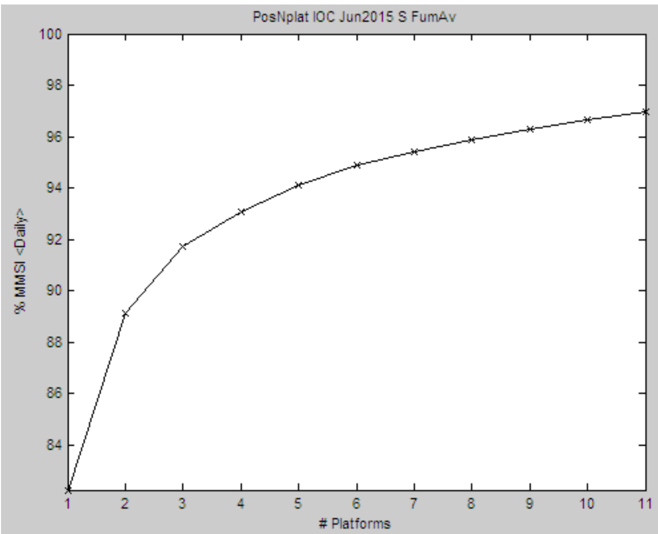

Figure 2-13. The average daily percentage of MMSI numbers seen over the month of June 2015 by the average combination of 1 AIS satellite out of 11, 2 out of 11, etc., up to all 11.

Additional contribution of terrestrial AIS

In order to quantify the contribution of the terrestrial AIS over that of the satellites, we count how many MMSI numbers are seen extra when adding terrestrial AIS to that of k satellites, for k = 1 to 11. The result is shown in Figure 2-14. The top line of that figure shows that on average around 160 MMSI numbers are seen each day in terrestrial AIS that were not also seen in satellite AIS, when using only one satellite; the points on the line are the average for all 11 satellites, like the black line was in Figure 2-12. That is 160 / 1,522 = 10.5 %. (As in the previous section, the MMSI numbers here exclude those that were only seen once in the whole month, position reports on land, and AIS message 4.) The second line from the top (purple) shows that when using a constellation of two AIS satellites, the terrestrial AIS only adds around 105 MMSI numbers on average; each point on this line is computed by averaging over all possible combinations of two satellites out of 11, like the purple line was in Figure 2-12. Etcetera. The lowest line shows that on average, terrestrial AIS adds about 39 MMSI numbers to those seen by the AIS from a constellation of 11 satellites; that corresponds to 39 / 1,522 = 2.6 %.

The diminishing additional value of terrestrial AIS over satellite AIS with the number of satellites is visualised in Figure 2-16 left.

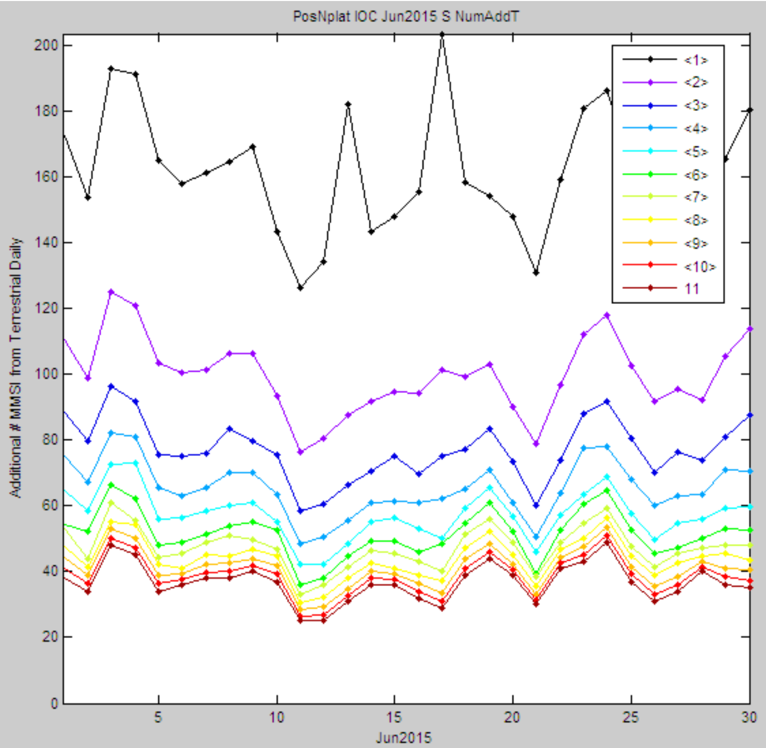

Figure 2-14. Each line (colour) shows for each day in June 2015, the number of MMSI numbers seen by terrestrial AIS that were not already seen by satellite AIS on that day. The top (black) line is when using one satellite; the second from top (purple) is for two satellites; etc. down to the lowest line (brown) that is for using all 11 satellites.

Additional contribution of LRIT

Just as for the terrestrial AIS, Figure 2-15 shows the additional number of MMSI numbers seen by LRIT over those seen by satellite AIS, for k = 1 to 11 satellites, averaged over all combinations of k satellites out of 11, for each day in June 2015. For LRIT, the use of one single satellite leaves on average about 27 MMSI numbers (1.8 %) unseen; whereas that number is reduced to about 11 MMSI numbers (0.7 %) when using 11 AIS satellites. The diminishing additional number of MMSI numbers seen exclusively in LRIT as a function of the number of AIS satellites is plotted in Figure 2-16 right.

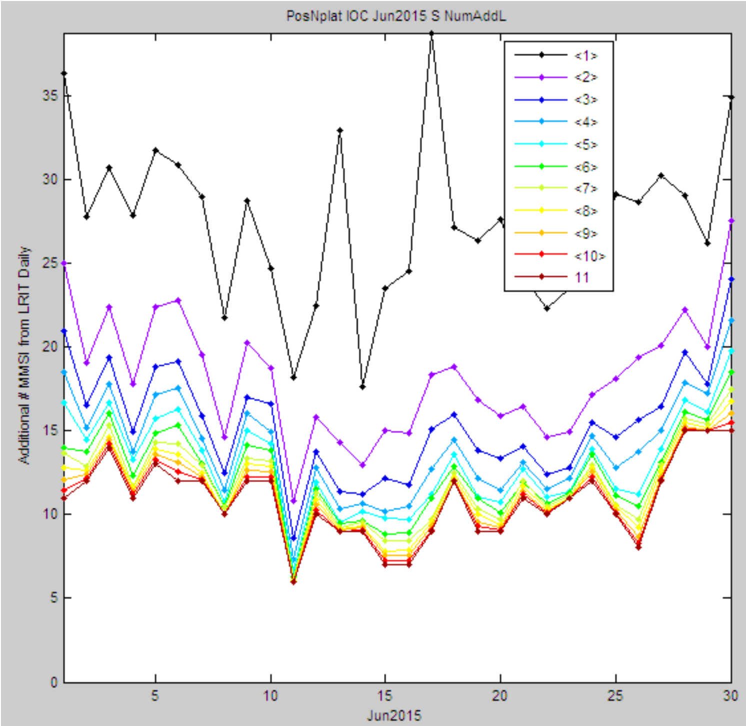

Figure 2-15. Each line (colour) shows for each day in June 2015, the number of MMSI numbers seen by LRIT that were not already seen by satellite AIS on that day. The top (black) line is when using one satellite; the second from top (purple) is for two satellites; etc. down to the lowest line (brown) that is for using all 11 satellites.

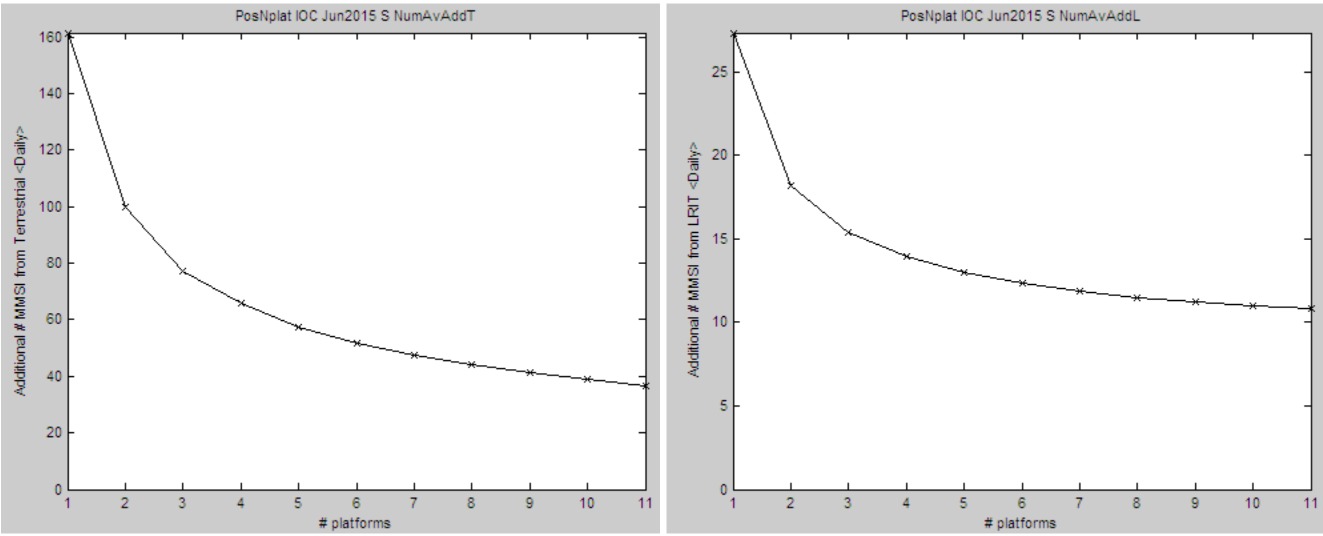

Figure 2-16. The number of MMSI numbers seen exclusively in terrestrial AIS (left) / LRIT (right), on a daily basis averaged over the month of June 2015, over and above that seen on satellite AIS, as a function of the number of AIS satellites used (1 to 11).

2.6 Completeness and timeliness of the MSP

Tracking

The previous sections have dealt with numbers of position messages and numbers of ships seen in the data per time period, as a measure of the completeness of the data. In order to make the MSP and to make unbiased ship density maps, and in general to use the data for maritime awareness, the ships in the AOI need to be tracked. This means that the position reports of a certain MMSI number need to be collated, ordered in time, and the resulting irregularly sampled series of positions must be interpolated to a regular time grid (e.g. every 15 min) to obtain the ship's track. That process, however, is not straightforward because the data contain errors and because sometimes, as mentioned, several ships share the same MMSI number.

The latter situation can be recognised by the tracking algorithm based on the fact that ships cannot move faster than a certain maximum speed – for sure not faster than e.g. 100 knots. If two consecutive positions of the same MMSI number are so far apart that a higher speed would be needed to connect them, they must belong to two different ships. With this principle, two ship tracks under the same MMSI number can be separated if they are sufficiently far apart in space, and sufficiently densely sampled in time. Even several (3, 4, 5, ….) tracks can be disentangled from the same MMSI number in this way, provided they are separated well enough. However, if the two ships are too close together or there are long gaps in the position reports, the outcome of this track disentanglement becomes less reliable.

Errors in the data can occur in any field of the message, in particular in the MMSI field, the longitude or latitude field, and the time stamp. These errors give rise to outlier points in the track, points that lie away from all the other points. The tracking algorithm has to recognise these and discard them. If it concerns single, isolated outliers, this is not too difficult. However, there are two sources of error that cause serious problems. The first is related to the time stamp; the second to faulty AIS equipment. The time stamp is, for the AIS message, not inside the message, but instead it is affixed to the message by the receiver upon reception. The receiver uses its own clock for that. It is found that some receivers have unstable or offset clocks. When the track of an MMSI number is compiled from several data sources, messages from a receiver with a clock offset are shifted with respect to the other messages, which leads to a jumping around of the track. This can further lead to a spurious splitting of the MMSI over two tracks. Concerning faulty AIS equipment, some MMSI numbers display positions that have a very high error rate, often in a systematic way, such as e.g. dropping the sign or one digit in the latitude. Such types of errors are more difficult to recognise automatically and also give rise to spuriously split tracks (more tracks under the same MMSI number).

The MSP is based on the tracks, because the MSP plots the ship positions at regular time intervals, as interpolated from their irregular position reports. In the case of the real-time MSP, obviously only past position reports are used to construct the track, followed by a prediction (extrapolation) step to get the MSP at the current reference time. However, when using historic data, the MSP ship positions can be interpolated if there are position reports on both sides of the time axis available. This reduces the uncertainty of the prediction. Nonetheless, if the gap between two position reports that bracket the desired MSP time is too large, interpolation is not the best option. For example, to calculate an MSP position 5 min after a position report when the next position report is 24 hours away, it is more reliable to extrapolate the previous position reports.

For the current evaluation, the ships tracks (and consequently the MSPs) have been calculated along the lines above:

- Make ordered time sequence of positions of a certain MMSI number;
- Remove outlier points;
- Check, by using a maximum speed criterion, if the track can be ascribed to a single vessel or should be split over multiple vessels;
- Interpolate the regularly time-sampled MSP positions from the series of irregularly time-sampled position reports of one track;
- However, instead of interpolate, extrapolate (forward or backward) in case the time gap between two position reports is too large.

This method leaves the choice of a number of parameters. In the present calculation, these parameters were set to:

- Time allowance for unsynchronised clocks from different receivers: 60 seconds;
- Maximum speed to allow in one track (else split into several tracks): 35 knots;
- Maximum time gap to interpolate (else extrapolate): 18 hours;
- Maximum extrapolation time (backward and forward): 9 hours;
- MSP sampling interval: 15 minutes.

MMSI numbers of which only a single position report was received by all data sources combined, during the entire month, have been omitted (as these are ascribed to errors). However, when constructing the tracking from a subset of data sources, it can be that a certain MMSI only has one single position report, or a few reports in one small time interval. Such vessels are present in the MSP for 18 hours, on account of the maximum extrapolation time of 9 hours used. The 9 hours are chosen as a compromise between on the one side bridging gaps in position reports of ship that move around in the AOI, as they are actually observed to occur, and on the other side preventing undue extrapolation of the presence of a ship that has turned off its AIS transmitted e.g. because it went into a port. Ships that are seen to exit the AOI are not tracked anymore, as are ships of which the positions are extrapolated on land.

Indicators for the MSP

The evaluation has been done for the month of June 2015. The resulting MSPs can be characterised by a set of parameters that quantify their completeness and accuracy:

- Total number of tracks seen in the entire month;
- Average number of tracks (vessels) seen in the MSP at any one time;
- Average number of position reports per track;
- Time from MSP to closest (earlier or later) position report, measured as median, mean and 85 % percentile.

The first two parameters quantify the completeness of the MSP: how many ships are tracked? The second two its accuracy: how well are the ship positions determined? The fourth property, time from MSP to nearest position report, is characterised by three parameters, because the distribution of its values over all MSP points is very asymmetric; most values are small, but a few are very large. The median gives the best indication of where most values are; the 85 % percentile is the value where 85 % of the MSP points remain below and gives a better indication of where the large values are.

The tracking and calculation of the MSPs has been done with the data for the month of June with various combinations of data sources, in order to quantify the impact of the various data types and the number of satellites used. The tracking calculation is relatively fast – it takes only 7.5 minutes on a simple laptop to calculate all the 15-min-spaced MSPs for the whole month, using the 1.7 million position messages of the month as input. Nonetheless, a calculation of all MSPs for all possible combinations of sub-sets of satellites was not done, not even when limiting to maximally 55 combinations per number of satellites used, as was done for Figure 2-12. That would still lead to a computation time of 463 x 7.5 minutes = 58 hours. Instead, the performances were calculated only once for each number of satellites, by taking one arbitrary satellite to start with, then adding a second (arbitrary) one, then a third, etc., up to 11. Furthermore, the performance was calculated for LRIT only, and for all satellites plus LRIT. Terrestrial AIS was not at all used in this evaluation, because clock offsets in that data set introduced too many errors.

Results for the MSP

The results are tabulated in Table 2-1, and plotted in Figure 2-17. Inspecting the series of 1 to 11 AIS satellites, it can be seen that the number of vessels, seen in the whole month (column 1) as well as in the MSP at any one instant (column 2), levels off for a higher number of satellites. This was also seen before when counting the average daily number of MMSI numbers, which is in between the instantaneous number and the monthly total. The average number of position reports per track (column 3) on the other hand increases more or less linearly with the number of satellites. The measures for the time from the MSP positions to their nearest position report (columns 4-6) decrease significantly with the number of satellites, but less so toward the high end. Their decrease is a bit jumpy. This is because if a satellite is added with a similar orbit as the previous ones, even though it may add many messages, most of those will come at the same time as the messages already considered, leaving long gaps unfilled; but if a satellite with a different orbit is added, it will pass over at different times, filling gaps in the track. On the whole, after 7 satellites, there is almost no more improvement.

Regarding LRIT (first line in the table, leftmost points in the figures), it tracks much fewer ships, and has much fewer reports per ship. But the performance in terms of mean time to nearest report (at 1.8 hours) is comparable to that of 3 AIS satellites. The performance in terms of the 85 % percentile (at 2.78 hr) is even comparable to that of 6 AIS satellites, meaning that even though 3 AIS satellites have a good performance averaged over all ships, they leave a part of the ships with long time gaps.

Column	1	2	3	4	5	6
Data	Total # tracks in month	Average # tracks in MSP	Average # reports per track	Median time to nearest report (hr)	Mean time to nearest report (hr)	<85 % time to nearest report (hr)
LRIT	727	165.2	25.3	1.633	1.800	2.783
1 sat	4940	1187.6	421.1	2.417	2.624	4.687
2 sat	5002	1226.1	607.5	2.177	2.407	4.418
3 sat	5056	1262.4	798.7	1.353	1.843	3.613
4 sat	5075	1275.6	811.9	1.304	1.791	3.548
5 sat	5119	1314.7	1165.4	0.746	1.396	2.871
6 sat	5123	1319.5	1343.6	0.724	1.375	2.859
7 sat	5153	1331.9	1542.7	0.355	1.034	2.348
8 sat	5157	1334.5	1705.6	0.328	1.005	2.290
9 sat	5160	1338.6	2008.6	0.300	0.972	2.226
10 sat	5162	1341.2	2193.8	0.279	0.943	2.161
11 sat	5164	1344.6	2441.6	0.257	0.912	2.091
11 sat + LRIT	5183	1354.4	2436.1	0.255	0.891	2.027

Table 2-1. Properties of the MSPs based on tracking using (from top to bottom) LRIT data only, one AIS satellite, two satellites, three, etc., to 11 satellites, and at the bottom 11 AIS satellites + LRIT. The columns are explained in the text.

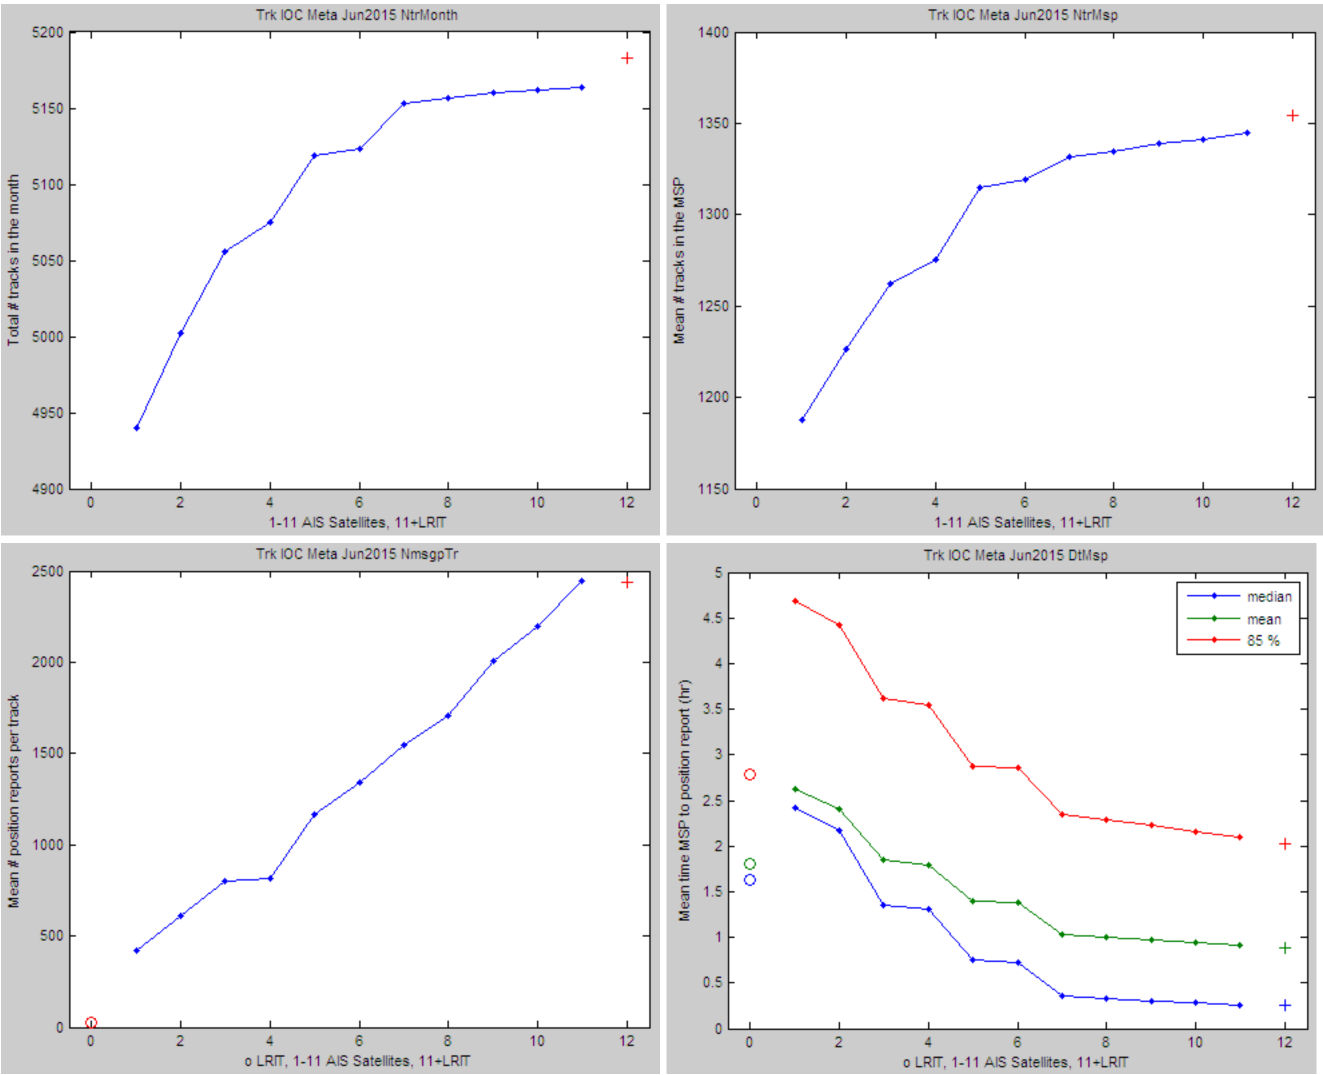

Figure 2-17. Performance indicators of the MSP after tracking, as a function of using: only LRIT data (at 0 on the horizontal axis, o symbol); using the data from 1 to 11 satellites; and using 11 satellites plus LRIT (at 12 on the horizontal axis, + symbol). Top left: The total number of tracks seen during the whole month of June 2015. Top right: the average number of tracks seen in the MSP at any one time. Bottom left: The average number of position reports

per track. Bottom right: The time between the MSP and the nearest position report in hours, as median (lowest values, blue); mean (intermediate values, green); and 85 % percentile (highest values, red).

Figure 2-18 shows how the number of tracks in the MSP, of which Table 2-1 column 2 gives the average, varies with time. The graph for the one-satellite case (top right) shows an oscillation that is related to the revisit periodicity of the satellite as dictated by its orbit. If the maximum extrapolation time of 9 hours would be extended, this oscillation would be further dampened, but that would also give rise to spurious vessel presence in the MSP. In the graph for 6 satellites (bottom left) the oscillation is much reduced as the various satellites are not fully synchronised. The LRIT (top left) does not show this regular oscillation, only some random noise (LRIT also does not contain Class B vessels).

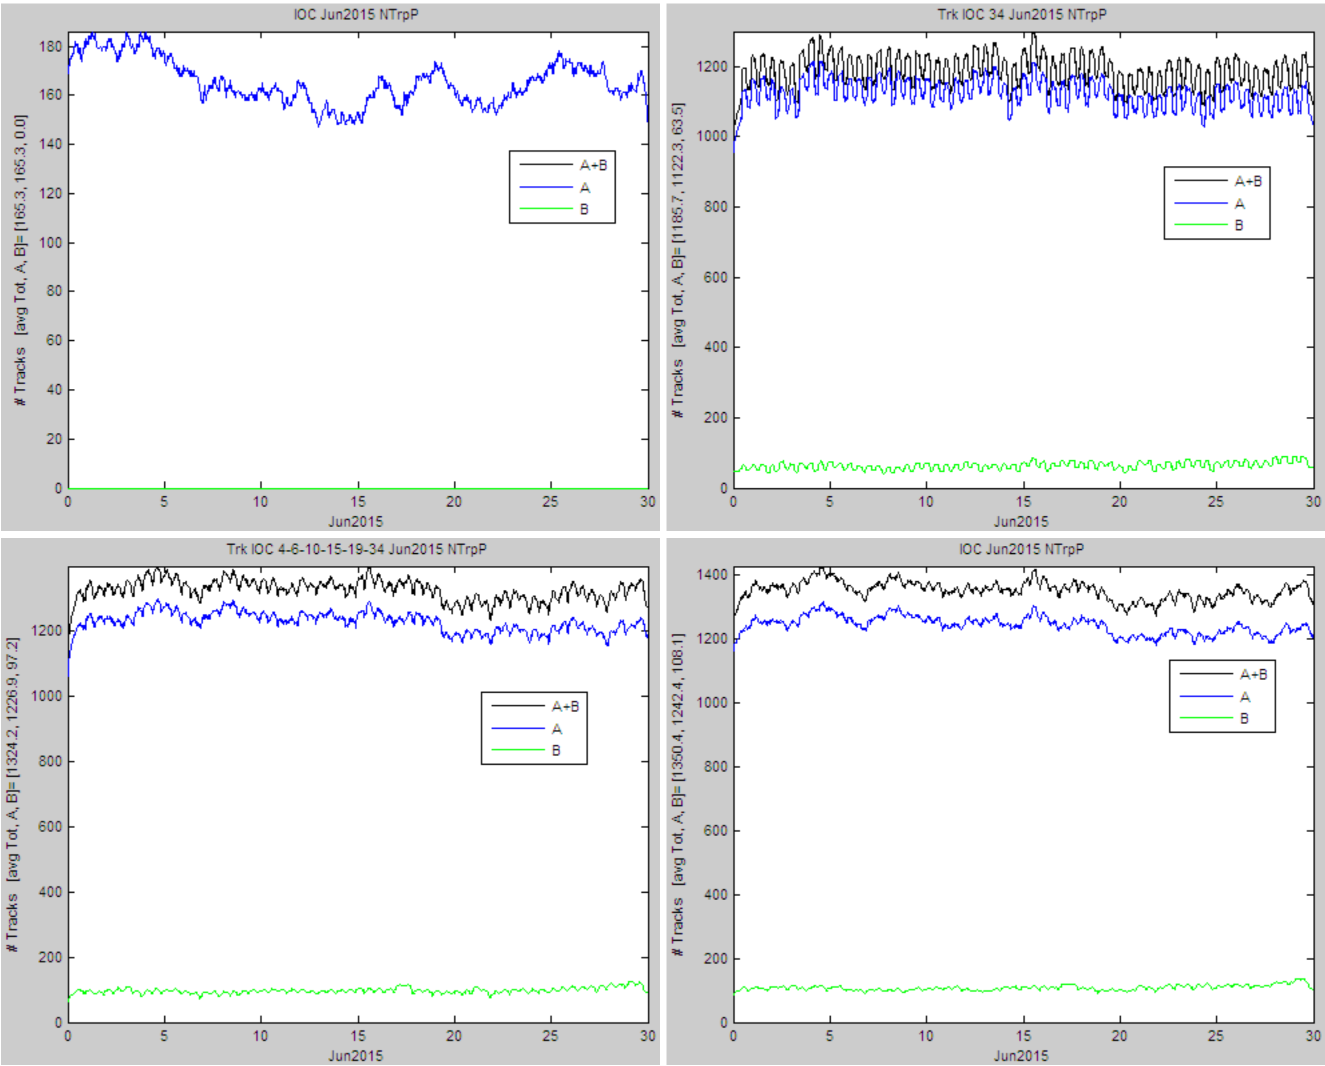

Figure 2-18. The number of tracks in each 15-minute MSP for the month of June 2015, split in Class A and Class B vessels and their sum. Top left: LRIT only. Top right: one AIS satellite. Bottom left: 6 satellites. Bottom right: 11 AIS satellites.

Figure 2-19 visualises the asymmetry of the distribution of the times from the MSP points to their nearest position report. It is from these distributions that the median, mean and 85 % percentile used above have been calculated. The histogram for LRIT (top left) directly reflects the 6 hourly reporting: most MSP points are between 0 and 3 hours away from their nearest position report. In some cases where an LRIT report is missed, it is between 0 and 6 hours. There is a tail up to 9 hours, the maximum time for extrapolation, when no more LRIT reports are received – but not due to leaving the AOI because then the tracking is stopped.

The histogram for one AIS satellite (top right) also reflects a periodicity due to the satellite orbit.

The histogram for 6 AIS satellites (bottom left) shows a peak left of 1 hr, a plateau left of 3.5 hr and a tail up to 6-9 hr. The distribution for all 11 AIS satellites (bottom right) shows that most MSP points are very close to a position report, with the peak to the left of around 0.5 hr; but a few are up to 3-4 hr away, and very few up to 9 hr. Adding the LRIT to 11 AIS satellites (not shown) does not change the shape of the distribution much.

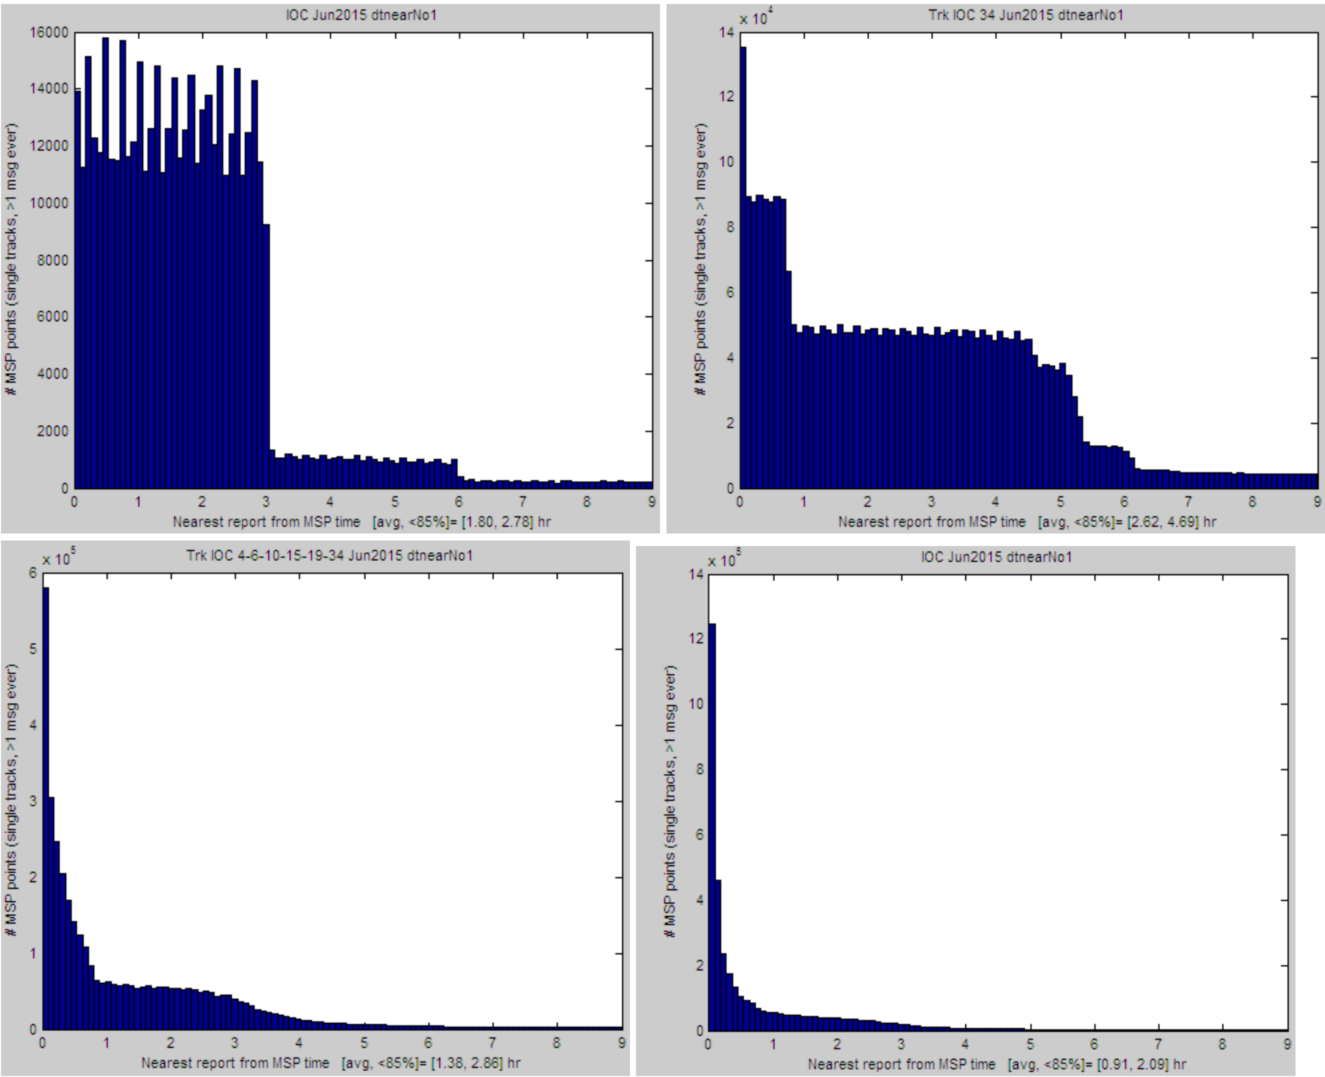

Figure 2-19. The distribution of the times between an MSP point and its nearest position report. The horizontal scale is in hours and runs from 0 to 9. The vertical scale counts MSP points, from all tracks taken together. Top left, LRIT only. Top right, one AIS satellite. Bottom left, 6 AIS satellites. Bottom right, all 11 AIS satellites.

Instead of making a histogram of the individual values of the time from MSP to nearest position report from all MSP points together, it is also possible to first average over all MSP points of one track, and thus get one mean value per track. These mean values can then be plotted in a histogram. The result is displayed in Figure 2-20. These distributions are narrower than the previous ones, because they are distributions of values that have already been averaged. It shows that for LRIT (top left), most tracks have a mean time of MSP to nearest position report between 1.5 and 2.1 hr. For one AIS satellite (top right) it is between 2 and 3 hr. For 6 satellites it is between 0.7 and 1.4 hr; and for 11 satellites between 0.2 and 1 hr.

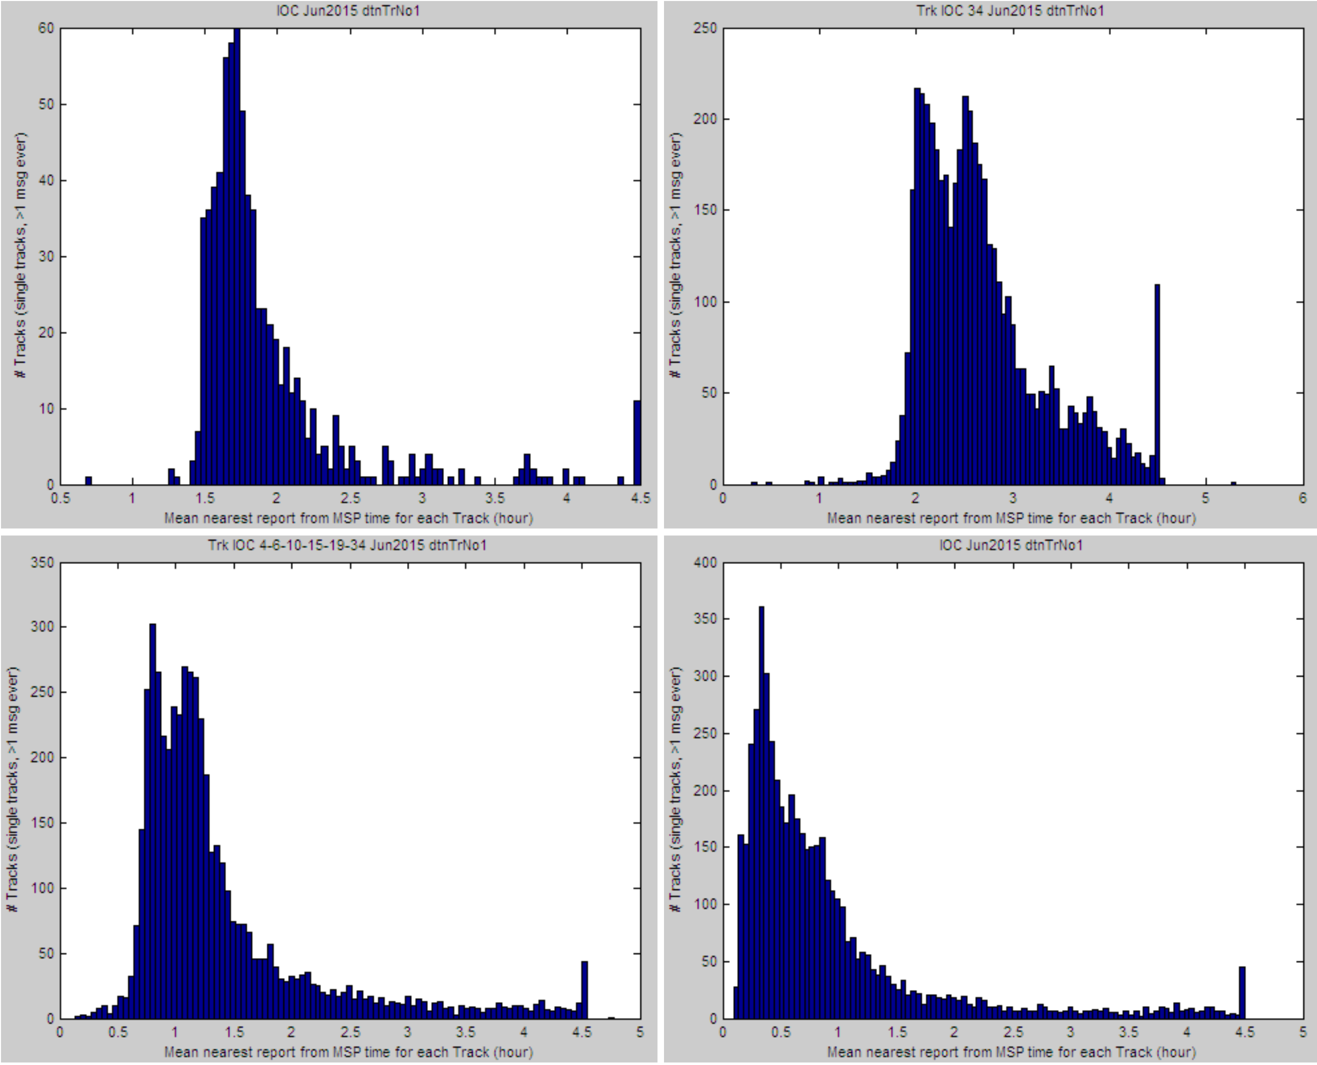

Figure 2-20. The distribution of the mean values for each track of the time from MSP to nearest position report. The horizontal scale is in hours. The vertical scale counts tracks. Top left: LRIT only. Top right: One AIS satellite. Bottom left: Six satellites. Bottom right: All 11 satellites.

3 Monthly ship density maps

3.1 Procedure

Ship density maps have been compiled on a monthly basis. The data that have been used are the satellite AIS data and the LRIT. Terrestrial AIS has not been used, to avoid a bias in the areas that are covered by that, which would prejudice the density maps because those areas are so unequally distributed (see Figure 2-4 bottom left). As a first step, the MSPs are constructed by the tracking process described in the previous chapter. The same parameters as mentioned there have been used in the tracking. Another reason why terrestrial AIS was not used is because the clock errors in there (discussed in the previous chapter) lead to erroneous tracks. Then, each of the MSP ship positions is counted in a grid with a cell size of 0.025 degrees, which is approximately 1.5 nautical miles or 2.8 km. The resulting 2-D density map is convolved with a boxcar of 6 x 6 grid cells, leading to a resolution of 0.15 degrees which is approximately 9 nautical miles or 16.7 km. The resulting numbers are normalised to obtain ship density in units of the average number of ships present at any one time per square degree. The density maps cover the entire AOI.

The tracking also allows to compute the speed of the ship, derived from a ship's consecutive MSP positions. This enables the separation of MSP positions by speed. This provides additional insight in the ship behaviour: fast moving ships are mostly in transit, while slow moving ships out on the open sea are mostly engaged in some activity, like fishing or exploration. Ships are also slow moving in and near ports. Some trials have shown that a speed limit of 6.5 knots gives a useful separation between transit and other activity. Therefore, the ship density maps have been split into these two domains: fast moving ships (> 6.5 knots) and slow moving ships (< 6.5 knots). The latter also include stationary ships such as found moored in ports or at anchor in port vicinity.

3.2 Results

The figures in this section show the ship density maps for each month between September 2014 and September 2015, split into fast moving and slow moving (including stationary) ships. The scale is logarithmic to take into account the large dynamic range of the densities, as previously discussed with Figure 2-4, and signifies the 10-logarithm of the average number of ships present at any one time per square degree unit of surface. The very high numbers are particularly reached in ports, in areas that are so small that they are difficult to recognise at the scale of the figures (small red dots on the coast). The size of the resolution cell of 0.15 degrees is printed here at a size of 0.28 mm.

The ship density maps of the fast moving ships are always on the left. They are dominated by the main ship traffic routes. They do not vary much from month to month. The density maps of the slow moving ships are always on the right. Most of the structure in those maps can be attributed to fishing activity. Linear features are due to a few ships transiting at a slow speed. These maps show much monthly variation, mostly reflecting the shifting of the fishing areas with the season. In many places, the ship density displays a strong correlation with EEZ boundaries (the pink lines are non-authoritative EEZ boundary indications, for orientation only).

The PMAR-MASE data stream worked at full capacity between 15 Sep 2014 and 15 Sep 2015. Nevertheless, some data were already coming in during the first half of Sep 2014 (for setting up and testing the system), and were still coming in during the second half of Sep 2015 (as some providers did not immediately switch off their data flow). Therefore, density maps of these entire months could be produced.

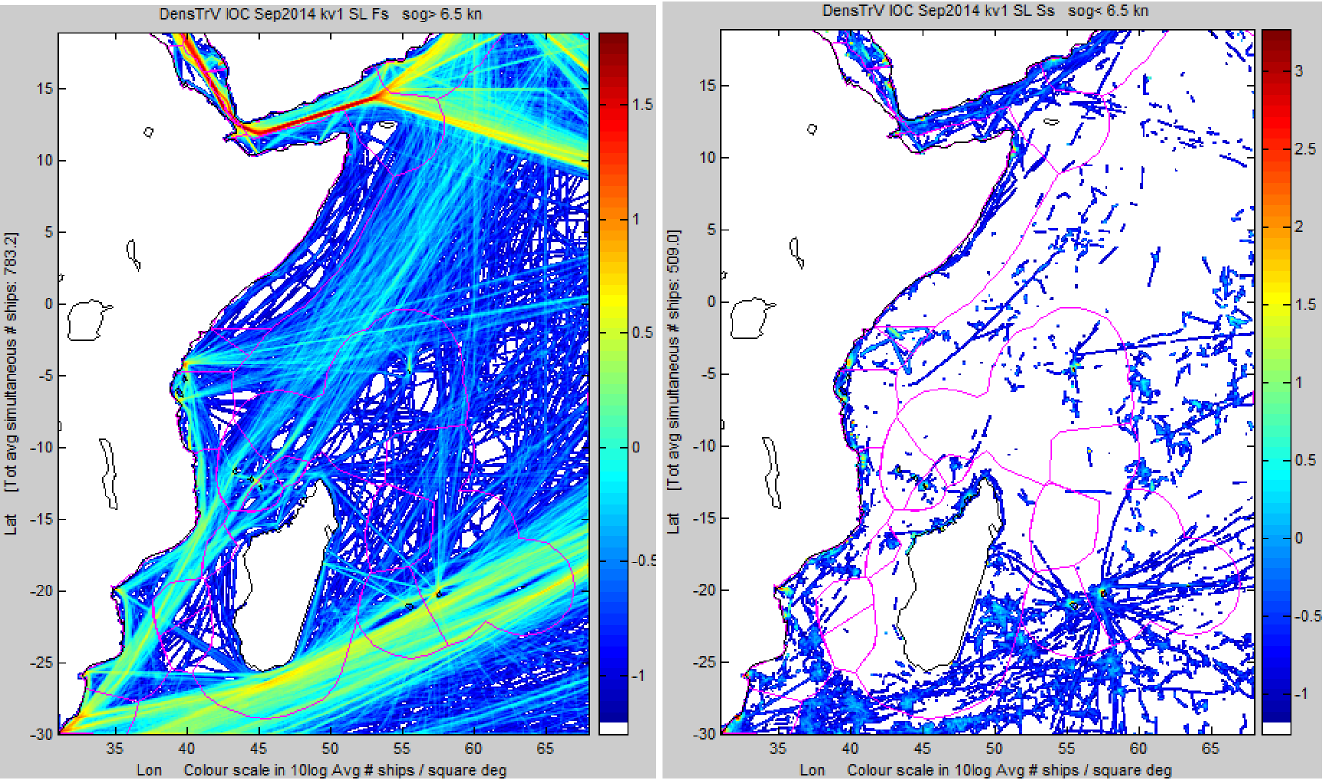

Figure 3-1. Ship density map for September 2014. Fast moving ships on the left, slow moving ships on the right.

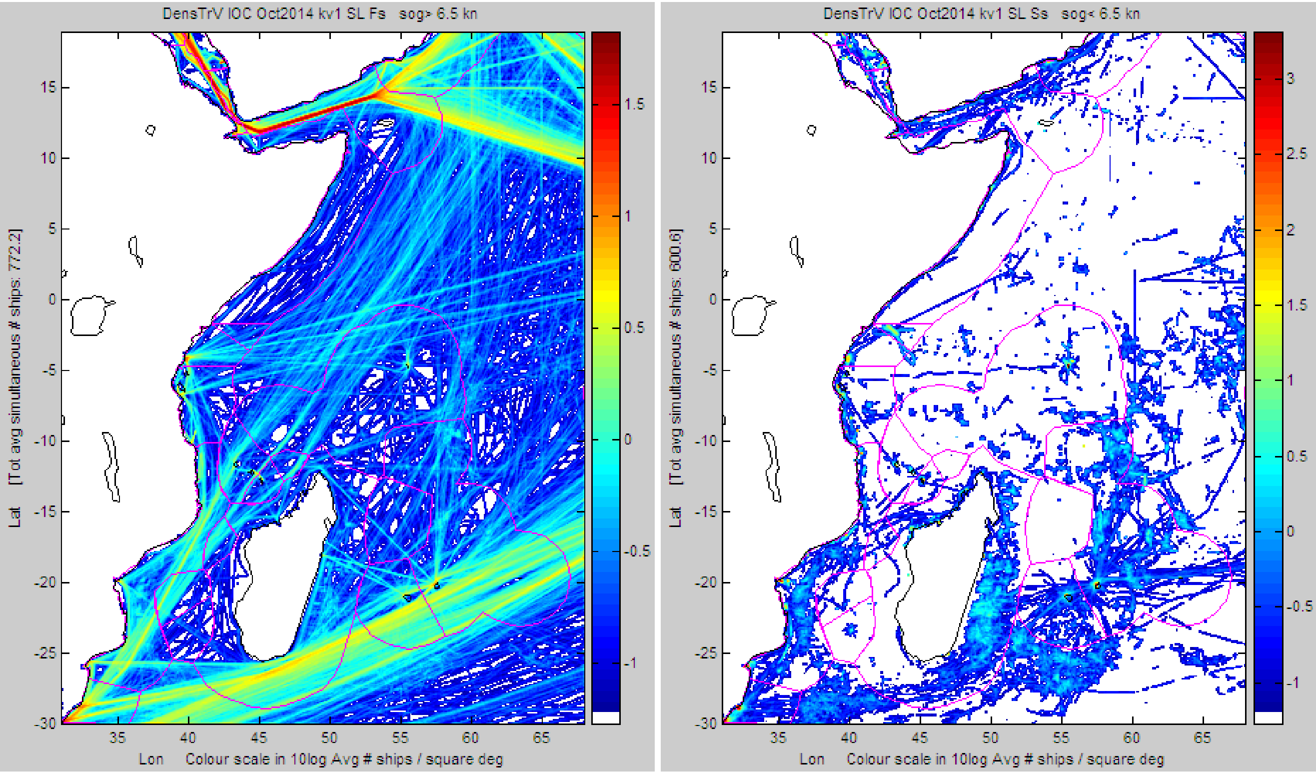

Figure 3-2. Ship density map for October 2014. Fast moving ships on the left, slow moving ships on the right.

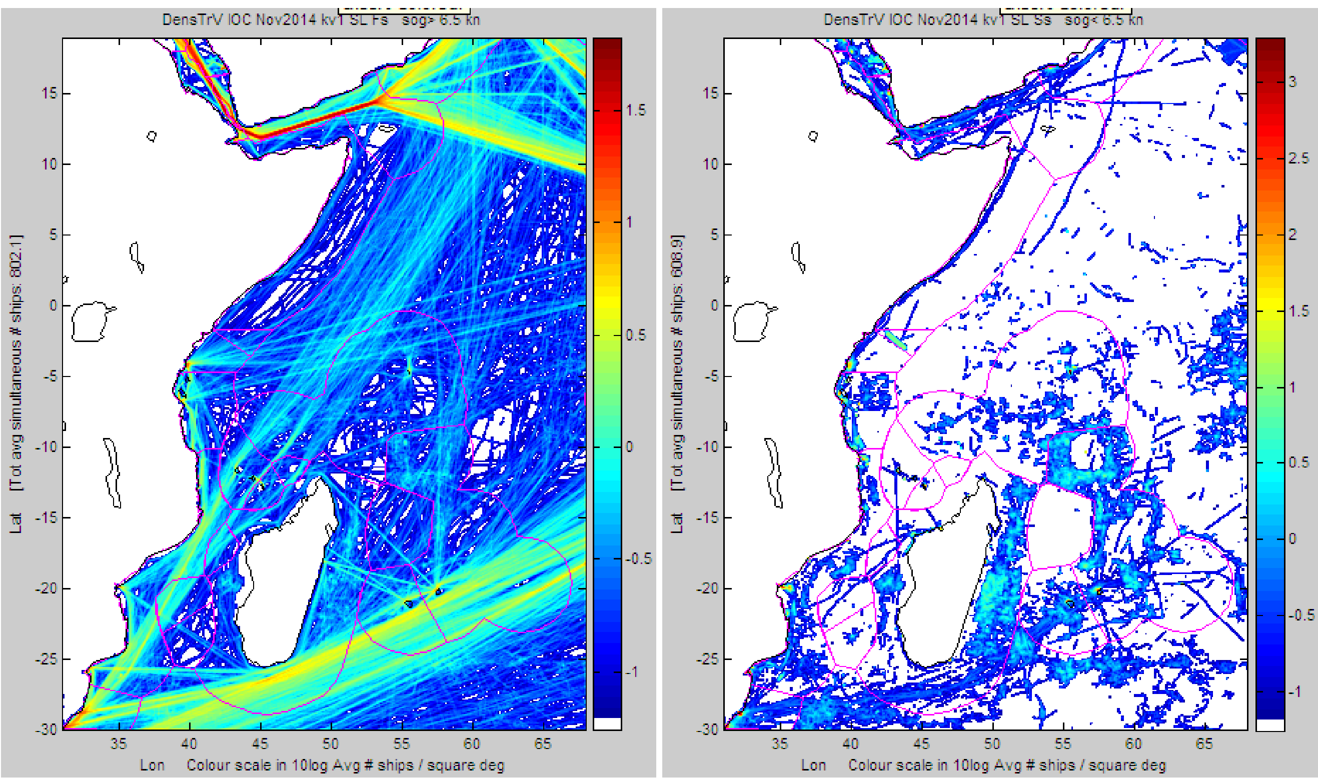

Figure 3-3. Ship density map for November 2014. Fast moving ships on the left, slow moving ships on the right.

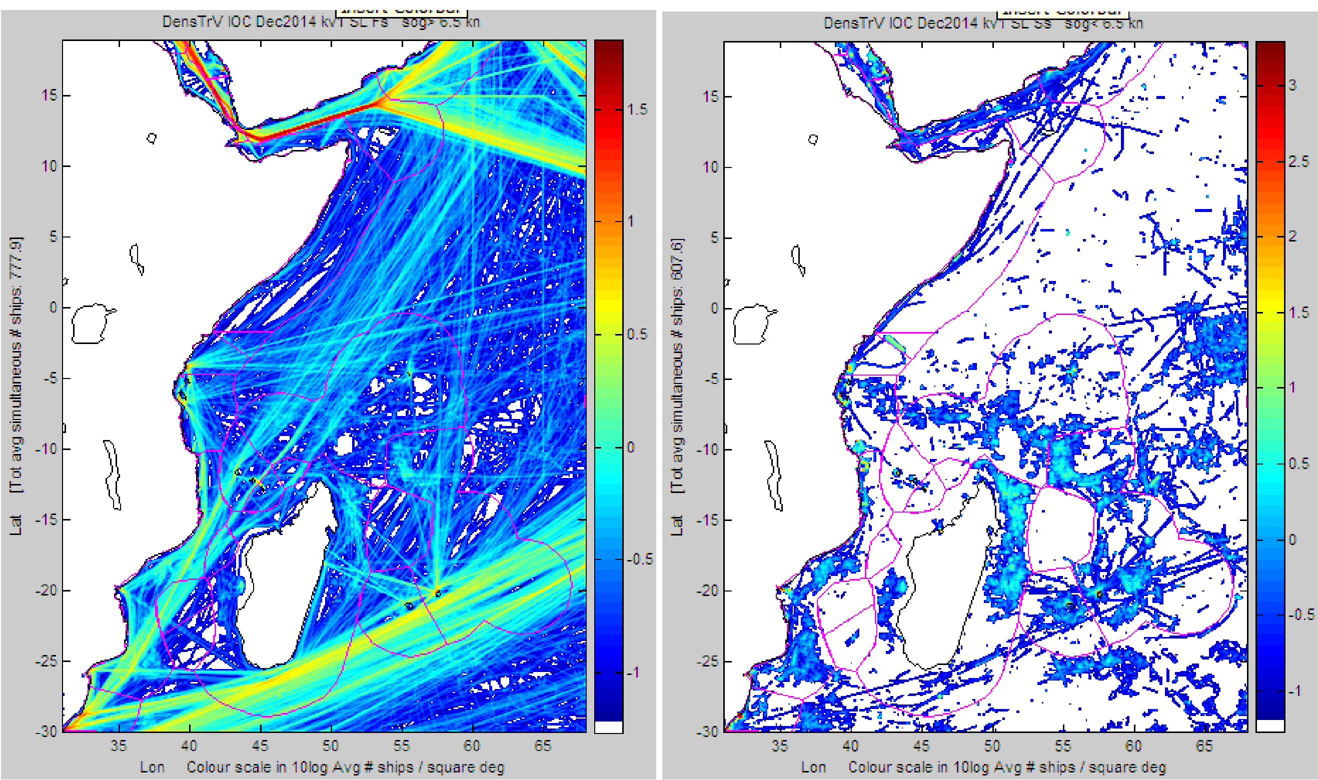

Figure 3-4. Ship density map for December 2014. Fast moving ships on the left, slow moving ships on the right.

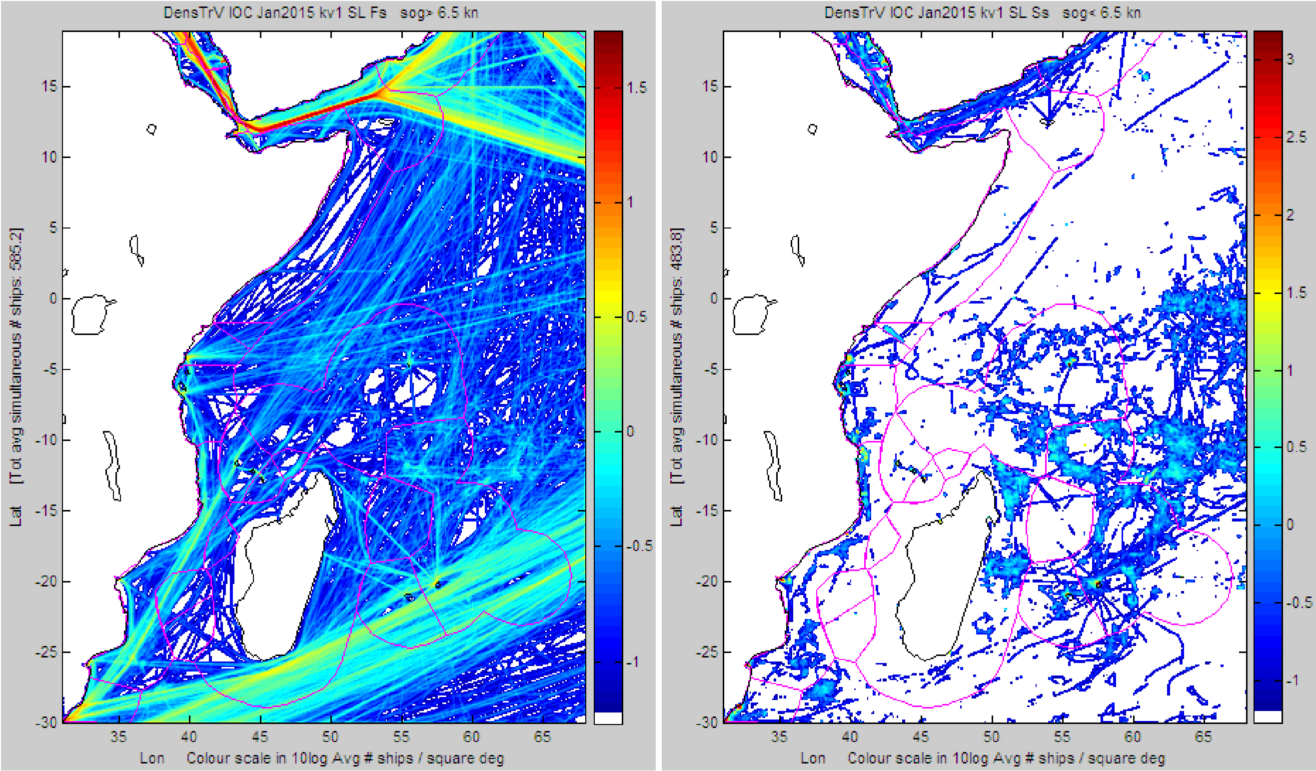

Figure 3-5. Ship density map for January 2015. Fast moving ships on the left, slow moving ships on the right.

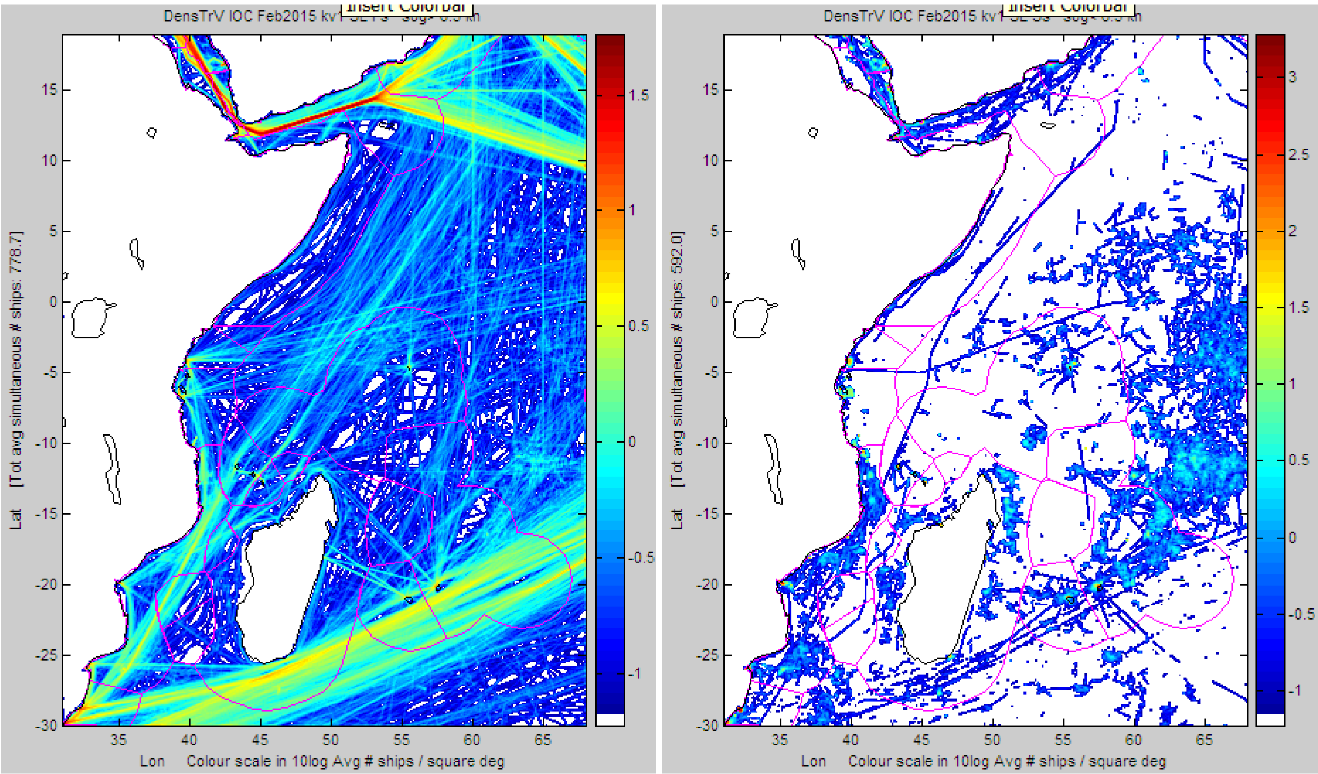

Figure 3-6. Ship density map for February 2015. Fast moving ships on the left, slow moving ships on the right.

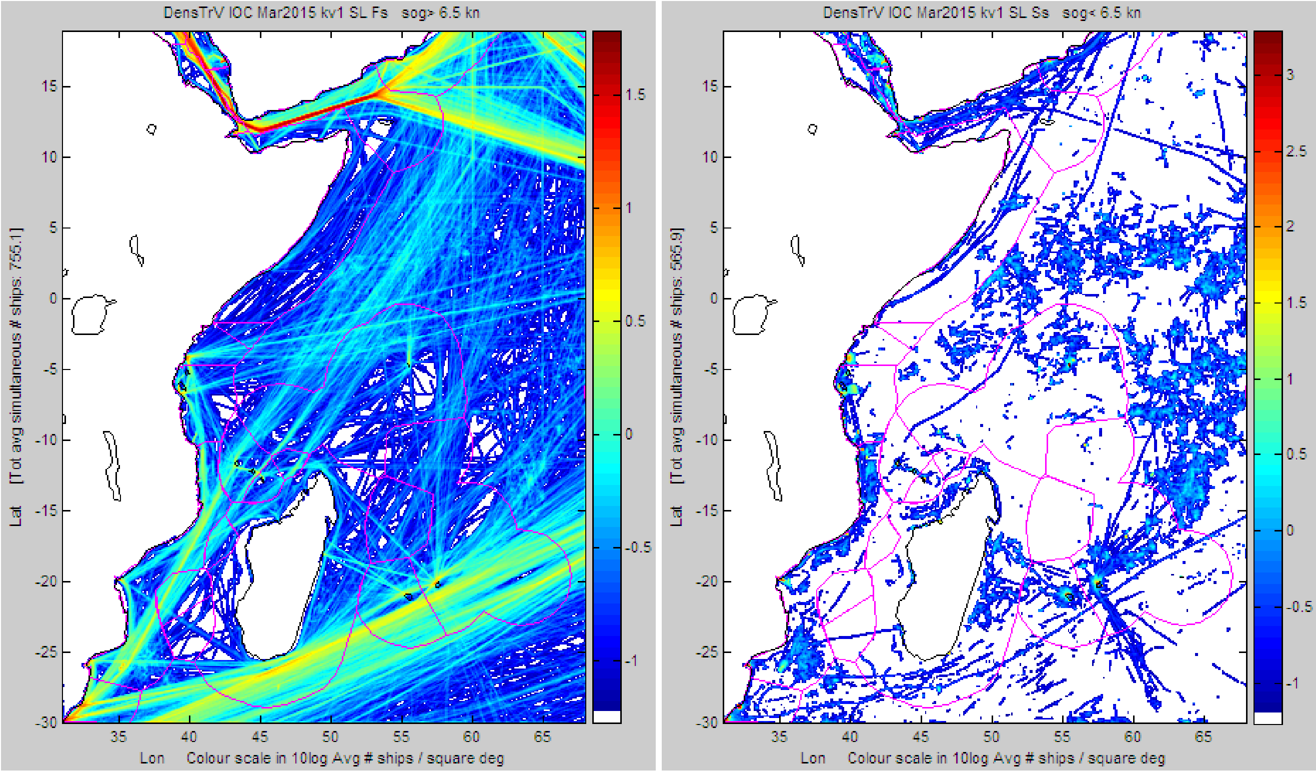

Figure 3-7. Ship density map for March 2015. Fast moving ships on the left, slow moving ships on the right.

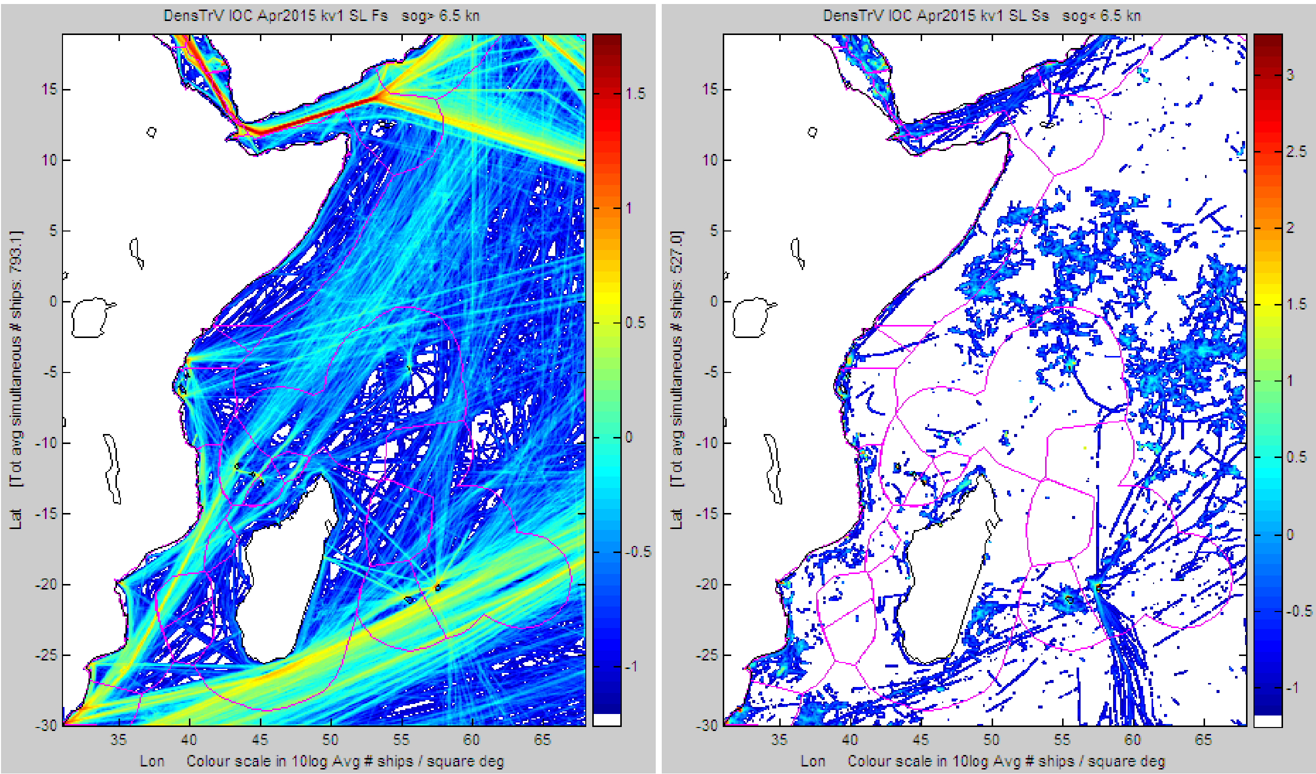

Figure 3-8. Ship density map for April 2015. Fast moving ships on the left, slow moving ships on the right.

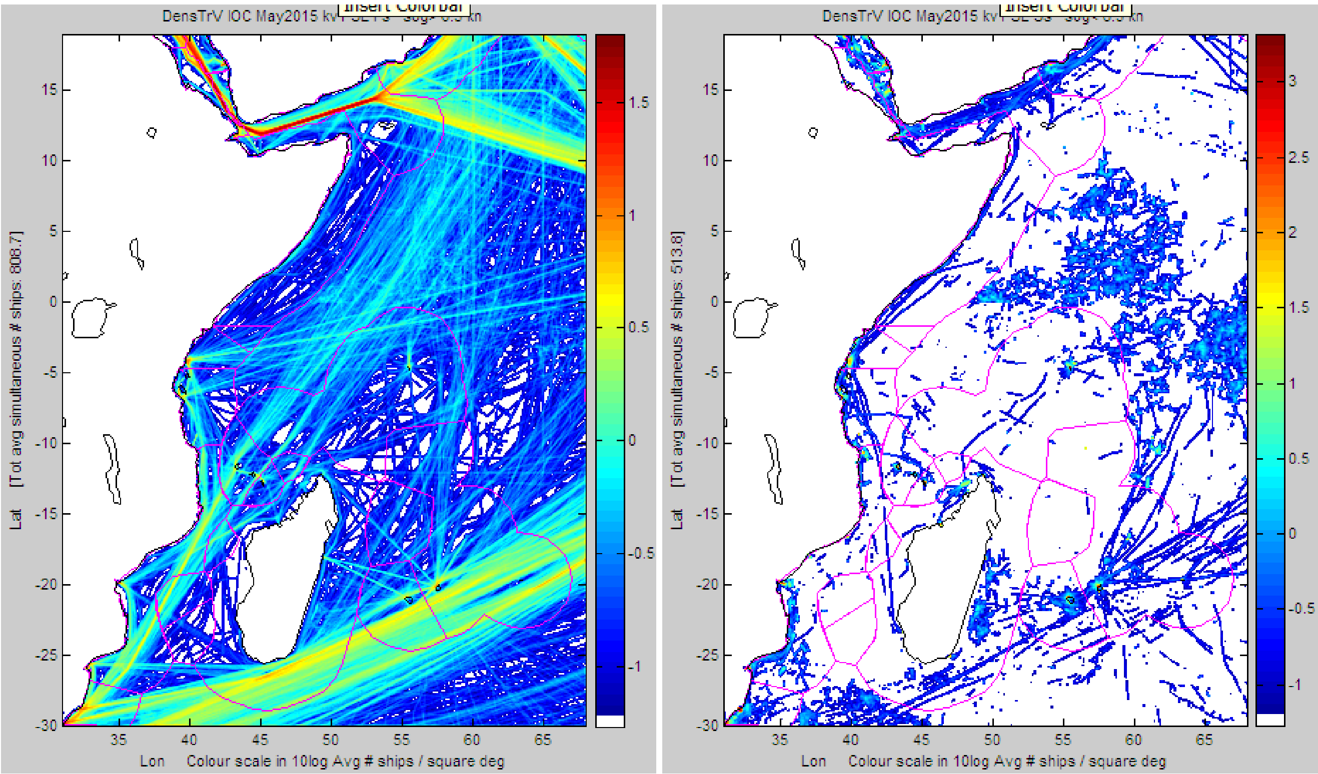

Figure 3-9. Ship density map for May 2015. Fast moving ships on the left, slow moving ships on the right.

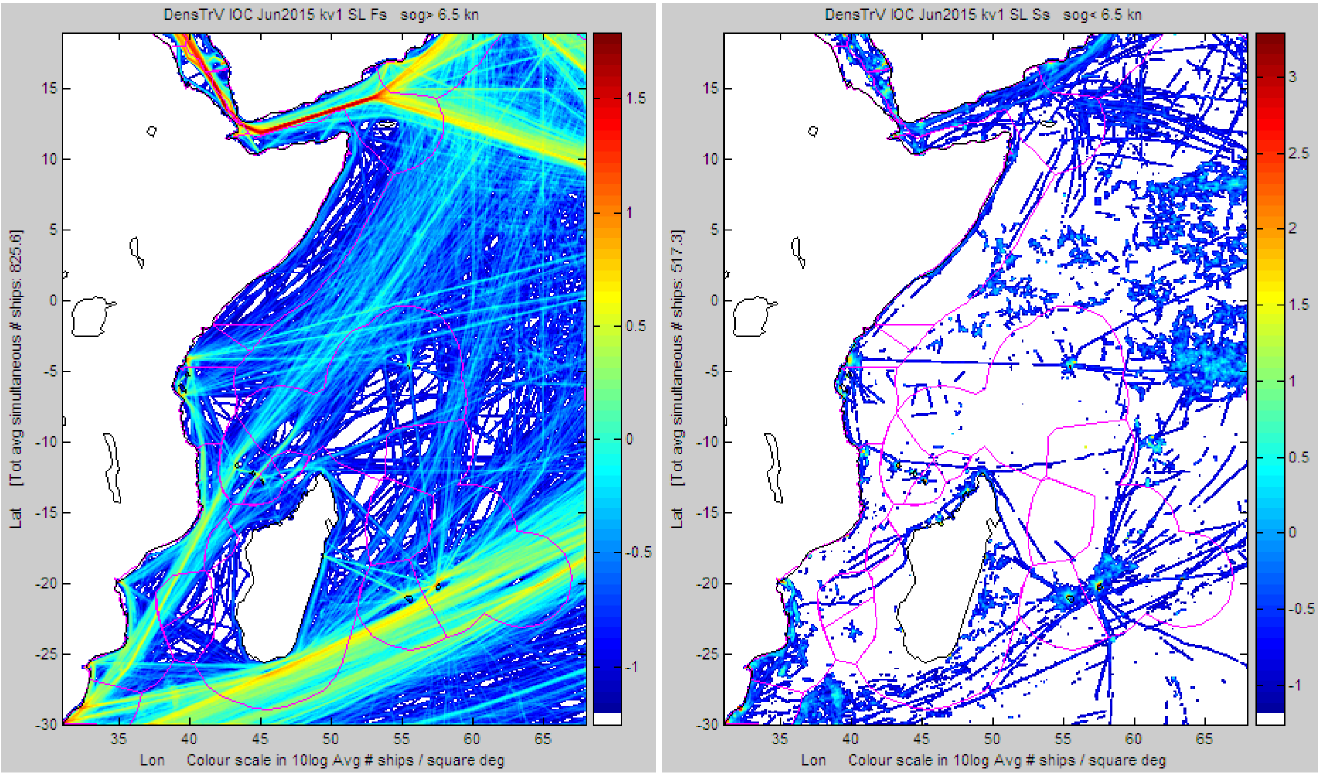

Figure 3-10. Ship density map for June 2015. Fast moving ships on the left, slow moving ships on the right.

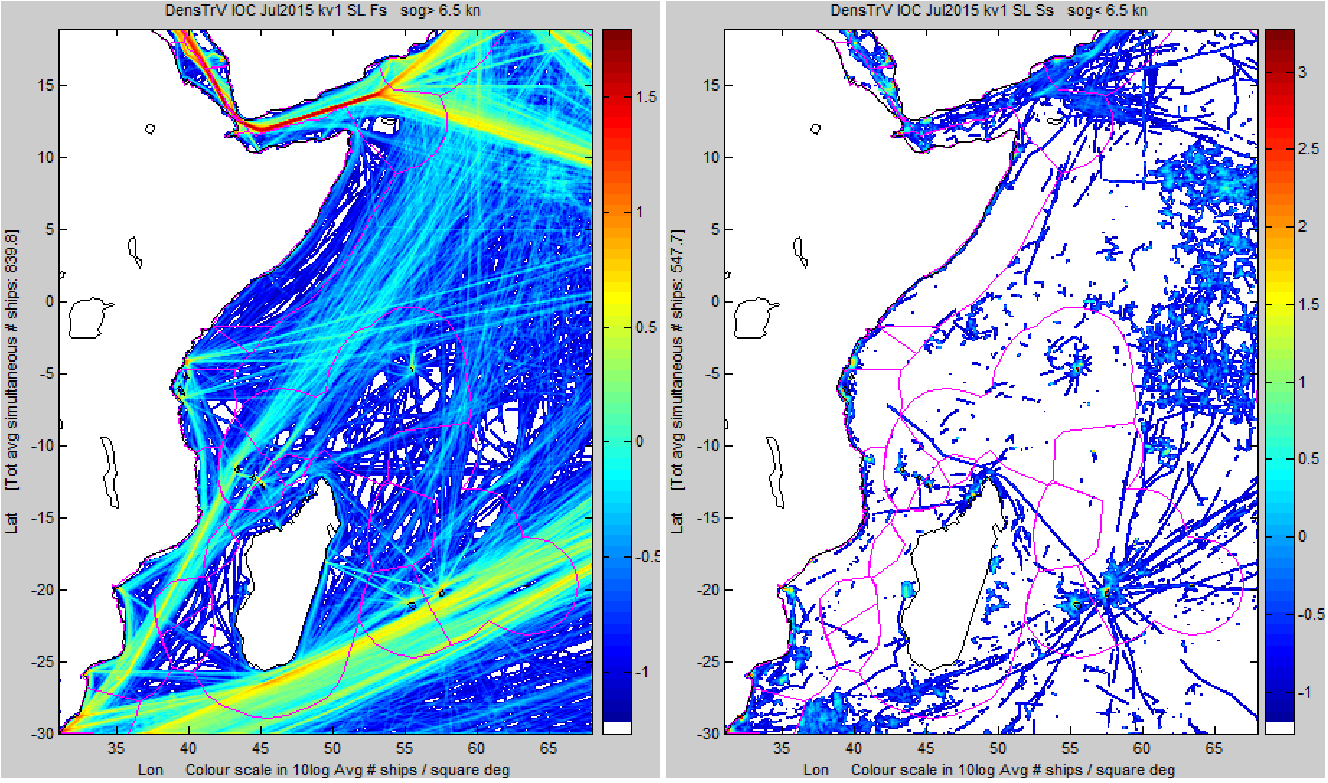

Figure 3-11. Ship density map for July 2015. Fast moving ships on the left, slow moving ships on the right.

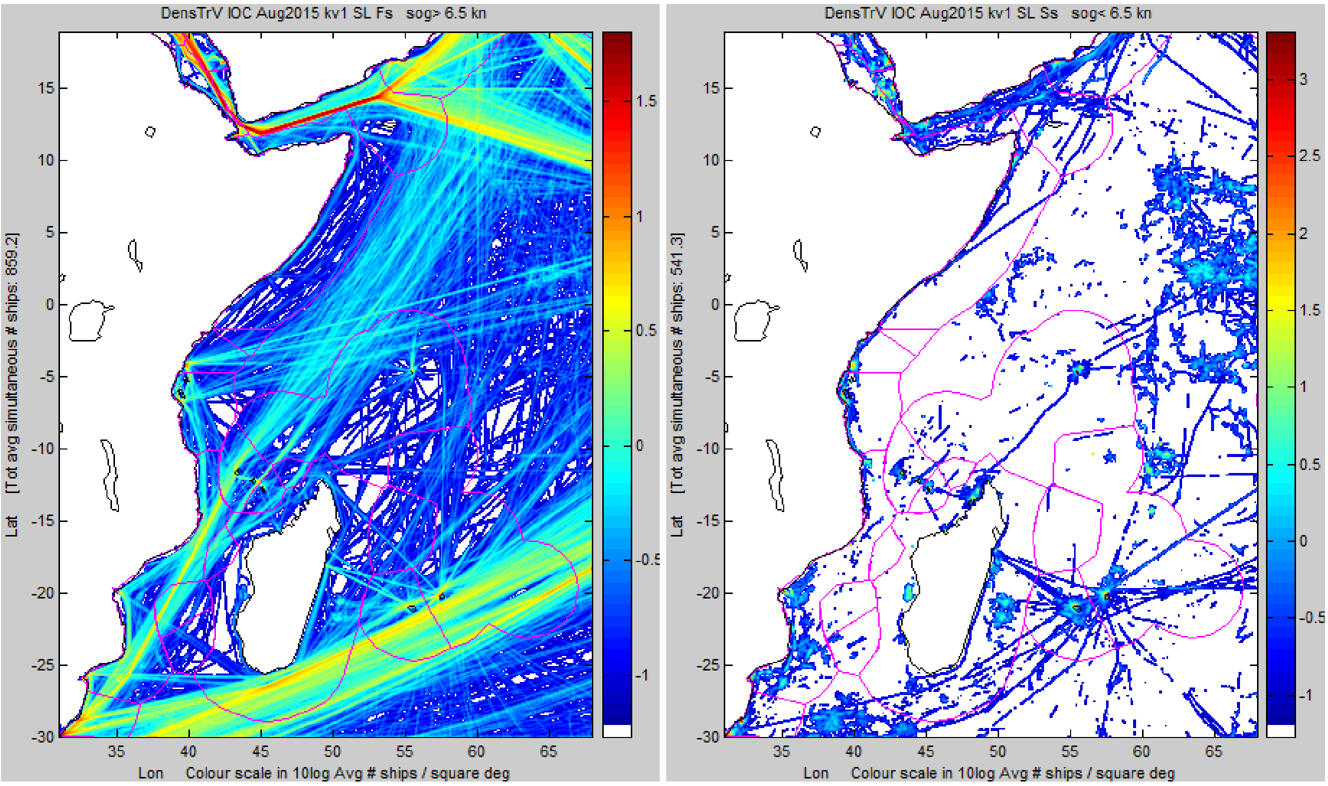

Figure 3-12. Ship density map for August 2015. Fast moving ships on the left, slow moving ships on the right.

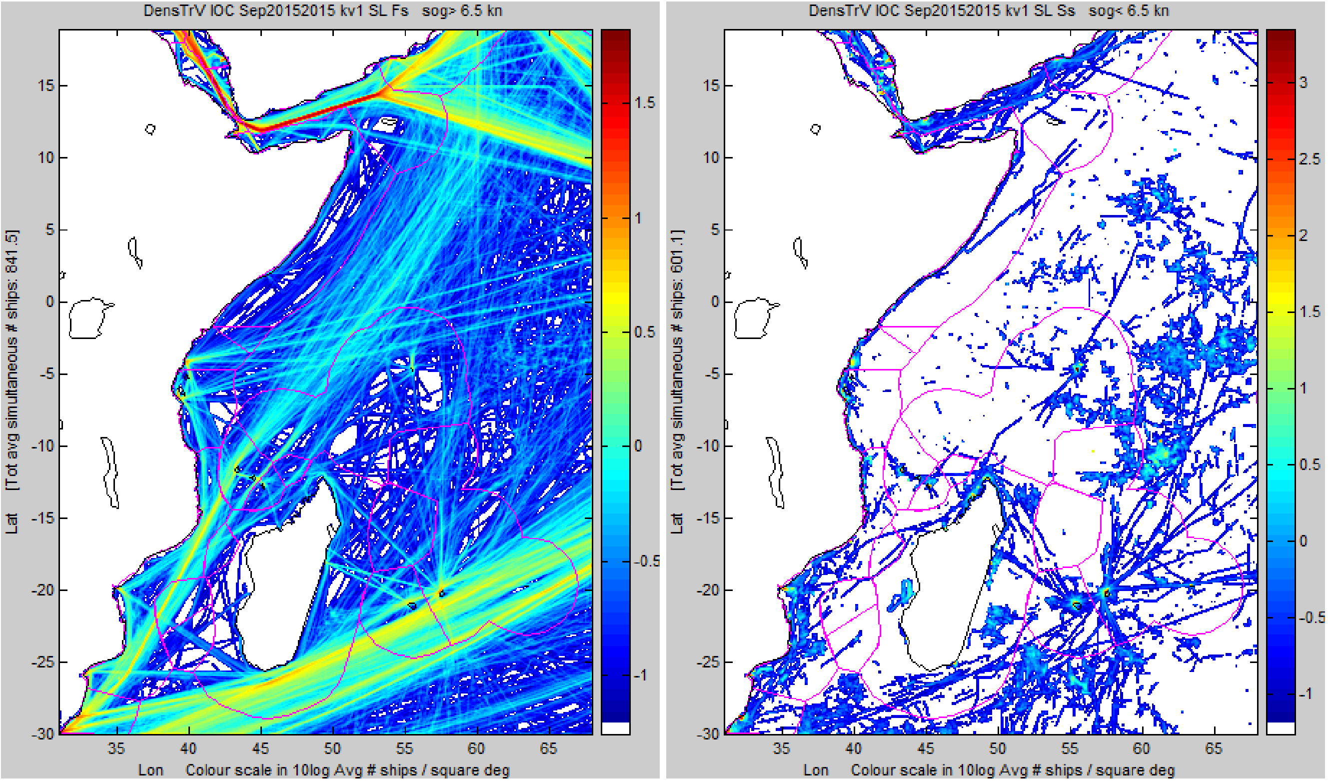

Figure 3-13. Ship density map for September 2015. Fast moving ships on the left, slow moving ships on the right.

4 Ship detections from satellite SAR

4.1 Technology and method

The ship reporting systems (AIS and LRIT) provide very rich information on ship presence and behaviour, but only for the reporting ships. Ships that are not transmitting on these systems are not seen. In order to find out more about the presence of such ships on the open sea away from coastal monitoring sites, Earth observation satellites can be used. Such satellites make images of the Earth surface, also over sea. Satellites carry optical cameras for images in the optical (visual) domain, but also radars for images in radar wavelengths. This study uses radar images to prevent the obscuration by clouds that hinders optical imaging from space. The type of radar used on satellites is the Synthetic Aperture Radar (SAR), which is a method to still get good resolution even from a long distance. The character of the SAR image data is very different from that of the ship reporting data. The satellite SAR provides an image of the sea surface of a limited size, in which the ships show up as bright points against a darker background. There is no possibility to identify the ships that are seen; just to notice their presence and obtain their geographic position. Small ships are not detected in the image, the exact size limit being dependent on many factors including image type and sea conditions. Other effects on the sea surface besides ships may cause false alarms, such as breaking waves or interference from nearby land. A satellite SAR image does not show any movement, and can only be taken once every so often; there is no possibility for continuous or persistent monitoring.

Satellite SARs can make images in different modes: wide-area images at low resolution, detecting only big targets; or small-area images at high resolution, where targets down to just a few meters can be detected if the circumstances are right. The latter are less useful for maritime surveillance, because their coverage is too small (of the order of 5-10 km wide). The wide-area images on the other hand can have coverages of 150 km to 400 km wide.

The targets detected in a satellite image, taken at a certain time, are compared to the ships that are known to be present at that time from the reporting data. In this way it is established for each ship detected in the SAR image whether it was already known from the reporting data, or whether it is a non-reporting ship. In addition, there could be reporting ships that are not detected in the SAR image. This process is called correlation (between the SAR ship detections and the reported ship positions). The correlation was done using the JRC's Blue Hub software.

The ship detections from the satellite SAR images are referred to as VDS (for Vessel Detection System). The ship positions of the reporting ships at the time of the satellite image are referred to as Interpolations, because these positions are interpolated from the bracketing ship position reports (i.e., the reports available from just before and just after the time of the satellite image). The outcome of the correlation process leads to three kinds of target:

- Correlated target, where a VDS position and an Interpolation match;
- Uncorrelated VDS, where a VDS position (satellite SAR detection) does not have a matching Interpolated reported position – so a non-reporting ship (or a false alarm);
- Uncorrelated Interpolation, where an interpolated position from a reporting ship is not matched with a VDS target, so a ship that is not detected in the satellite SAR image.

4.2 Satellite SAR data used

Two different satellite SARs were used, Radarsat-2 and Sentinel-1.

Radarsat-2 is commercially operated by the Canadian company MDA. Its images are commercially offered through European resellers. The PMAR-MASE project has bought 29 Radarsat-2 images. Of those, 21 were "ScanSAR Narrow B" (SNB) mode that has 300 km x 300 km image size at 50 m resolution; and 8 were "Wide Fine" mode that has 150 km x 150 km size at 8 m resolution. All images were taken in the period 15 Aug – 9 Sep 2015. Those images were analysed with the JRC's "SUMO" software for ship detection. The analysis is semi-automatic; first, the images are subjected to an automatic target detection and correlation process, and then the results are verified by an operator.

In addition, ship detection results were received through a scientific exchange from Defence Research and Development Canada (DRDC) from 5 Radarsat-2 images in MSSR-DVWF mode, which is an extra wide mode for ocean surveillance that was not commercially available (at the time – it is now). It has a 450 km swath at 40 m resolution. These images were acquired in the period 24 Feb – 10 Mar 2015.

Sentinel-1 is a relatively new SAR satellite under the EU's Copernicus program, operated by the European Space Agency (ESA) on behalf of the European Commission. Its images are available for free. However, unlike for the commercial satellite SARs, it is not possible to request the collection of Sentinel-1 images in specific locations and times. So one can only use what has been collected, if indeed anything of interest has been collected. Consulting ESA's Sentinel-1 archive for images within the PMAR-MASE AOI and time frame, many hundred image were found. Most are on the coast. All these images have been downloaded, plus some more that were acquired further south, just outside the AOI; in total 1,500 Sentinel-1 images. Also these images have been analysed with the SUMO ship detector, but it is not possible to do the manual verification for so many images. Therefore, most of these 1,500 images were analysed in the so-called 'batch' mode, fully automatic. Their results are therefore less reliable than the manually verified ones. Only 12 images from ESA's Sentinel-1 archive were selected for manual verification; these were acquired over the Somali coast. They are all of the "Interferometric Wide" (IW) mode that has 250 km swath width at 20 m resolution. Most of these 12 were acquired between 24 Dec 2014 and 3 Jan 2015, and two on 12 Sep 2015.

The correlation of the ships detected in the satellite SAR images (the VDS) was done with the AIS and LRIT collected during 15 Sep 2014 – 15 Sep 2015 as described in the first chapters of this report.

4.3 Results

Radarsat-2 SNB and Wide Fine

Figure 4-1 shows the locations and dates of the commercially obtained Radarsat-2 images; the ScanSAR Narrow B mode images on the left and the Wide Fine mode on the right. The acquisitions have been in strips of 2-6 images long.

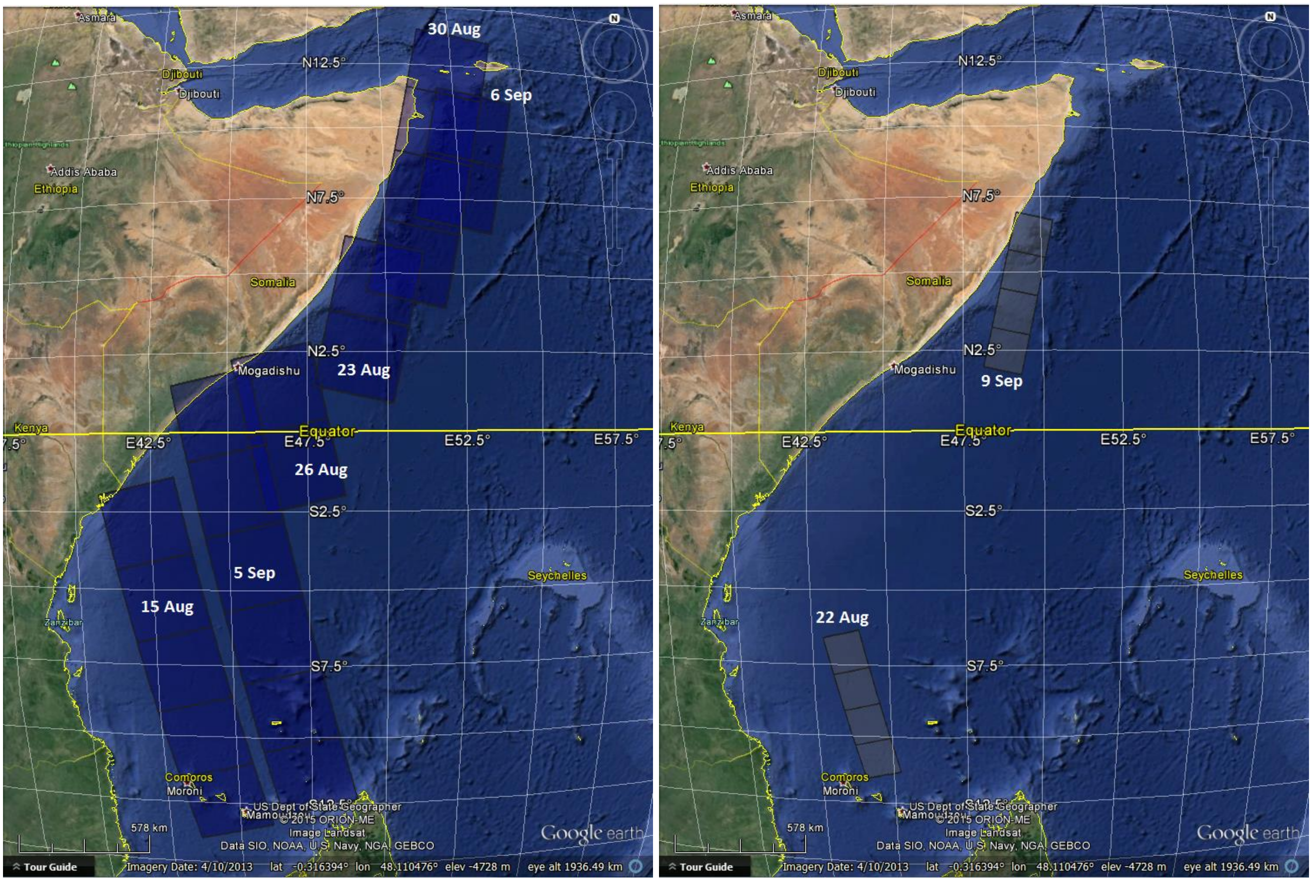

Figure 4-1. Locations and dates of the 29 commercially procured Radarsat-2 images. Left: ScanSAR Narrow B mode (300 km wide). Right: Wide Fine mode (150 km wide). (Google Earth map backgrounds.)

Figure 4-2 shows the results of the ship detection and the correlation. Only 38 ships have been detected in the 29 images, which is not a high number (average 1.3 ship per image). These were compared to 39 Interpolated reporting ship positions. This led to 34 correlated targets, 4 uncorrelated VDS (non-reporting ships), and 5 uncorrelated Interpolations (undetected ships).

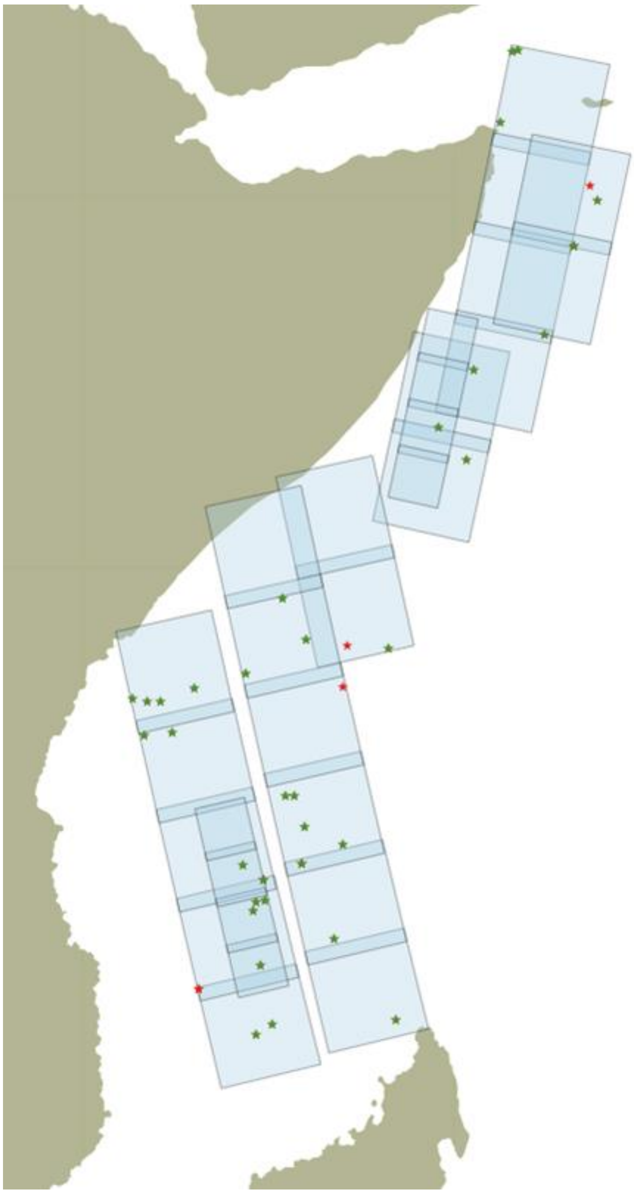

IMAGE	#VDS ☆	#INT	#CORR	#VDS uncor	#INTP Uncor
15-08-2015 15:30:15	3	4	2	1	2
15-08-2015 15:30:55	2	2	2	0	0
15-08-2015 15:31:35	0	0		0	0
15-08-2015 15:32:16	2	2	2	0	0
15-08-2015 15:32:56	4	4	4	0	0
22-08-2015 15:26:44	1	1	1	0	0
22-08-2015 15:27:04	2	2	2	0	0
22-08-2015 15:27:25	1	1	1	0	0
22-08-2015 15:27:46		0		0	0
23-08-2015 02:23:40	1	1	1	0	0
23-08-2015 02:24:20	1	1	1	0	0
26-08-2015 15:12:27	2	1	1	1	0
26-08-2015 15:13:07		1		0	1
30-08-2015 02:17:17	3	3	3	0	0
30-08-2015 02:17:57		2		0	2
30-08-2015 02:18:37		0		0	0
30-08-2015 02:19:18		0		0	0
05-09-2015 15:17:55	1	1	1	0	0
05-09-2015 15:18:35	2	2	2	0	0
05-09-2015 15:19:15	4	4	4	0	0
05-09-2015 15:19:56	1	0		1	0
05-09-2015 15:20:36	3	3	3	0	0
05-09-2015 15:21:17		0		0	0
06-09-2015 02:13:46	3	2	2	1	0
06-09-2015 02:14:26	1	1	1	0	0
09-09-2015 02:27:37		0		0	0
09-09-2015 02:27:57		0		0	0
09-09-2015 02:28:18	1	1	1	0	0
09-09-2015 02:28:39		0		0	0
29	**38**	**39**	**34**	**4**	**5**

Figure 4-2. Left: The 29 commercial Radarsat-2 images (blue squares) with the correlated ships (34 green stars) and the non-reporting ships (4 red stars). Right: Tabular overview per image, the left column specifying the image date and start time (UTC).

Radarsat-2 MSSR-DVWF

Figure 4-3 shows the results for the Radarsat-2 MSSR-DVWF mode images of which the ship detections were provided by DRDC. The number of detected ships is much higher, at 216 in total (average 43.2 per image). One reason for this is that the MSSR images are bigger than the ScanSAR Narrow B and Wide Fine images. But even more than that, the MSSR images have been taken in areas of higher ship density, around the islands of Mauritius and La Réunion and between Aden and Djibouti.

Another reason why these images originally had more detections was that, upon closer inspection, they were found to contain many false alarms caused by coastal features (see Figure 4-4). These have been removed, however, and are not anymore included in the detection counts and plots.

Figure 4-5 shows some details of two of the images with a high number of detections.

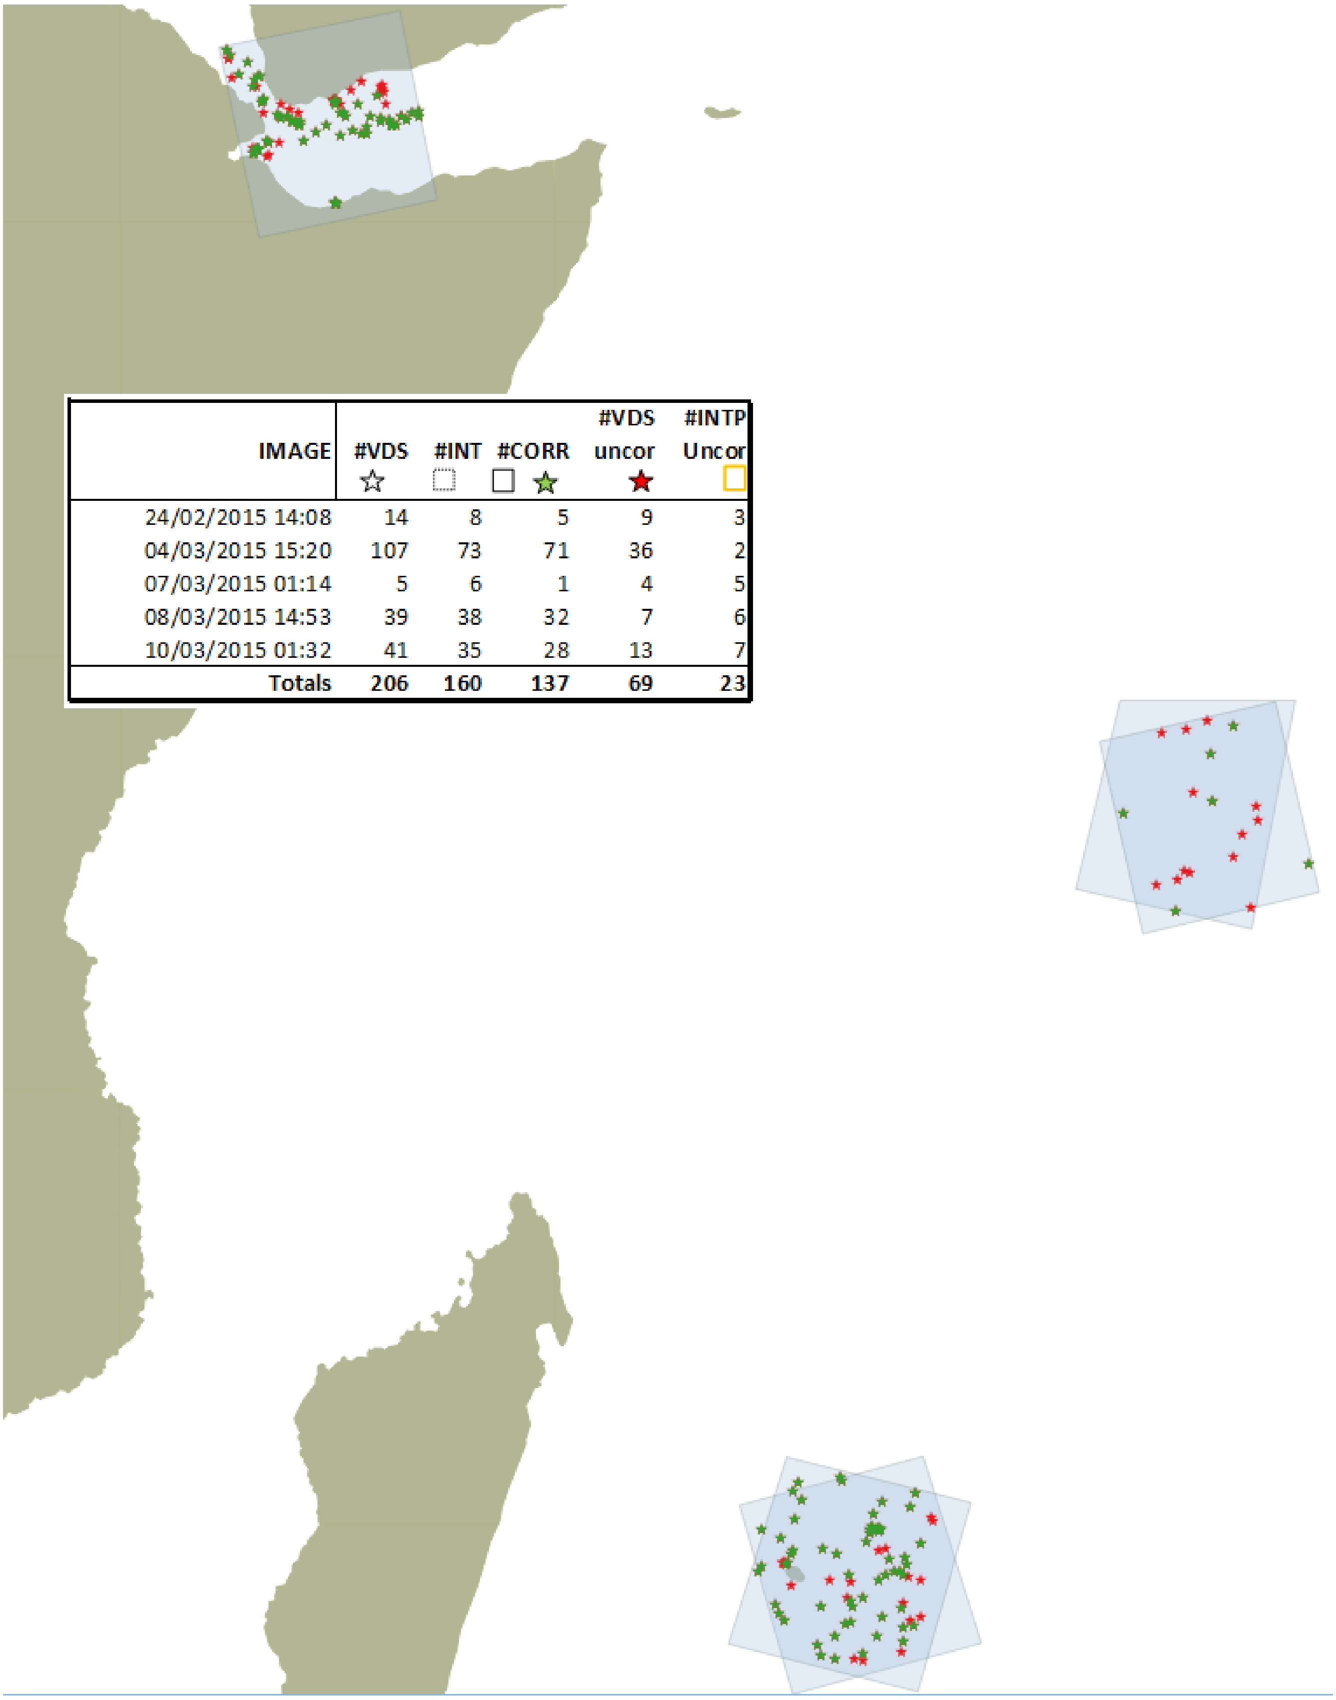

IMAGE	#VDS	#INT	#CORR	#VDS uncor	#INTP Uncor
24/02/2015 14:08	14	8	5	9	3
04/03/2015 15:20	107	73	71	36	2
07/03/2015 01:14	5	6	1	4	5
08/03/2015 14:53	39	38	32	7	6
10/03/2015 01:32	41	35	28	13	7
Totals	**206**	**160**	**137**	**69**	**23**

Figure 4-3. The correlation results of the five Radarsat-2 images in MSSR-DVWF mode of which the ship detections were provided by DRDC. A total of 206 targets were detected, and compared with 160 Interpolations, which led to 137 correlated detections (green stars) and 69 uncorrelated ones (red stars, non-reporting ships).

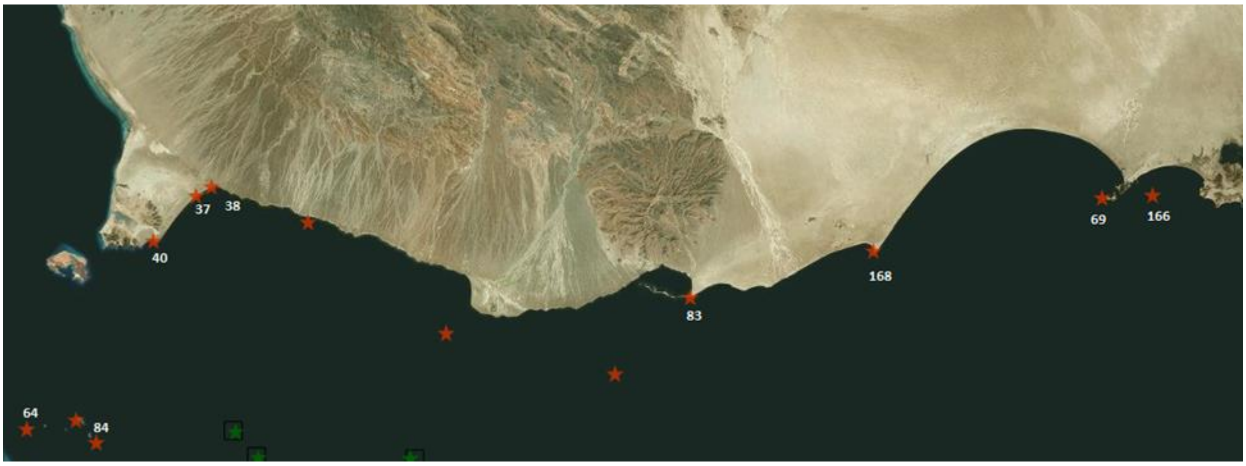

Figure 4-4. Detections (stars) in the Radarsat-2 MSSR-DVWF image of the next figure (left side), overlaid over an optical satellite image background. The detections that touch the coast are attributed to false alarms due to features on the coast. These have been removed from the count of correlated / non-correlated ships. Optical background image from ESRI World Imagery.

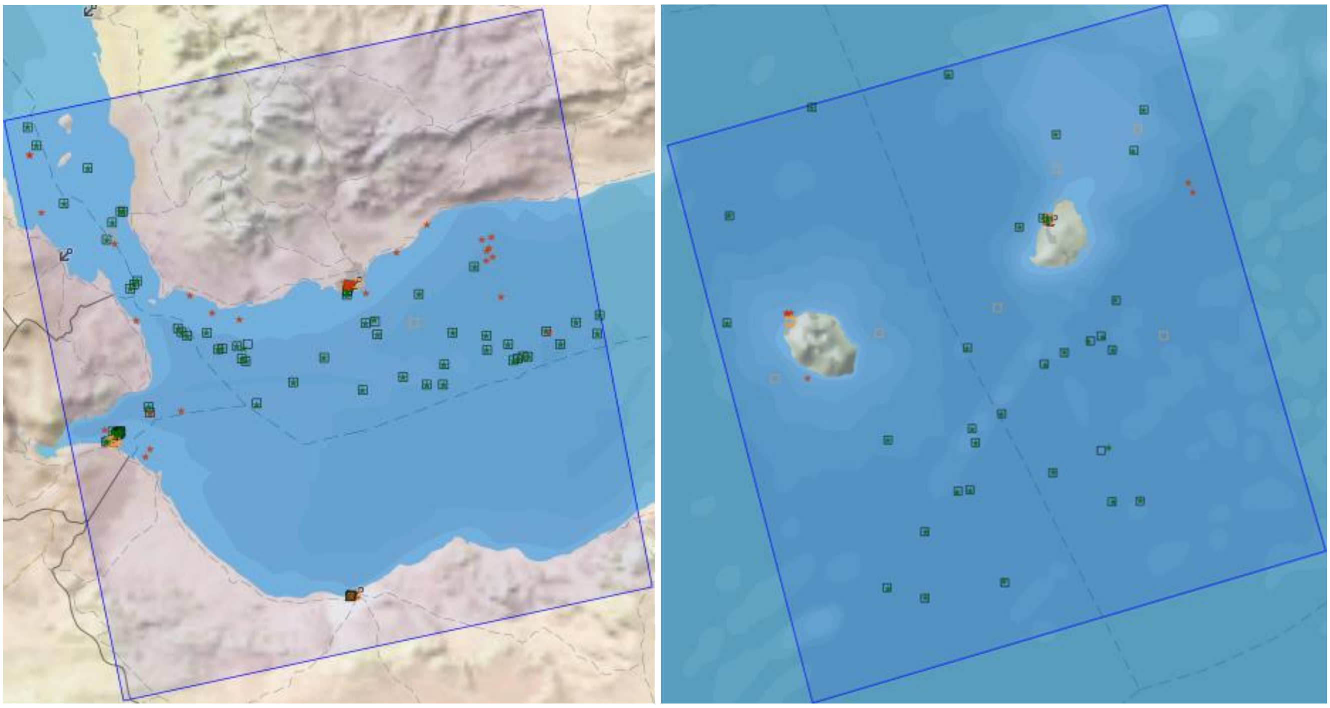

Figure 4-5. Correlations for two of the Radarsat-2 MSSR-DVWF images that have many targets. Left, Aden and Djibouti. Right, La Réunion and Mauritius. Correlated VDS are green stars, correlated Interpolations are black squares, uncorrelated VDS are red stars, uncorrelated Interpolations are yellow squares. Map background from Natural Earth.

Sentinel-1 IW manually verified

The results of the ship detection and correlation of the manually verified Sentinel-1 images on the Somali coast are shown in Figure 4-6. In the 12 images, a total of 28 VDS targets are detected (average 2.3 target per image). These are compared with 10 Interpolated reporting ship positions. This leads to 9 correlated targets and 19 uncorrelated ones.

Many detected targets are very close to the coast. A closer inspection has shown that several of them are in fact ship wrecks. Two examples are shown in Figure 4-7. These are not counted under the 28 VDS targets.

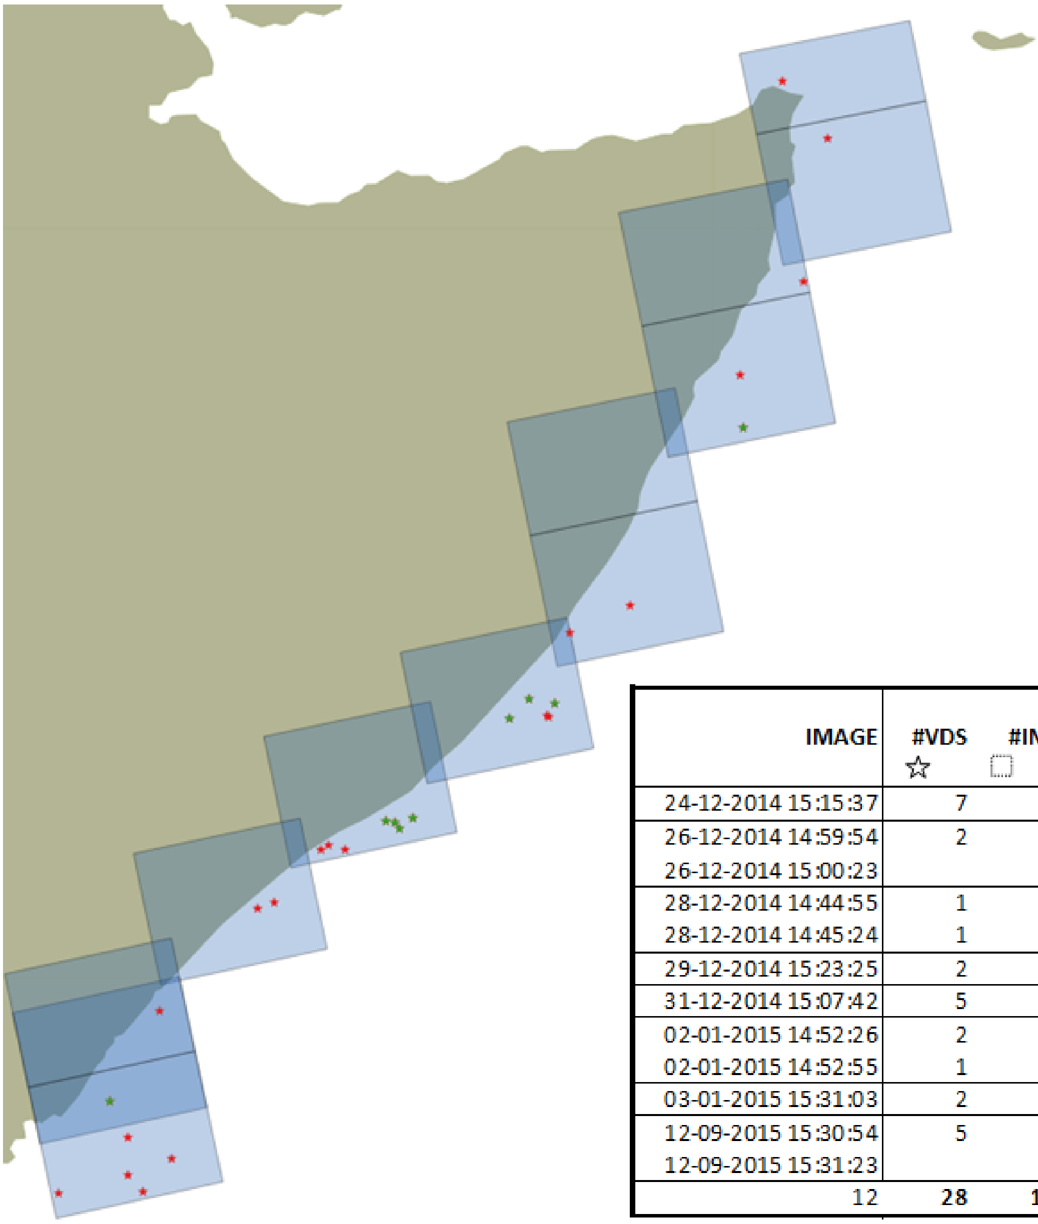

IMAGE	#VDS	#INT	#CORR	#VDS uncor	#INTP Uncor
24-12-2014 15:15:37	7	4	4	3	0
26-12-2014 14:59:54	2			2	0
26-12-2014 15:00:23				0	0
28-12-2014 14:44:55	1			1	0
28-12-2014 14:45:24	1			1	0
29-12-2014 15:23:25	2			2	0
31-12-2014 15:07:42	5	3	3	2	0
02-01-2015 14:52:26	2	1	1	1	0
02-01-2015 14:52:55	1			1	0
03-01-2015 15:31:03	2	1	1	1	0
12-09-2015 15:30:54	5			5	0
12-09-2015 15:31:23		1		0	1
12	**28**	**10**	**9**	**19**	**1**

Figure 4-6. Twelve Sentinel-1 images over the Somali coast (the lowest ones extending over the Kenya coast). Correlated targets are green stars (9) and uncorrelated targets are red stars (19).

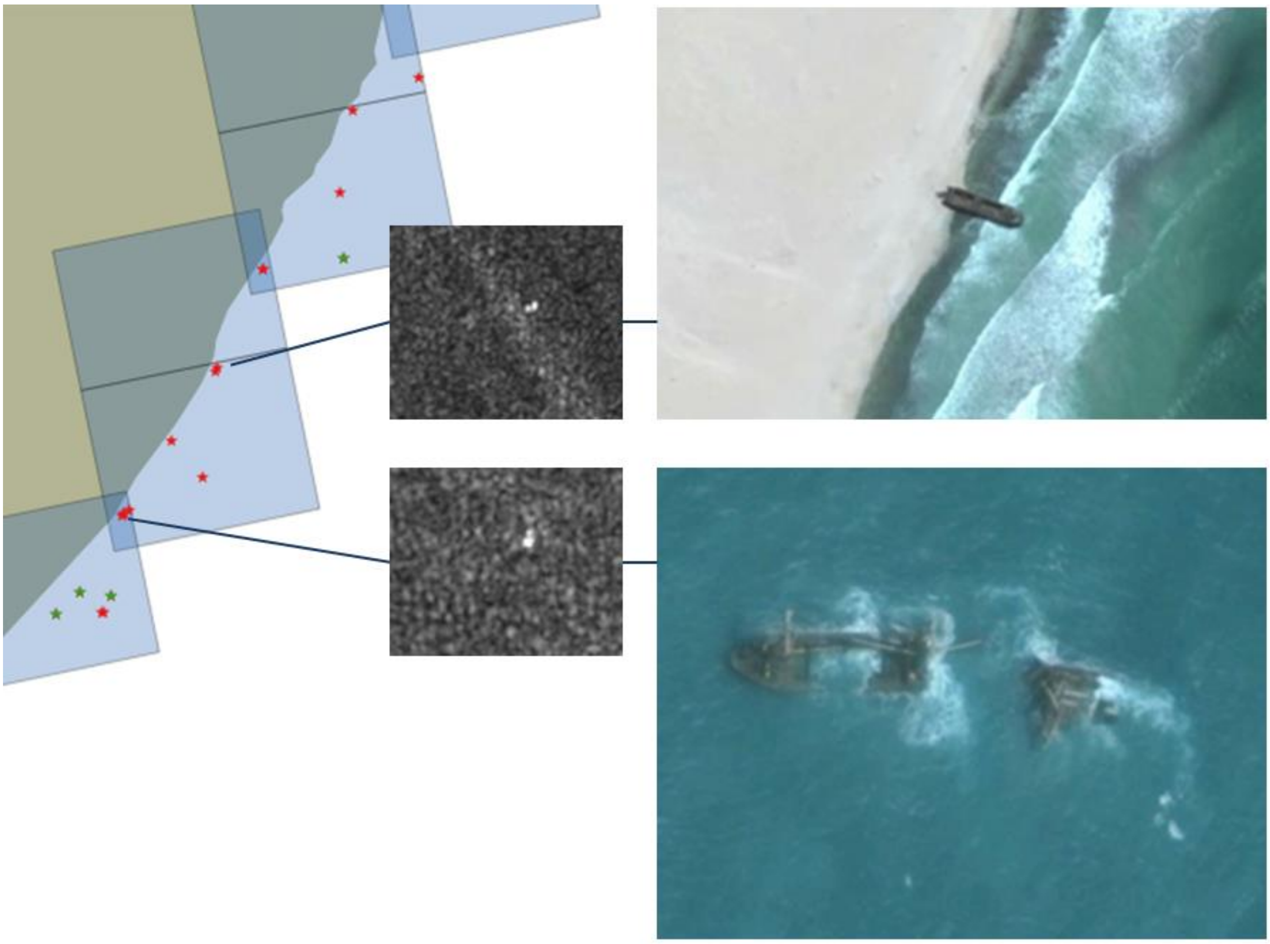

Figure 4-7. Two examples of VDS detections close to the Somali coast that are found to be ship wrecks. The map on the left is a part of the previous figure, except that here the detections near the coast are still included as red stars. The two b/w image chips in the middle are extracts from the Sentinel-1 image around the two detections in question. The two optical images on those locations are taken from Google Earth.

Overall numbers of correlation and non-reporting ships for the verified images

The results of all verified SAR images are put together inTable 4-1. This compilation brings out the variation in the number of detections per unit area (5^{th} column), and in the fraction of uncorrelated detections (9^{th} column). The Radarsat-2 SNB and Wide Fine images were taken in relatively empty areas, and the fraction uncorrelated detections (non-reporting ships) is low, 0-12 %. The Sentinel-1 images over the coast of Somalia have a high number of uncorrelated detections, even after the detections very close to the coast (some of which are confirmed wrecks) have been removed. Finally, the Radarsat-2 MSSR images were taken over areas of much higher ship density, where also an increased presence of smaller, non-reporting ships may be expected. The fraction of non-reporting ships of 33 % seems reasonable there. The overall fraction of non-reporting ships, taking all data together, is 34 %.

Sensor	**Mode** **Swath Resolution**	**# Imag**	**# VDS**	**# VDS /** **10^4 km^2**	**# Interpol**	**# Corr**	**# VDS uncorr**	**% VDS uncorr**	**# Interpol uncorr**
Radarsat-2	SNB 300 km 40 m	21	33	0.17	34	29	4	12 %	5
Radarsat-2	Wide Fine 150 km 8 m	8	5	0.28	5	5	0	0 %	0
Radarsat-2	MSSR-DVWF 450 km 40 m	5	206	2.03	160	137	69	33 %	23
Sentinel-1	IW 250 km 20 m	12	28	0.37	10	9	19	68 %	1
Overall		46	272	0.71	209	180	92	34 %	29

Table 4-1. Ship detection and correlation results from all verified satellite SAR images.

Sentinel-1 unverified

The results of the ship detection and correlation of all Sentinel-1 images, analysed in batch mode without manual verification, are shown in Figure 4-8. A total of 18,000 targets were detected; 14,000 of which inside the AOI. Of those targets, about 40 % could not be correlated with known / reporting ships. This number likely overestimates the percentage of non-reporting ships, because the fully automatic, unverified results are bound to contain many false alarms. But in order of magnitude it compares with the manually verified overall 34 %.

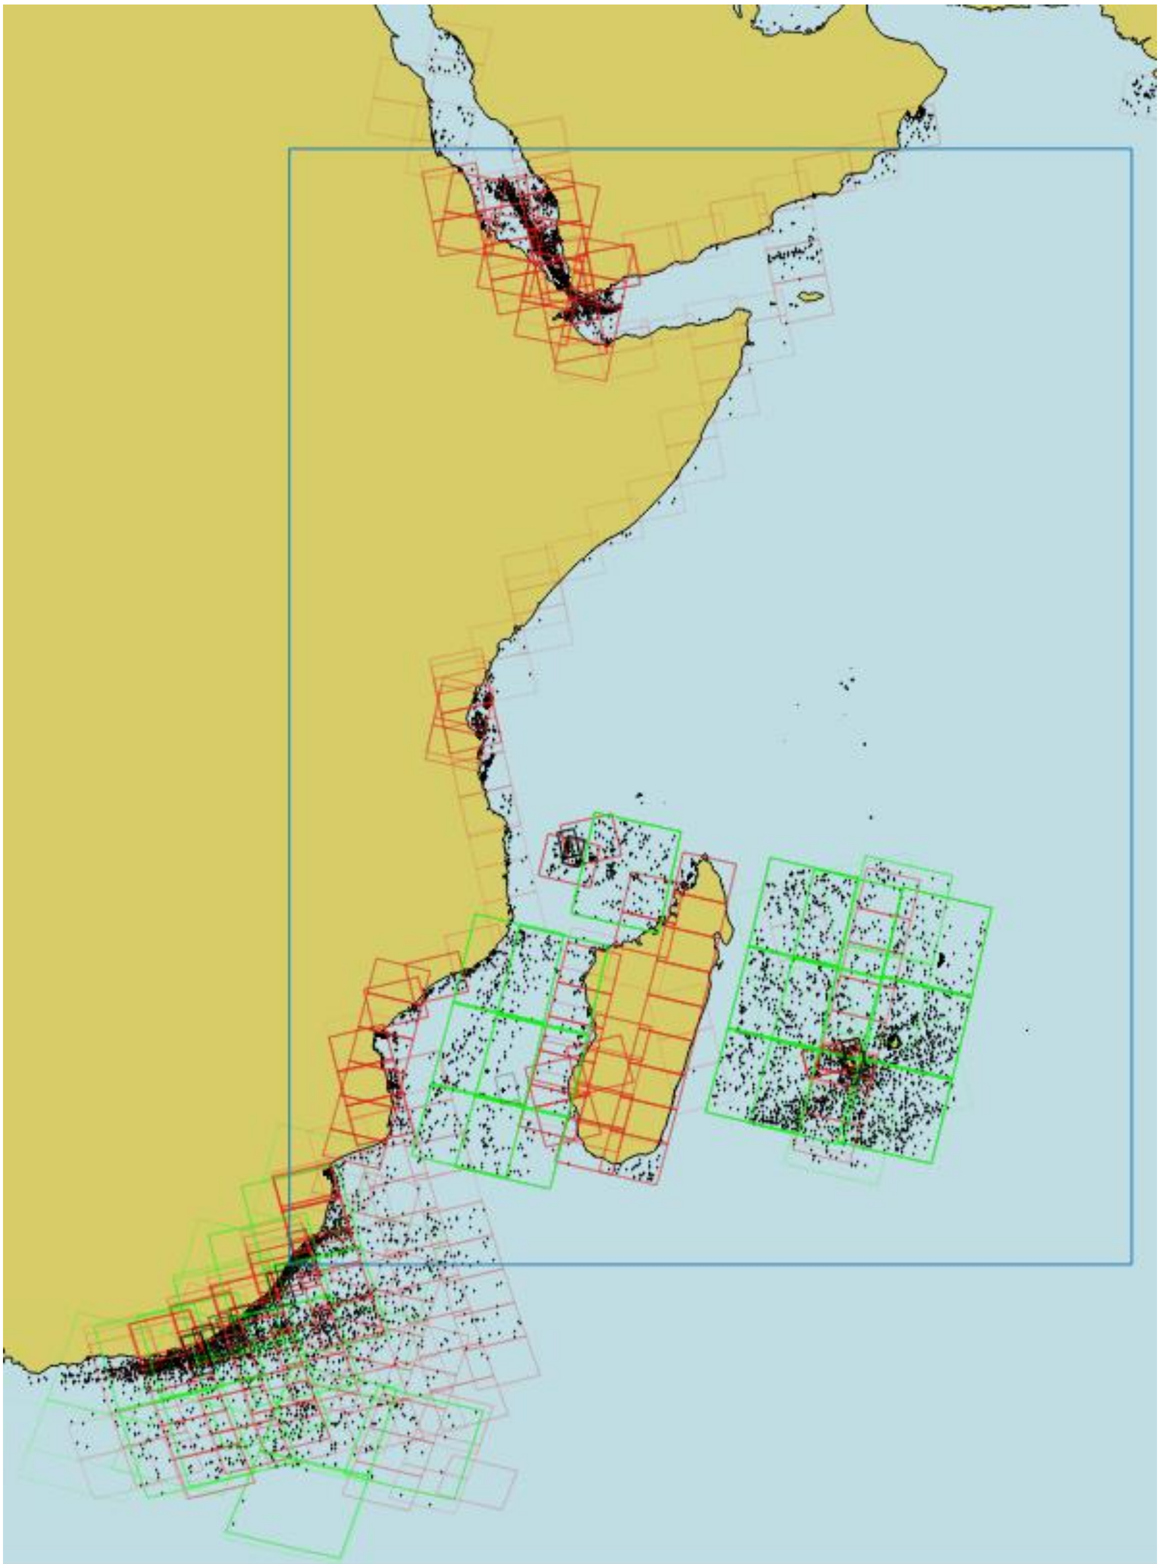

Figure 4-8. The outlines (squares) of all Sentinel-1 images downloaded from the ESA archive up to 15 Sep 2015, and the targets (dots) detected in them. A green outline means that the image has been acquired only once; a red outline means that several images over that location have been collected. The large blue box is the PMAR-MASE AOI.

Oil pollution

It is well known that SAR images of the sea surface are suitable to reveal oil spills. In one Sentinel-1 image, a dark trail is seen behind a ship that has the characteristics of an oil slick (see Figure 4-9). At the tip of the dark oil trail, a bright spot is seen that signifies a ship. Comparing with the ship reporting data, there is a reporting ship whose position matches that of the bright spot and whose track closely matches the oil trail. The oil trail has drifted slightly off the ship track due to the effect of ocean currents and wind. The Sentinel-1 image in question is of the Extra Wide (EW) mode that has wider swath but lower resolution than the IW mode images that were used for the ship detection.

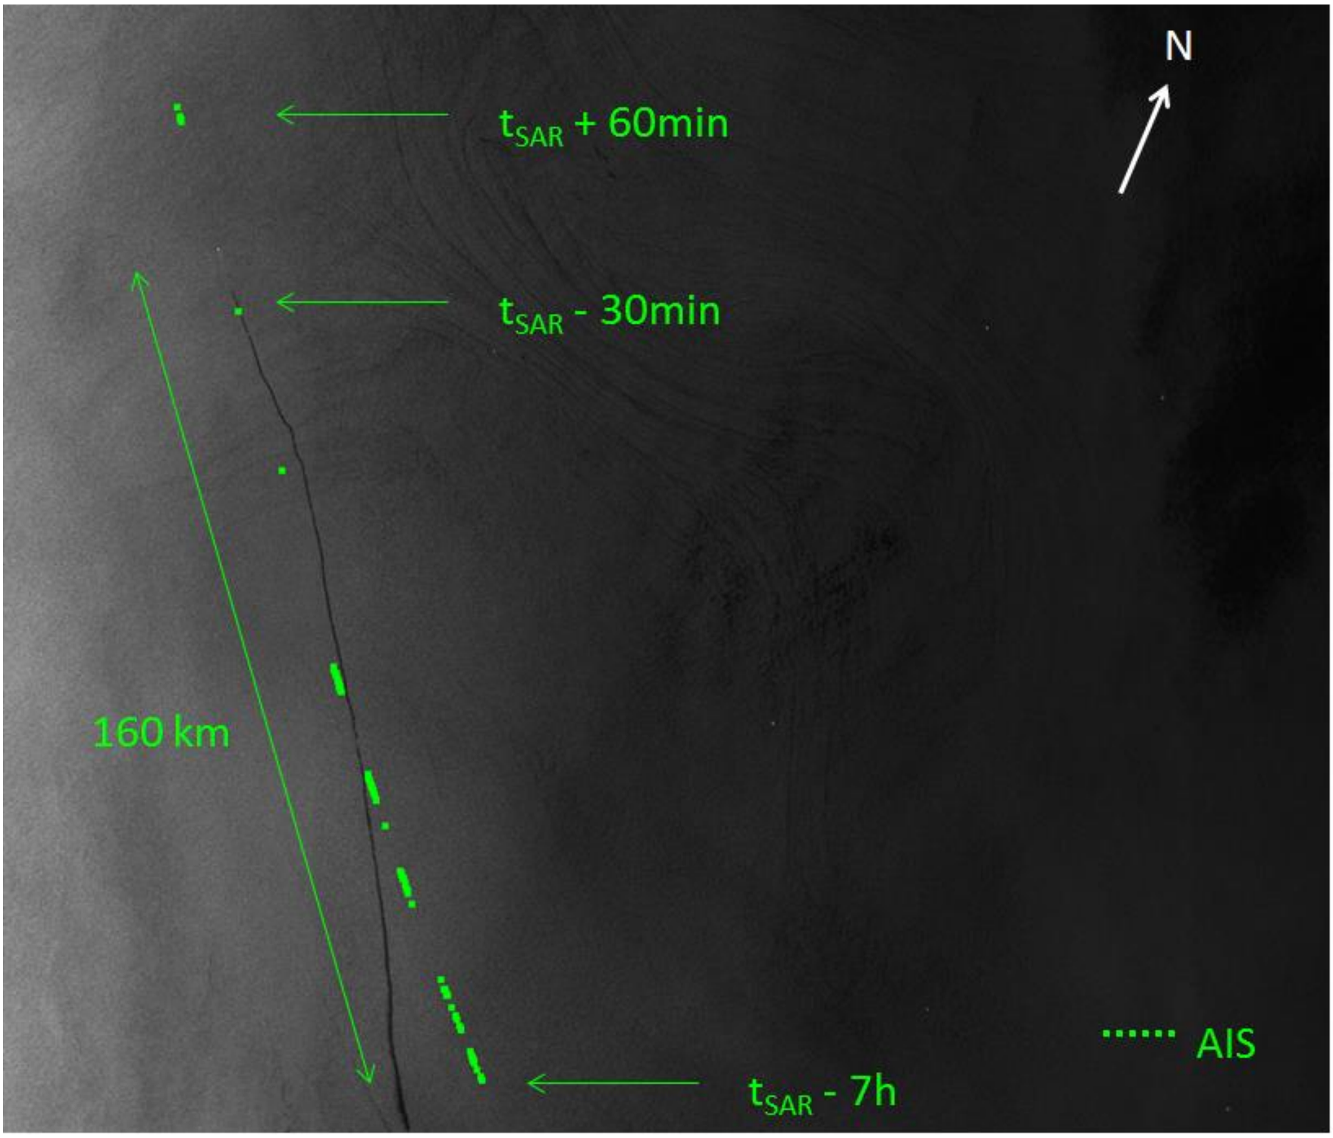

Figure 4-9. Part of a Sentinel-1 EW-mode image off Mozambique on 28 Aug 2015, with in green the AIS positions of a ship overlaid. The dark trail is likely an oil spill. The ship is seen as a bright point at the tip of the trail. Sentinel-1 image © Copernicus 2015.

Acknowledgments

LRIT data were obtained courtesy of the National Competent Authorities of the Flags that participate in the EU LRIT Data Centre, with the help of EMSA (European Maritime Safety Agency).

Satellite AIS data from NORAIS, AISSat-1 and AISSat-2 were obtained courtesy of the Norwegian Coastal Administration and the Norwegian Defence Research Establishment (FFI).

MSSIS data were obtained courtesy of the Volpe Center of the U.S. Department of Transportation, the U.S. Navy and SPAWAR.

Ship detections from RADARSAT-2 MSSR mode images were provided by Paris W. Vachon of Defence Research and Development Canada.

A coastline data base of OpenStreetMap, © OpenStreetMap contributors, was used.

Map backgrounds were used of Google Earth; of Environmental Systems Research Institute (ESRI – 'World Ocean', 'World Imagery' and 'World Topographical'); of TRANSAS Web Map Services (Nautical Charts); and of Natural Earth (free vector and raster map data @ naturalearthdata.com).

Sentinel-1 images are © Copernicus 2014-2015.

Radarsat-2 images are © MDA 2015 and were purchased from MDA / e-Geos.

Satellite AIS data were purchased from exactEarth, Orbcomm / LuxSpace and SpaceQuest.

European Commission
EUR 27612 EN – Joint Research Centre – Institute for the Protection and Security of the Citizen

Title: Maritime Awareness Systems Performance in the Western Indian Ocean 2014-2015

Author(s): Harm Greidanus, Marlene Alvarez, Vincenzo Gammieri, Carlos Santamaria,
Alfredo Alessandrini, Pietro Argentieri

Luxembourg: Publications Office of the European Union

2015 – 49 pp. – 21.0 x 29.7 cm

EUR – Scientific and Technical Research series – ISSN 1831-9424

ISBN 978-92-79-54033-2

doi:10.2788/420868

Abstract

The PMAR-MASE project aimed to provide practical experience to authorities in the Eastern-Southern Africa / Indian Ocean region on wide-area maritime awareness. Data from automatic ship reporting systems (AIS and LRIT) have been received during one year over a large Area of Interest (AOI) in the Western Indian Ocean. The data were used to construct the Maritime Situational Picture (MSP) over the AOI, i.e. the map with the current positions of all reporting ships. The MSP has been served continuously in real time to two operational users in the region, between 15 Sep 2014 and 15 Sep 2015. The data have also been used to make monthly ship traffic density maps. Finally, a number of satellite images have been collected to probe for the presence of ships that are not seen on the reporting systems.

This report assesses the quality (completeness and timeliness) of the ship reporting data and the resulting MSP; it presents the derived monthly ship density maps, split between slow moving and fast moving traffic in order to understand ship activity; and it present results on the fraction of non-reporting ships based on the analysis of the satellite images.

The ship reporting data that were used are AIS data received by up to 16 satellites of four different operators, AIS data from various coastal networks in the region, and LRIT data from the Flags that use the EU LRIT Data Centre ("EU Flags"). Fusing the data from these systems – that amounts to more than 566,000 incoming ship position messages per day (figures are given here for the month of June 2015) – it is found that on average, a total of 1,522 ships are seen in the AOI on a daily basis. Of those, 11 (0.7 %) are only seen by LRIT, in spite of the fact that LRIT contributes only a tiny 614 position reports per day (0.1 %) on average.

A sampling of the non-reporting ship traffic by satellite imaging has shown that 34 % of the ships detected in satellite radar images are not seen in AIS or LRIT. However, this fraction varies much from one location to another; among the limited locations sampled, it is highest off the Somali coast.

LB-NA-27612-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:10.2788/420868

ISBN 978-92-79-54033-2