JRC Technical Notes

IMAGE-2006 Mosaic: Automatic Seam Line Delineation

Pierre Soille and Conrad Bielski

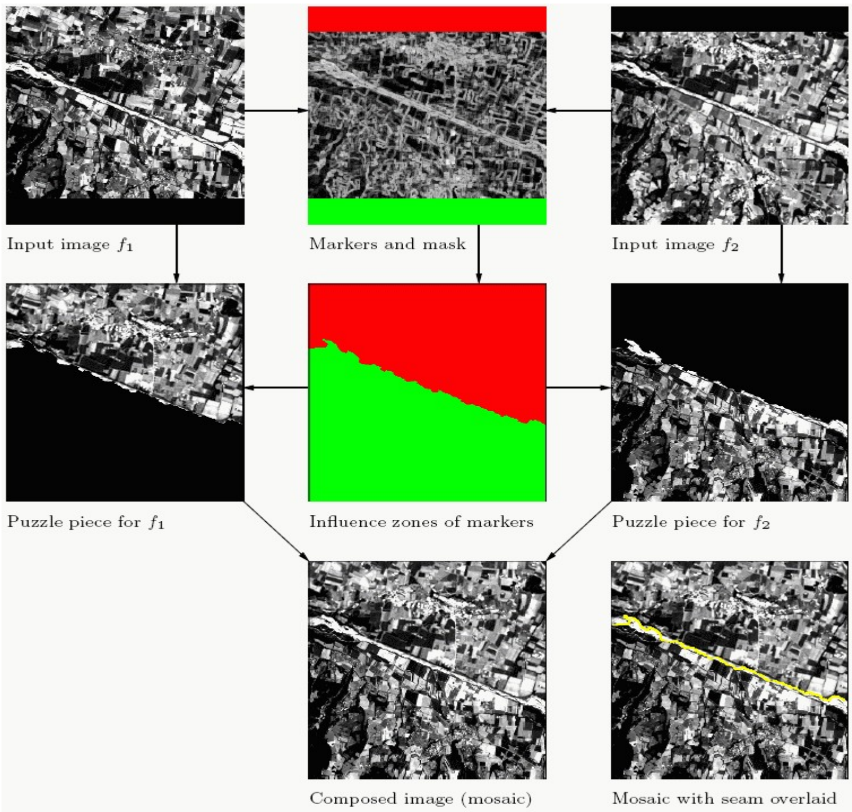

EUR 23636 EN

JRC
EUROPEAN COMMISSION
Institute for Environment and Sustainability

The mission of the Institute for Environment and Sustainability is to provide scientific-technical support to the European Union's Policies for the protection and sustainable development of the European and global environment.

European Commission
Joint Research Centre
Institute for Environment and Sustainability
Spatial Data Infrastructures Unit

Contact information
Address: via E. Fermi 2749, I-21027 Ispra (Italy)
E-mail: Pierre.Soille@jrc.ec.europa.eu
Tel.: int+39-0332 785 068
Fax: int+39-0332 786 325

http://ies.jrc.ec.europa.eu/
http://www.jrc.ec.europa.eu/

JRC 49563

EUR 23636 EN
ISBN 978-92-79-20957-4
ISSN 1831-9424
doi:10.2788/50299

Luxembourg: Publications Office of the European Union, 2011

IMAGE-2006 Mosaic: Automatic Seam Line Delineation

Pierre Soille and Conrad Bielski
COSIN Action
IES-Spatial Data Infrastructures Unit

December 23, 2008

Abstract

This report details the methodology developed for automatically composing the IMAGE-2006 imagery. The method is based on the automatic definition of seam lines in the regions where two or more images overlap. This is achieved thanks to morphological image compositing. A parallel algorithm enables the timely mosaicing of large data sets on a grid engine.

Contents

1 Introduction

To construct an image of larger field of view than what could be obtained with a single image acquisition, two or more image acquisitions need to be combined. The resulting image is called a mosaic in the sense that it is made of fragments of different images. For example, with IRS or SPOT imagery, a country such as Belgium cannot be covered by a single image. In addition, if one desires cloud cover not to exceed a given percentage, some regions may need to be covered more than once. Figure 1 shows that 15 IRS and SPOT scenes were necessary in the first coverage of IMAGE-2006 to cover the territory of Belgium while securing that cloud cover does not exceed 5%. In this figure, the images have been loaded in an arbitrary order so that the value of any given pixel corresponds to the value given by the last image providing values for this pixel. That is, one would obtain the same result by superimposing printed copies of the images using the same order as the one used for loading the images. Because the image have been orthorectified and projected beforehand, data values occupy only a portion of the coloured rectangle frame of each image (small frames for SPOT, large frames for IRS images).

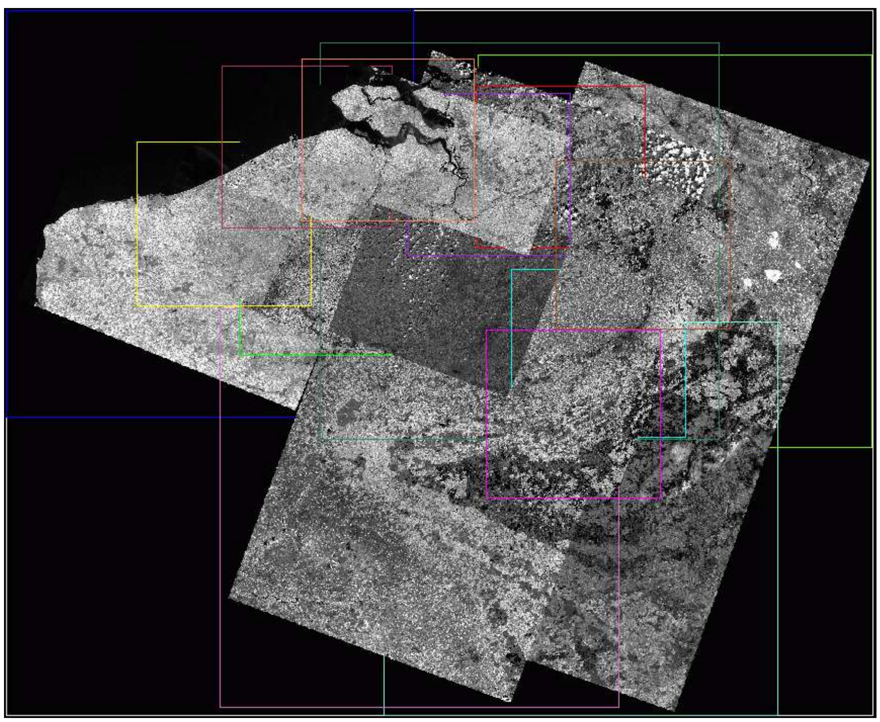

Figure 1: The 14 scenes delivered for Belgium composed according to the (arbitrary) order in which they have been loaded. The frames indicate the relative positions of these overlapping scenes. The displayed data corresponds to raw digital numbers of band 4.

The mosaicing of a set of spatially overlapping images comes down to determine a unique value for each pixel that is covered by more than one image. Beyond trivial (and arbitrary) methods such as the one used for creating Fig. 1, this can be achieved by pixel or region based methods. With pixel based methods, the output value of a pixel is determined with the sole knowledge of the input values of this pixel. By contrast, with region based methods, the output value of a pixel depends on some spatial relationships with other pixels.

Among pixel based methods, those relying on a selection procedure are often performing better than those relying on linear combinations of the available values. A simple per pixel selection method is to define as output value of a pixel the minimum or maximum value of the available input values. This approach is of limited interest for creating a seamless mosaic because any variation of reflectance values due to atmospherical or seasonal effects leads to blocky structures in the output mosaic. Still, the point-wise minimum and maximum composition rules are of practical interest. For instance, the minimum composition rule generates a mosaic showing the least possible amount of cloud cover because clouds have usually higher reflectance values than most land cover classes in the optical domain. Also, because atmospheric effects such as haze and thin clouds increase the top-of-atmosphere (TOA) reflectance values, the minimum mosaic tends to select scenes that are less corrupted by atmospheric effects. Similarly, shadows and water bodies as well as all objects having low reflectance values are highlighted in the minimum mosaic. The maximum mosaic, on the other hand, reveals all clouds occurring in all overlapping images and tend to give precedence to TOA reflectance values most affected by atmospheric conditions. Nevertheless, the minimum and maximum mosaics are useful for visual control purposes. For instance, any cloudy area appearing in the minimum mosaic can be considered as an area permanently covered by clouds given the available imagery, i.e., it will appear in any other mosaic based on the same set of input images. The maximum mosaic is also useful for visually assessing the performance of a cloud detection algorithm: the union of all cloud masks extracted for each individual image should match all clouds visible in the maximum mosaic. This principle was used for creating the cloud galleries presented in [6]. All effects described above can be observed in the point-wise minimum and maximum mosaics displayed in Fig. 2. In this figure, the mosaics have been rendered using the Top of Atmosphere (TOA) reflectance values of the 4th, 3rd, and 1st bands for, respectively, the red, green, and blue channels.

Region based methods rely on a partition of the domains of overlap. Each segment of the partition is associated with an index value indicating from which image the values of the mosaic must be selected from within this segment. Hence, each segment with its associated index value can be considered as a decision region. Because decisions are taken at the level of a region, region based methods lead to mosaics showing a higher spatial coherence than pixel based methods.

The IMAGE-2006 mosaics were produced thanks to a region based method called morphological image compositing [5]. This method partitions automatically the domains of overlap. It is best introduced in the simplest case consisting of two partly overlapping images that need to be composed to form a unique image. In this case, the problem comes down to partitioning their domain of overlap into two regions indicating which image should be selected within these regions. The interface between these two regions define a cut or seam line because on one side of this line an image is used while the other one is used on the other side. Therefore, the seam line should be placed so as to diminish as much as possible the visual detection of the transition from one image to

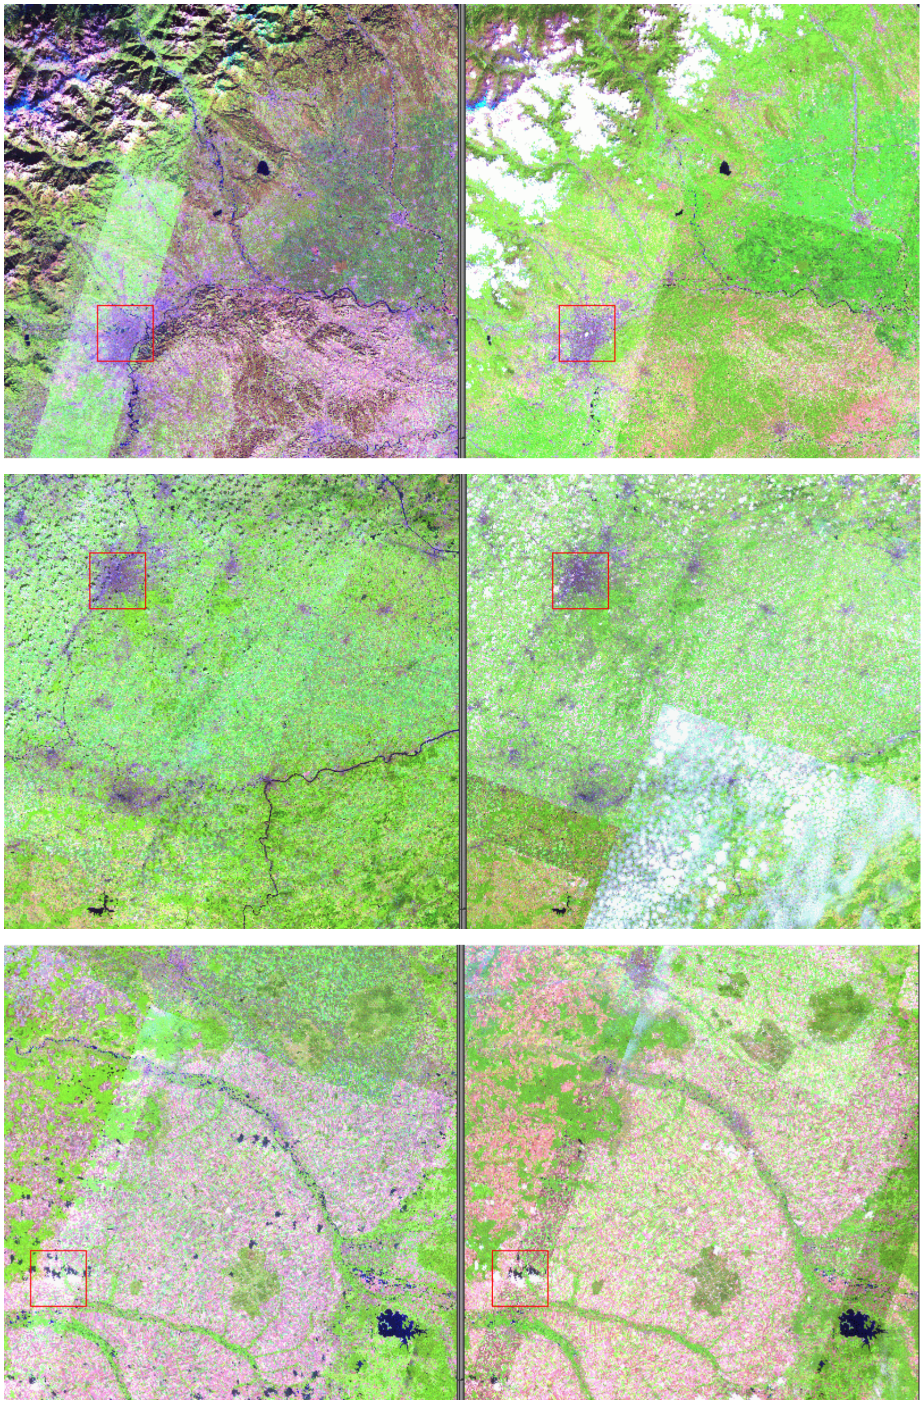

Figure 2: Comparison between the point-wise minimum (left) and the point-wise maximum (right) mosaics computed on the basis of all imagery of the first coverage of IMAGE-2006. The areas framed by a red box represent Turin (top), Brussels (middle), and a permanent cloud South of the Seine river (bottom).

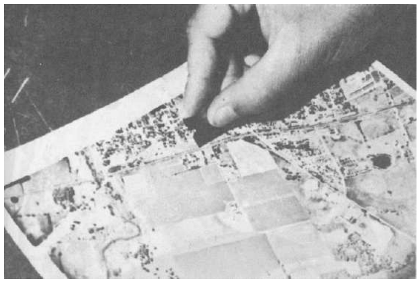

Figure 3: Making the razor cut or seam line by following lines of definite tone demarcation. Source: Wolf [10, p. 241].

another in the resulting mosaic. That is, cuts revealing noticeable mismatch in tones and position should be avoided. In the early days of image mosaicing, this was achieved by making a razor cut following terrain structures such as roads, railroads, edges of fields, or other lines of definite tone demarcation [10, p. 242]. This old-fashioned method is illustrated in Fig. 3 and was still used in the first continental mosaic, namely the first Landsat mosaic of North America in 1976 [2] consisting of 569 Landsat images pieced together by technicians during 4 months. Morphological image compositing [5] mimics human operators by automatically defining seam lines following the most salient object boundaries occurring within the domain of overlap. This is achieved thanks to a region growing process [11]. The growth is initiated by defining two seeds corresponding to the regions where there is no overlap (one seed for each image). The actual growth proceeds within the domain of overlap until it is completely covered. The domains reached by each seed indicate which image should be used within these domains. To ensure that the boundary of these domains matches salient image structures, the growth is constrained by an image whose intensity values are proportional to variations in tone within a neighbourhood centred at each pixel of the overlapping domain. That is, when variations are high, the speed of the growth process is decreasing and vice versa. The propagation is computed in such a way that growing regions actually meet along lines of high tonal variations corresponding to the most salient image structures. Morphological image compositing can be extended to an arbitrary number of images while minimising the occurrence of undesirable objects such as clouds.

The goal of this report is to present a detailed description of the methods developed to this aim. It is organised as follows. Section 2 presents a formal description of morphological compositing for the direct (in place) processing of the images to compose. The direct processing of a large number of images is rapidly limited by the available computer random access memory. Section 3 shows that this limitation can be addressed thanks to sequential processing. For a very large data sets such as the IMAGE-2006 imagery, sequential processing would require too much time. Section 4 addresses this issue by introducing a parallel algorithm enabling the distribution of the mosaicing task to a series of computers. Concluding remarks are presented in Sec. 5.

2 Direct processing

Mosaicing consists in merging a sequence of spatially overlapping images so as to create a unique image whose spatial extent is equal to the smallest rectangle enclosing all input images. Direct processing assumes that enough memory

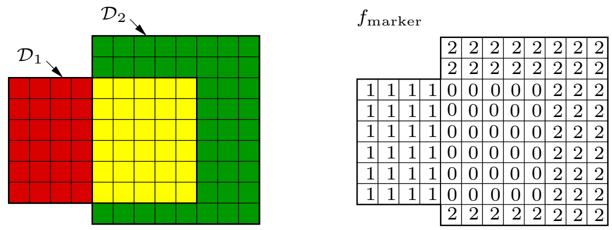

Figure 4: Left: definition domains $\mathcal{D}_1$ and $\mathcal{D}_2$ of two overlapping images. The yellow region defines the domain of overlap (combination of the fundamental colours used for each individual definition domain). Right: the corresponding marker image f_{marker} with its numerical values.

is available to perform all computations on temporary images whose extent is equal to that of the output mosaic. The fundamental principles of morphological image compositing were first introduced in the case of direct processing. They are recalled hereafter, first for the simplest case of 2 images to compose and then for an arbitrary number of images.

2.1 Composing 2 images

Let $\mathcal{D}_1$ and $\mathcal{D}_2$ denote the definition domain of two digital images f_1 and f_2. Let assume that these images partly overlap, i.e., $\mathcal{D}_1 \cap \mathcal{D}_2 \neq \emptyset$, and $\mathcal{D}_2 \neq \mathcal{D}_1$. The goal is to create an image f whose definition domain $\mathcal{D}_f$ equals $\mathcal{D}_1 \cup \mathcal{D}_2$. Note that, in practice, this image is represented in computer memory by an image whose extent is equal to the domain of the smallest rectangle enclosing $\mathcal{D}_1 \cup \mathcal{D}_2$. The morphological composition method is based on a region growing process starting from markers and driven by a mask image.

2.1.1 Marker image

The marker image whose definition domain is equal to $\mathcal{D}_f$ is defined as follows for each pixel $\mathbf{x}$:

$$\left[f_{\text{marker}}\right](\mathbf{x}) = \begin{cases} 1, & \text{if } \mathbf{x} \in \mathcal{D}_1 \text{ and } \mathbf{x} \notin \mathcal{D}_2, \\ 2, & \text{if } \mathbf{x} \in \mathcal{D}_2 \text{ and } \mathbf{x} \notin \mathcal{D}_1, \\ 0, & \text{otherwise (i.e., no marker).} \end{cases}$$

The definition domains of two overlapping images and the corresponding marker image are illustrated in Fig. 4. That is, when creating the mosaic, for the pixels with value 1 (resp. 2) of the marker image, there is no choice but select the value of the image f_1 (resp. f_2) since only one value is available for these pixels. The problems reduces to determine from which image the values of the 0-valued pixels should originate. Indeed, two input values are available for these pixels. Two overlapping images and their corresponding marker image are illustrated in Fig. 5. Note the discrepancies in land cover use in the overlapping domain (temporal difference of 13 months between the two images). It follows that any mosaicing whose seam lines would not follow the boundaries of objects visible in *both* images would introduce undesirable discontinuities.

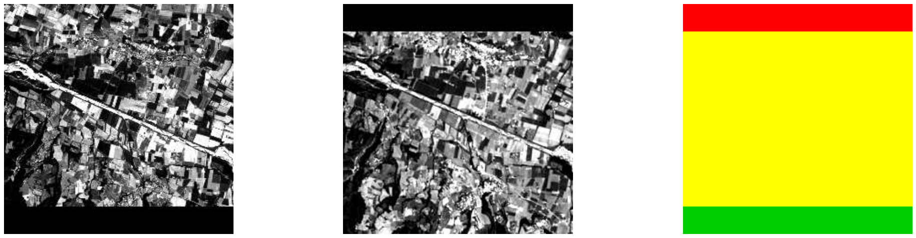

Figure 5: Two overlapping images and the corresponding marker image. Left: excerpt from `20060818-1041 IL3 FR 3963524448-BB`. Middle: excerpt from `20050715-1041 SP5 FR 3931424292-DB`. Right: corresponding marker image.

2.1.2 Mask image

The mask image is used to direct the growth of the regions where there is no overlap (i.e., regions defined with non-zero values in the marker image described in the previous section). By definition, the growth of the markers must be restricted to the domain of overlap so that the mask image needs to be defined solely within this domain. If it would be simply defined as an image with constant values, the growth would be isotropic and lead to regions corresponding to the (geodesic) influence zones [4] of the markers within the domain of overlap. That is, in the example of Fig. 5, the red marker would grow downwards and the green one upwards so that they would fill the rectangular overlapping domain (yellow) and meet exactly along its horizontal axis. Obviously, this would generate a visible seam line in the corresponding mosaic because it would not follow the boundaries of actual image objects. This can be solved by defining a mask image highlighting the boundaries of the image objects. Because object boundaries correspond to regions of high tonal variations, they can be enhanced by computing the difference between the highest and lowest grey level values within an elementary neighbourhood centred on each pixel. This definition is known as the morphological gradient ρ within a neighbourhood B [3]. To ensure that only tonal variations occur in *both* images at the same locations, the morphological gradient of both images are computed in parallel and then combined through the point-wise minimum operation $\wedge$:

$$f_{\text{mask}} = \rho_B(f_1) \wedge \rho_B(f_2).$$

This is illustrated in Fig. 6. When computing the gradient of each individual image, pixels that are outside the data ROI of the image must be ignored to avoid spurious high gradient values along the boundaries of the data ROIs.

2.1.3 From marker and mask images to decision regions

Once the marker and mask images have been defined, the overlapping domain needs to be partitioned into two decision regions, one for each marker. This can be achieved by a morphological transformation known as the watersheds from markers [8, 9]. That is, the watersheds of the mask image are computed using the markers as attractors. The resulting catchment basins partition the domain of overlap into two influence zones (decision regions) indicating which image should be used in each zone. The resulting puzzle pieces are then combined to create the desired mosaic. All these steps are illustrated in Fig. 7.

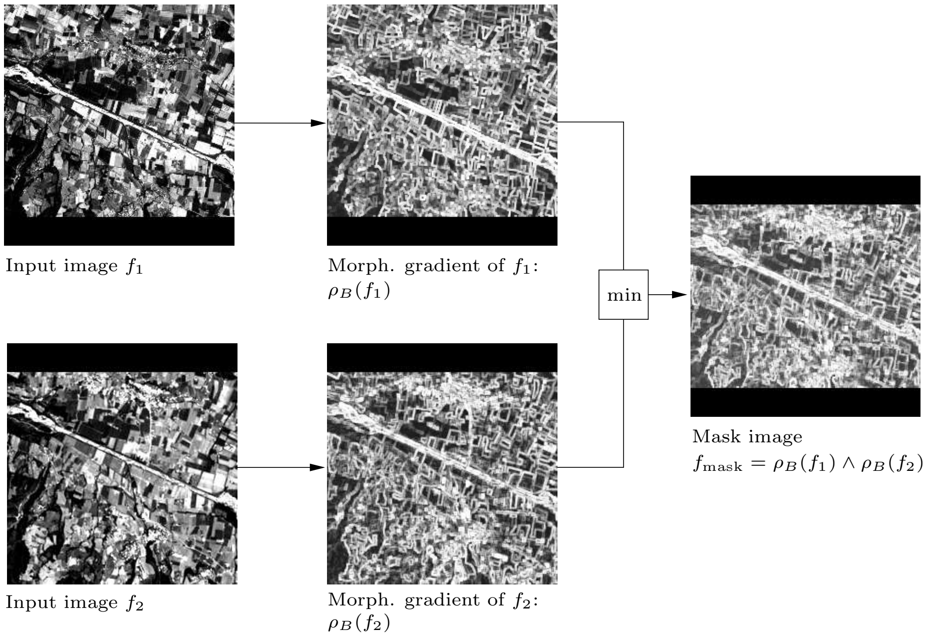

Figure 6: Generation of the mask image used for directing the growth of the markers.

In mathematical terms, composition of the input images is achieved with the following decision rule:

$$f(\mathbf{x}) = f_i(\mathbf{x}),$$

where the value of the subscript i is defined by the label value of the influence zone IZ at position $\mathbf{x}$:

$$i = \Big[\mathrm{IZ}_{f_{\text{marker}}}(f_{\text{mask}})\Big](\mathbf{x}).$$

2.1.4 Removing specific objects

The methodology can be easily extended to remove specific object such as clouds. The first step is to detected in each image all objects that are undesirable. For instance, cloud detection on SPOT and IRS imagery is detailed in [6]. The marker image is then updated to take into account the binary masks of the undesirable objects. For example, if a cloud is detected in an image, its extent defines an additional marker for the second image (from which the region growing procedure is also initiated). By doing so, provided that no cloud appears at the same location in the second image, one makes sure that the cloud will not appear in the mosaic.

Formally, denoting by M_i the binary mask of undesirable objects occurring in the image f_i, the marker image is defined as follows:

$$\Big[f_{\text{marker}}\Big](\mathbf{x}) = \begin{cases} 1, & \text{if } \mathbf{x} \in (\mathcal{D}_1 \cap \overline{\mathcal{D}_2}) \text{ or } \Big[\mathbf{x} \in (\mathcal{D}_1 \cap \mathcal{D}_2) \text{ and } \mathbf{x} \in (M_2 \cap \overline{M_1})\Big], \\ 2, & \text{if } \mathbf{x} \in (\mathcal{D}_2 \cap \overline{\mathcal{D}_1}) \text{ or } \Big[\mathbf{x} \in (\mathcal{D}_1 \cap \mathcal{D}_2) \text{ and } \mathbf{x} \in (M_1 \cap \overline{M_2})\Big], \\ 0, & \text{otherwise (i.e., no marker),} \end{cases}$$

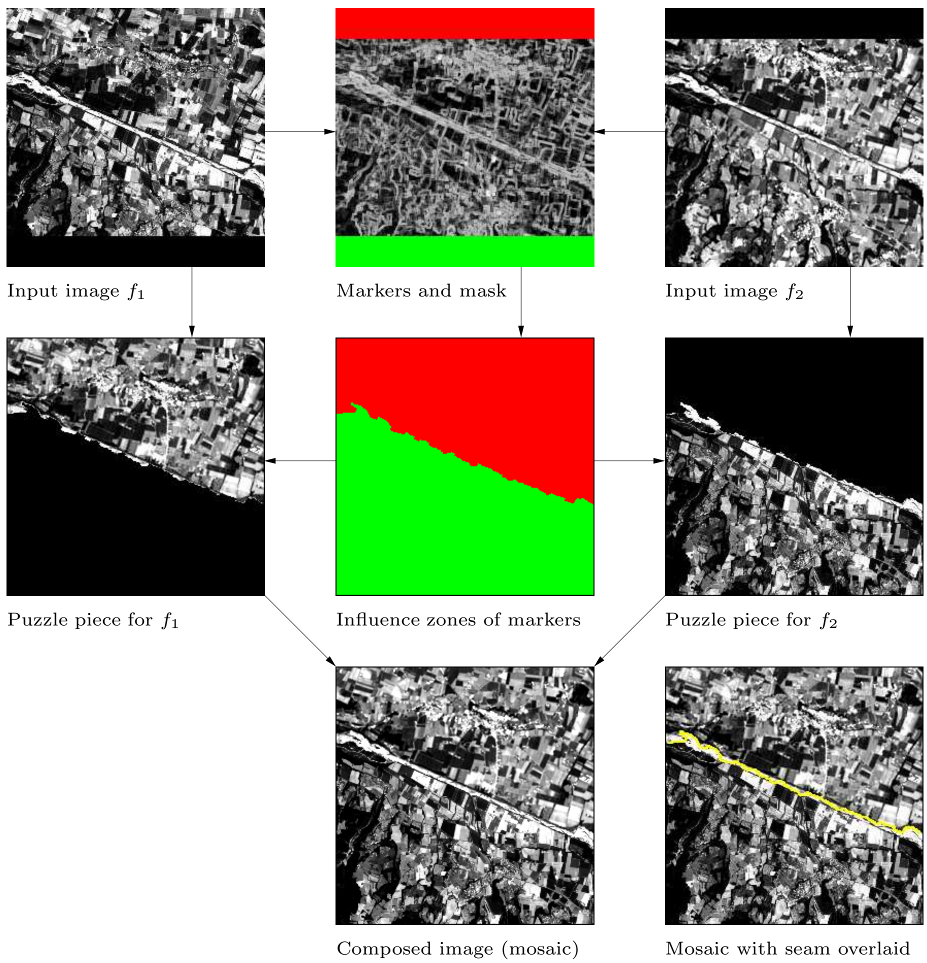

Figure 7: Morphological image compositing of two overlapping images: flow diagram (see text for details).

where the line over a set refers to the complement of this set. That is, in the case of satellite images, clouds (or any other undesirable object) occurring in an image are suppressed by considering markers of the other image at their positions (provided that clouds do not occur at the same position in the other image). The watershed-based propagation of the markers ensures that the resulting seam lines are following the boundaries of relevant image objects present in both images.

2.2 From two to an arbitrary number of images

The procedure described in Sec. 2.1 for the simplest case of 2 images can be extended to the mosaicing of an arbitrary image number n. This extension is described in the following subsections.

2.2.1 The overlap level image g

The key to the extension from 2 to n images consists in processing the domain of overlaps by increasing overlapping level. The overlap level of a given pixel indicates the number of images covering this pixel.

Let n denote the number of spatially overlapping images. Let $\mathcal{D}_f$ denote the definition domain of the composed image: $\mathcal{D}_f = \cup_i \mathcal{D}_i$. The overlap level image g indicates how many images are overlapping any given pixel $\mathbf{x}$ of $\mathcal{D}_f$:

$$g(\mathbf{x}) = \text{card}\{i \mid \mathbf{x} \in \mathcal{D}_i\}. \tag{1}$$

For example, the left diagram of Fig. 8 shows the definition domains of 3 overlapping images. Fundamental colours and their combinations have been used to highlight the domains of overlap. The g image is equal to one for the fundamental colours, 2 for the 3 compositions of 2 fundamental colours, and 3 for the central white domain where all 3 images overlap.

2.2.2 Mask image

The mask image is defined as the point-wise minimum between the morphological gradients of all input images:

$$f_{\text{mask}} = \bigwedge_i \rho_B(f_i). \tag{2}$$

By doing so, cut lines will naturally follow the the boundaries of object that are visible in *all* available images. When computing the gradient of each individual image, pixels that are outside the data ROI of the image must be ignored to avoid spurious high gradient values along the boundaries of the data ROIs.

2.2.3 Iterative procedure

When composing an arbitrary number of images, markers grow through an iterative procedure from an initial state (matching the regions where there is no overlap) to a final state where they partition the domains of overlap into a series of decision regions.

The initial marker image is analogous to that introduced for the composition of an image pair. The markers correspond to the regions where there is no

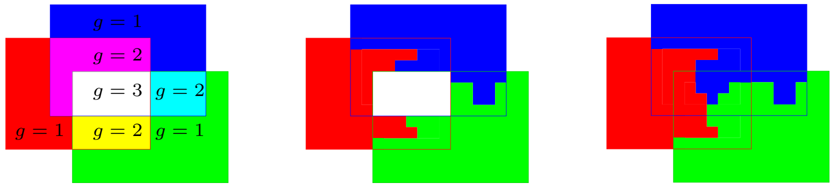

Figure 8: Three overlapping images. Left: overlap level image g. The yellow, magenta, cyan, and white regions define the domain of overlap (combination of the colours used for each individual definition domain where no o overlap occurs). Middle: after region growing in the domains where two images overlap. Right: after subsequent region growing in the domain where all 3 image overlap. Adapted from [5, Fig. 3].

overlap (red, green, and blue region in the left image of Fig. 8). They are then propagated within the regions of overlap level 2 (see middle image of Fig. 8), and then to the subsequent levels until the maximum degree of overlap is reached. In the example of Fig. 8, this occurs at level 3 (see right image). At each step, the propagation is directed by the mask image defined in the previous section.

This iterative procedure is illustrated in Fig. 8 and can be formalised as follows. Let k refers to an arbitrary iteration level ($1 \leq k < g_{\max}$). The k-th iteration processes the pixels that have on overlap level equal to k. Accordingly, we denote by $f^{(k)}$ the values of f that are defined at the end of the k-th iteration. Consequently, the definition domain of $f^{(k)}$ is equal to those pixels of the overlap level image whose values are less than or equal to k:

$$\mathcal{D}_{f^{(k)}} = \{\mathbf{x} \mid g(\mathbf{x}) \leq k\}.$$

Let $\mathcal{D}_i^{(k)}$ refer to those pixels of $\mathcal{D}_{f^{(k)}}$ whose values originate from the input image f_i:

$$\mathcal{D}_i^{(k)} = \{\mathbf{x} \in \mathcal{D}_{f^{(k)}} \mid f^{(k)}(\mathbf{x}) \leftarrow f_i(\mathbf{x})\}.$$

Initially, $k = 1$ and therefore $\mathcal{D}_i^{(1)}$ is trivially defined since there is only one image covering each pixel $\mathbf{x}$. For actual overlaps (i.e., $k > 1$, the definition of $\mathcal{D}_i^{(k)}$ is achieved thanks to the computation of the decision regions by the marker-controlled segmentation described in Sec. 2.1. The marker image at iteration $k > 1$ is defined by the decision regions already determined during the previous iteration(s):

$$f_{\text{marker}}^{(k)}(\mathbf{x}) = \begin{cases} i, & \text{if } \mathbf{x} \in \mathcal{D}_i^{(k-1)}, \\ 0, & \text{otherwise (i.e., no marker).} \end{cases}$$

The decision rule is then be formulated as follows:

$$f^{(k)}(\mathbf{x}) = f_i(\mathbf{x}) \text{ where } i = \left[\mathrm{IZ}_{f_{\text{marker}}^{(k)}}(f_{\text{mask}})\right](\mathbf{x}),$$

the propagation if the markers being restricted to $\mathcal{D}_{f^{(k)}}$. Note that during the computation of the influence zones, it is necessary to check that markers only propagates within the definition domain of the image they correspond to. The iterative procedure stops when k reaches the largest overlap level $g_{\max}$.

2.2.4 Removal of specific objects

The procedure for the removing specific objects in the case of two overlapping images (see Sec. 2.1.4) is more complex when a higher number of images overlap. Indeed, assuming that a given pixel belongs to an object detected solely in one image and that should not appear in the composed image, there is no direct selection rule to choose among the remaining $n-1$ pixels values. This problem can be solved by remapping the original image indices $i \in \{1,\dots,n\}$ to 2^{i-1} so as to produce unique values when summing two of more indices (the sum is then equivalent to a bit-wise OR operation). If an undesired object such as a cloud-shadow complex occurs in a region of overlap, one needs to create a marker whose spatial extent corresponds to this object and whose value equals the bit-wise OR operation between the remapped indices of all images overlapping this object, except the one containing this cloud-shadow complex. The resulting index is referred to as *composite* index because it represents the union of two or more image indices. For instance, in the example depicted in Fig. 8, in case a cloud-shadow complex of the second image would occur in the region where all three images overlap, a marker with a composite index value equal to 5 would be generated where this complex occurs. Indeed, the sum of the powers of 2 of the index values corresponding to the first and third images equals $2^{1-1}+2^{3-1}=5$. In the sequel, the set of indices used for generating a composite index is called its generator set. Indices referring to a unique image are called *plain* indices. By definition, plain indices are always powers of 2. A formal definition of the marker image follows:

$$f^{(k)}_{\text{marker}}(\mathbf{x}) = \begin{cases} c, & \text{if } \mathbf{x} \in \mathcal{D}_c^{(k-1)}, \\ \sum_i \{2^i \mid \mathbf{x} \in \mathcal{D}_i \text{ and } x \notin M_i\}, & \text{if } g(\mathbf{x}) = k \text{ and } \exists\, j \mid \mathbf{x} \in M_j, \\ 0, & \text{otherwise (i.e., no marker).} \end{cases}$$

$c \in \{1,\dots,2^n-1\}$ and $\mathcal{D}_c^{(k-1)} = \Big\{\mathbf{x} \in \mathcal{D}_f^{(k-1)} \;\Big|\; \big[\mathrm{IZ}_{f^{(k-1)}_{\text{marker}}}(f_{\text{mask}})\big](\mathbf{x}) = c\Big\}$. At each iteration, the composite label values of the influence zones originating from specific structures to remove must be mapped to plain label values. This is achieved by using a majority rule: each influence zone with a composite label value is scanned and is set to the plain label value of its adjacent influence zone with which it share the largest number of boundary pixels. If such a region is not found (e.g., all adjacent regions are labelled with composite labels) then the smallest plain label forming the composite label is selected.

3 Sequential processing

Sequential processing enables the processing of one image at a time without the need to hold in random across memory an image whose size is equal to the smallest rectangle encompassing the domains of all input images. Still, it is desirable to produce results that are independent of the processing order. Indeed, order independence ensures that the same results are obtained whatever the order in which the images are processed so that results are exactly reproducible. The pursuit of this goal leads to the idea of sequential order independent image compositing. This idea was first briefly presented in [1] and is summarised hereafter.

3.1 The concept of overlap matrix

When processing a given image, it is convenient to know at any time the set of images whose definition domains overlap this image. This can be achieved by computing once for all the so-called *overlap matrix*. This is a symmetrical $n \times n$ matrix indicating whether the definition domains of an arbitrary image pair overlap or not:

$$m_{i,j} = \begin{cases} 1, & \text{if } \mathcal{D}_i \cap \mathcal{D}_j \neq \emptyset, \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

In general, owing to the arbitrary shape of the definition domains $\mathcal{D}_i$, the overlap matrix cannot be constructed on the sole knowledge of the frame size and positioning of the upper left corner in the reference coordinate system. Indeed, the footprint of a georeferenced satellite image does not cover the full image frame (see example in Fig 1). The domain (footprint) covered by data values is called data region of interest (data ROI). It has usually a trapezoidal shape or more complex shape in case the data values values are deemed reliable only in sub-domain of the full data ROI domain. Hence, for all pairs of images whose frames intersect, we need to compute the intersection between their actual definition domains to assess whether these intersect or not.

3.2 Ordered propagation

Beyond the information available through the overlap matrix, the number of overlapping images (called *overlap level* hereafter) for any given pixel must be known. Indeed, seam lines are detected in an ordered fashion, starting with all regions whose overlap level is equal to 2 and proceeding to the subsequent level until the maximum number of overlap is reached.

We then proceed as follows. The $n \times n$ overlap matrix is scanned row by row. The index of the row defines the current *anchor image*. Actual processing is restricted to the region defined by its frame: we assume that the actual definition domain (i.e., data ROI) is buffered by the image frame). Within this region, the least overlap level greater than 1 defines the *current overlap level* denoted by k. We also track whether an overlap level higher than the current level occurs. This is used to determine whether the overlap matrix needs to be scanned again later on. The morphological compositing routine [5] is then called while restricting its effect to the processing of those regions whose overlap level is equal to k. The routine assigns each pixel of these regions to a unique source image so that the definition domain of the anchor image and those of the images intersecting it can be updated accordingly. This update can only remove some parts of these input definition domains since it concerns regions where more than one image was competing for the same domain. This procedure secures order independence in the sense that identical results are obtained whatever order the anchor images are processed.

4 Parallel processing

The sequential algorithm described in the previous section allows for the processing of an arbitrary number of scenes. However, because these scenes must be processed one after the other, this is not a solution for processing very large image data sets ($n > 100$) such as continental or global coverages. In this section, we propose an order independent parallel algorithm allowing for well chosen individual images to be processed in parallel. The concepts of overlapping matrix and ordered processing by increasing level of overlap level originally developed for sequential processing are also used for parallel processing. Section 4.1 introduces the notion of independent sets of anchor images. The parallel algorithm as such is described in Sec. 4.2.

4.1 Independent sets of anchor images

The key to parallel processing is to recognise that two scenes can be processed in parallel if and only if they are independent, i.e., if the processing of the first one at the current processing level is independent of the outcome of the processing of the second one at the same processing level. This happens if there is an empty intersection between (i) the definition domain (ROI) of the first scene unioned with the definition domains (ROIs) of all its overlapping scenes and (ii) the definition domain of the second scene unioned with the definition domains (ROIs) of all its overlapping scenes. A list of independent images is such that any pair of images of the list are independent. A largest list of independent images is any list of independent images that cannot be extended without violating the independence condition. The search for *the* largest list(s) of independent images that can be extracted from the image list may prove very difficult (a brute-force search would require testing all permutations of the input image list). Therefore, in practice, it is acceptable to generate *a* largest list. Such a list can be obtained by scanning the image list and insert the current image in the independent list if it is independent of all images already inserted in the independent list. The resulting list of independent images is not sufficient because it does not contain all input images. Therefore, the procedure must be repeated so as to create additional lists of independent images until all images are selected. That is, as soon as an image is selected in a list, it is flagged as not selectable during the creation of subsequent lists. In addition, because images must be processed in increasing order of overlap level, the creation of lists of independent images must be performed for each successive processing level. To speed up the creation of the lists, when creating the overlapping matrix, the maximum overlap level of each image is computed. This prevents from inserting any image whose maximum overlap level exceeds the current overlap level.

4.2 Ordered processing

Ordered processing is achieved by processing images in parallel for increasing overlap levels. Initially, all regions where only 2 images overlap are processed. This step is itself achieved by processing in parallel a largest list of independent images and then a subsequent list of independent images until all images have been processed. The algorithm proceeds by processing all regions where 3 and only 3 images overlap, and so forth until the maximum level of overlap $g_{\max}$ is reached.

This algorithm can be run on a batch queueing system. This type of systems distributes automatically the composition of the independent images held in the list plist to the available processors. Once all images of the independent list have

Figure 9: Making the razor cut with a computer farm consisting of 16 blades containing each 2 CPUs and allowing for the definition of cuts lines of more than 1,500 images in less than 12 hours. Compare with Fig. 3.

been processed, the subsequent list is sent to the queueing system and so forth until all images have been processed for all overlap levels. Figure 9 shows a grid engine composed of 16 blades with two CPUs each. This system allows for the definition of cut lines for a pan-European coverage at 25m resolution (more than 1,500 images) in less than 12 hours. This modern razor can be compared with the one used in the seventies (see Fig. 3).

5 Conclusion

The morphological image compositing proposed in [5] allows for the automatic definition of seam lines during mosaicing. In this report two algorithms enabling the processing of large data sets have been put forward. The first permits sequential processing so that it is not necessary to hold in random access memory temporary images whose size are equal to that of the target mosaic. The second enables the processing of several images in parallel. This latter algorithm is an asset for mosaicing large data sets such as the IMAGE-2006 imagery covering the entire territory of the European Union plus 11 additional countries. The resulting cloud free mosaics are described in [7].

References

[1] C. Bielski and P. Soille. Order independent image compositing. *Lecture Notes in Computer Science*, 3617:1076–1083, 2005. doi:10.1007/11553595_132.

[2] B. Bishop. Landsat looks at hometown Earth. *National Geographics*, 150 (1), July 1976.

[3] J.-F. Rivest, P. Soille, and S. Beucher. Morphological gradients. *Journal of Electronic Imaging*, 2(4):326–336, October 1993. doi:10.1117/12.159642.

[4] P. Soille. *Morphological Image Analysis: Principles and Applications*. Springer-Verlag, Berlin Heidelberg New York, 2nd edition, 2003.

[5] P. Soille. Morphological image compositing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(5):673–683, May 2006. doi:`10.1109/TPAMI.2006.99`.

[6] P. Soille. IMAGE-2006 Mosaic: Cloud detection on SPOT-4 HRVIR, SPOT-5 HRG, and IRS-LISS III. Technical report, European Commission, Joint Research Centre, 2008.

[7] P. Soille and C. Bielski. IMAGE-2006 Mosaic: Product description. Technical report, European Commission, Joint Research Centre, 2008. In preparation.

[8] P. Soille and L. Vincent. Determining watersheds in digital pictures via flooding simulations. In M. Kunt, editor, *Visual Communications and Image Processing'90*, volume 1360, pages 240–250, Bellingham, 1990. Society of Photo-Instrumentation Engineers. doi:`10.1117/12.24211`.

[9] L. Vincent and P. Soille. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6):583–598, June 1991. doi:`10.1109/34.87344`.

[10] P. Wolf. *Elements of Photogrammetry (with Air Photo Interpretation and Remote Sensing)*. McGraw-Hill, New York, 1974.

[11] S. Zucker. Region growing: childhood and adolescence. *Computer Graphics and Image Processing*, 5:382–399, 1976.

European Commission

EUR 23636 EN – Joint Research Centre – Institute for Environment and Sustainability
Title: IMAGE-2006 Mosaic: Automatic Seam Line Delineation
Author: Pierre Soille and Conrad Bielski
Luxembourg: Publications Office of the European Union
2011 – 16 pp. – 21.0 x 29.7 cm
EUR – Scientific and Technical Research series – ISSN 1831-9424 (online), 1018-5593 (print)
ISBN 978-92-79-20957-4
doi:10.2788/50299

Abstract
This report details the methodology developed for automatically composing the IMAGE-2006 imagery. The method is based on the automatic definition of seam lines in the regions where two or more images overlap. This is achieved thanks to morphological image compositing. A parallel algorithm enables the timely mosaicing of large data sets on a grid engine.

The mission of the JRC is to provide customer-driven scientific and technical support for the conception, development, implementation and monitoring of EU policies. As a service of the European Commission, the JRC functions as a reference centre of science and technology for the Union. Close to the policy-making process, it serves the common interest of the Member States, while being independent of special interests, whether private or national.

LB-NC-23636-EN-N

EUROPEAN COMMISSION

Publications Office

ISBN 978-92-79-20957-4