First Analyses of Sentinel-1 Images for Maritime Surveillance

Harm Greidanus, Carlos Santamaria

2014

Abstract

Sentinel-1 is the European Synthetic Aperture Radar (SAR) satellite operational since 3 October 2014. The SAR's characteristics should make it suitable for maritime surveillance (ship detection), and it will routinely collect a large amount of maritime imagery over European and global seas. After its launch in April 2014, preliminary data have been made available to limited users in the satellite's commissioning phase, and since the start of the operational phase data are available to the general public. These early data have been used to assess the quality of Sentinel-1 images and their suitability for ship detection. This was partly done by using the JRC's ship detection software SUMO, after adaptation to ingest and process Sentinel-1 data. It is found that the sensor lives up to its specifications, thereby making it very useful for maritime surveillance thanks to its combination of wide swath and low noise at the medium resolution with which it will mostly be operated ("IW" and "EW" modes).

Front cover: Part of Sentinel-1 image of the Gulf of Riga, SM mode, VH polarisation green, VV polarisation blue. Taken 22 June 2014 during the commissioning phase when the system was still being tuned.

Executive Summary

Sentinel-1 is the Synthetic Aperture Radar (SAR) satellite of the EU's Copernicus Earth Observation program. SAR is type of sensor that is suitable for ship detection, and Sentinel-1 combines a wide swath at medium resolution with a significant amount of routine daily acquisitions to form a potentially powerful instrument for maritime surveillance with a global reach. Sentinel-1 has three main modes of operation that are relevant for maritime surveillance: Stripmap (SM) with 80 km swath at 9 m resolution, Interferometric Wide (IW) with 250 km swath at 21 m resolution, and Extra Wide (EW) with 410 km swath at 50 m resolution. Products with other resolutions and noise levels can be derived from these modes, but the quoted resolutions are the main ones.

In contrast to most previous satellite SARs that only acquire images on demand, Sentinel-1 will routinely collect data in pre-defined modes and areas, with the aim to cover the European and Canadian seas, Arctic and Antarctic waters, and some global coasts. The IW and EW modes will be mostly used, with the SM mode reserved mainly for emergencies and disasters. So from Sentinel-1, we may expect routine sampling of the medium- and large-sized ship traffic in the mentioned areas.

Using the early data that have become available in a restricted way before, and more broadly after the 3 October 2014 date of reaching operational status, the sensor has been assessed for its image quality and its suitability for ship detection. Generic (radar) image properties of quantisation, resolution, thermal noise, speckle noise and geo-location accuracy have been analysed, using only scene backgrounds and targets of opportunity. This mostly confirms the specifications of the system. Some problems remain in using the geo-location tie points grids provided in the image metadata to yield sufficient accuracy, possibly depending on the way the grids are applied, but they can be circumvented by using the provided orbit data instead of the grids for the geo-location. The quality of the products was already improved by October following the findings of the commissioning phase, and further improvements are expected until the end of the ramp-up by May 2015.

Specifically for ship detection, the JRC's "SUMO" software, that can already ingest data from most of the common satellite SAR systems, was adapted to also ingest and analyse Sentinel-1 images. Applying SUMO to Sentinel-1 confirms the sensor's good performance for ship detection, as the sensor provides well-focussed images at low thermal and speckle noise levels with few artefacts. The only unexpected image artefacts are weak ghosts (re-appearances) of bright targets at some distance away in azimuth direction (in addition to the usual azimuth ambiguities). These can be a problem because they can result in false alarms. Their cause is not yet known. For the rest, Sentinel-1 displays the same instrument-related false alarm types as expected for a SAR, namely azimuth and range ambiguities. Side lobes, another instrument-created false alarm cause, are low, but that is as expected at medium resolution.

For accurate ship detection results, it is still needed to perform a manual verification on the outputs of automatic ship detection software. An estimation is made of how much data over seas and coasts will be acquired by Sentinel-1 on a daily basis, and it is concluded that fully automatic ship detection with SUMO can be performed on such a data stream in a sustained way, but that there is no time for manual verification. Possible strategies to deal with that are raising of the detection threshold, and the cataloguing of stationary features (an important false alarm cause) that re-occur in images that are acquired over the same area in the same mode with the satellite's repeat interval of 12 days.

The next steps to be done in order to further utilise Sentinel-1 for maritime surveillance include: setting up a processing infrastructure to routinely process the maritime images for ship detection; improve the automatic detection algorithms to better deal with false alarms, especially related to sea ice and, as mentioned above, stationary features; develop the capability to deal with Single Look Complex (SLC) data of Sentinel-1, which can be used for experimental purposes but probably not for bulk application in ship detection.

Contents

1 Introduction

Under the Copernicus program for Earth Observation implemented by the EU and ESA, Sentinel-1 (S-1) is the Synthetic Aperture Radar (SAR) satellite. Its images are useful for maritime surveillance, because medium and large ships can be found in them. S-1 will routinely collect maritime images over the European seas including Arctic waters, seas around Canada, around the Antarctic, and some of the world's coasts.

S-1 has four image modes, with different swath widths and resolutions. Each mode can be processed to different products. The table below lists the modes with their main products [1, 2].

Mode	Swath width (km)	Resolution of SLC product (m)	Resolution of GRDH product (m)	ENL of GRDH product
EW - Extra Wide	410	20 x 43	50 x 50	2.8
IW - Interferometric Wide	250	5 x 22	20 x 22	4.4
SM - Stripmap	80	5 x 5	23 x 23	29.7
WV - Wave	20	5 x 5	-	-

Table 1.1. Main parameters of the four S-1 image modes, for SLC (Single Look Complex) and GRDH (Ground Range Detected High-resolution) products. Resolutions are ground range x azimuth (also for SLC). Resolution and ENL (Equivalent Number of Looks) are indicative as they may change across the image. The ENL of the SLC product is by definition 1. The ENL determines the speckle noise level of the image: the higher the ENL, the lower the speckle noise, starting from 1 with the highest noise level.

The SLC (Single Look Complex) product retains the most detailed information, but is difficult to work with. It is not assessed in detail in this report, but some experiences are listed in Section 6.

The Wave mode is designed to monitor the global wave climate. It might conceivably be used for ship detection, but it has not yet been assessed for that purpose.

The Stripmap mode has a choice of 6 beams (S1 – S6) with incidence angles going from near to far. The IW and EW modes are multi-beam. The IW mode is made up of 3 beams that are combined into one wide image. The EW mode is made up of 5 beams.

The Stripmap mode has a Full resolution product (GRDF) with 9 x 9 m resolution and ENL 3.7. It is suitable for maritime surveillance, but up to now very few of these products have been made available. Therefore little analysis could be done on this product.

The EW, IW and SM modes all have a Medium resolution product with about 90 m resolution. This type of product is less suitable for ship detection.

S-1 images can have single polarisation HH or VV; or dual polarisation HH+HV or VH+VV.

The three products that are discussed in detail in this report are therefore EW-GRDH, IW-GRDH and SM-GRDH; in dual and single polarisation. The following three figures show examples of these products. Throughout the report, each figure caption mentions the image identification. This identification starts with the name of the satellite (S1A), followed by the mode, followed by the product, followed by the polarisation, followed by the start and stop times, followed by three more codes, the last one of which is a unique 4-character identifier by itself. The start and stop times are of the form yyyymmddThhmmss in UTC.
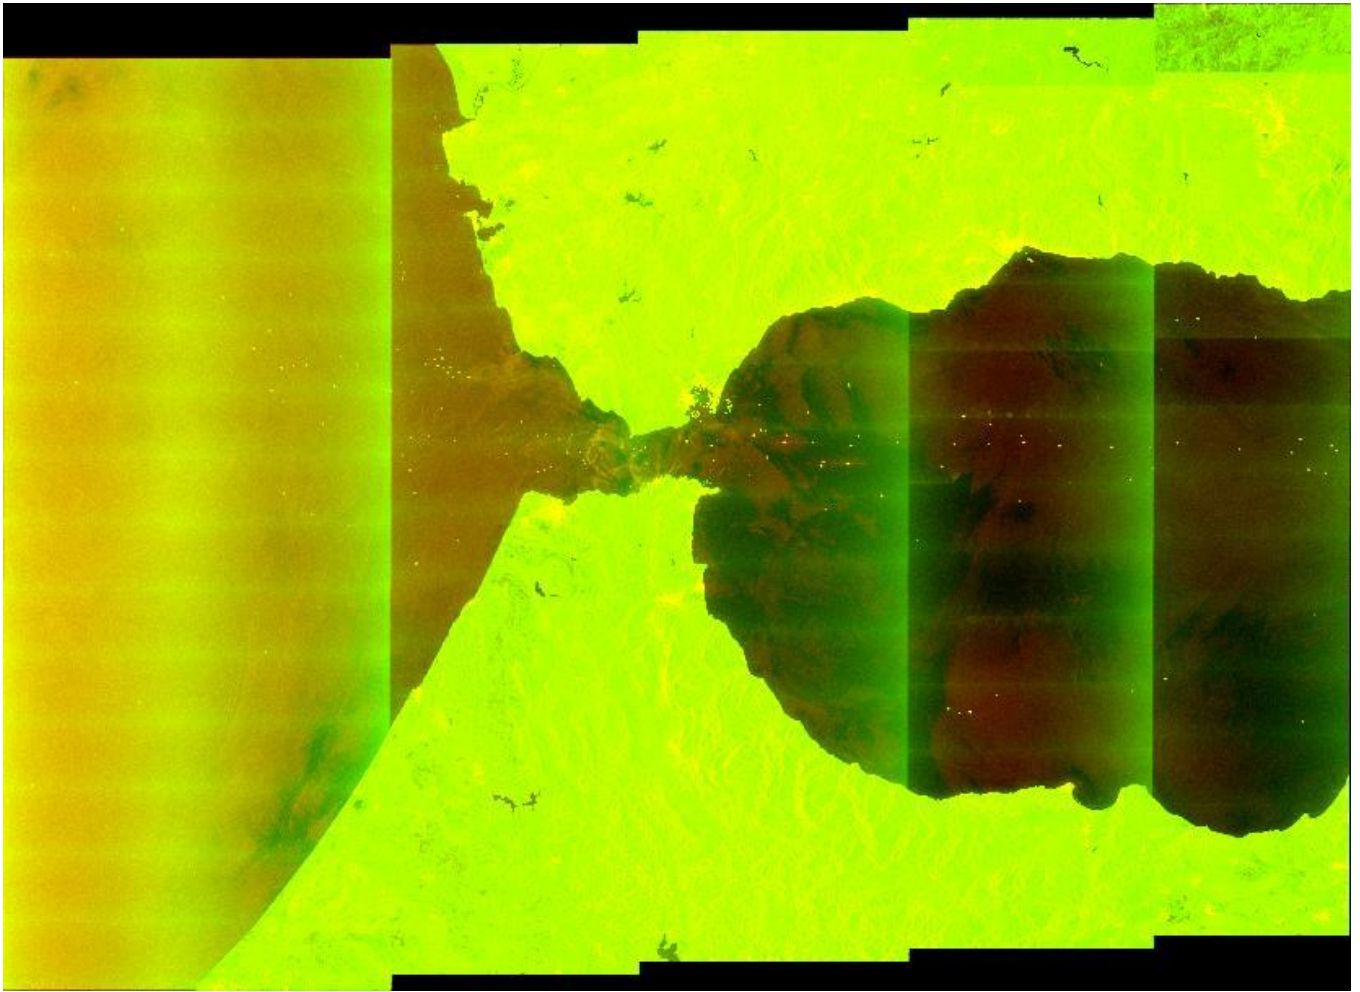
Fig. 1.1. Extra Wide mode (EW) image, dual polarisation with HH colour coded in red and HV in green, over the Strait of Gibraltar. Ships (there are many) show up as bright dots.
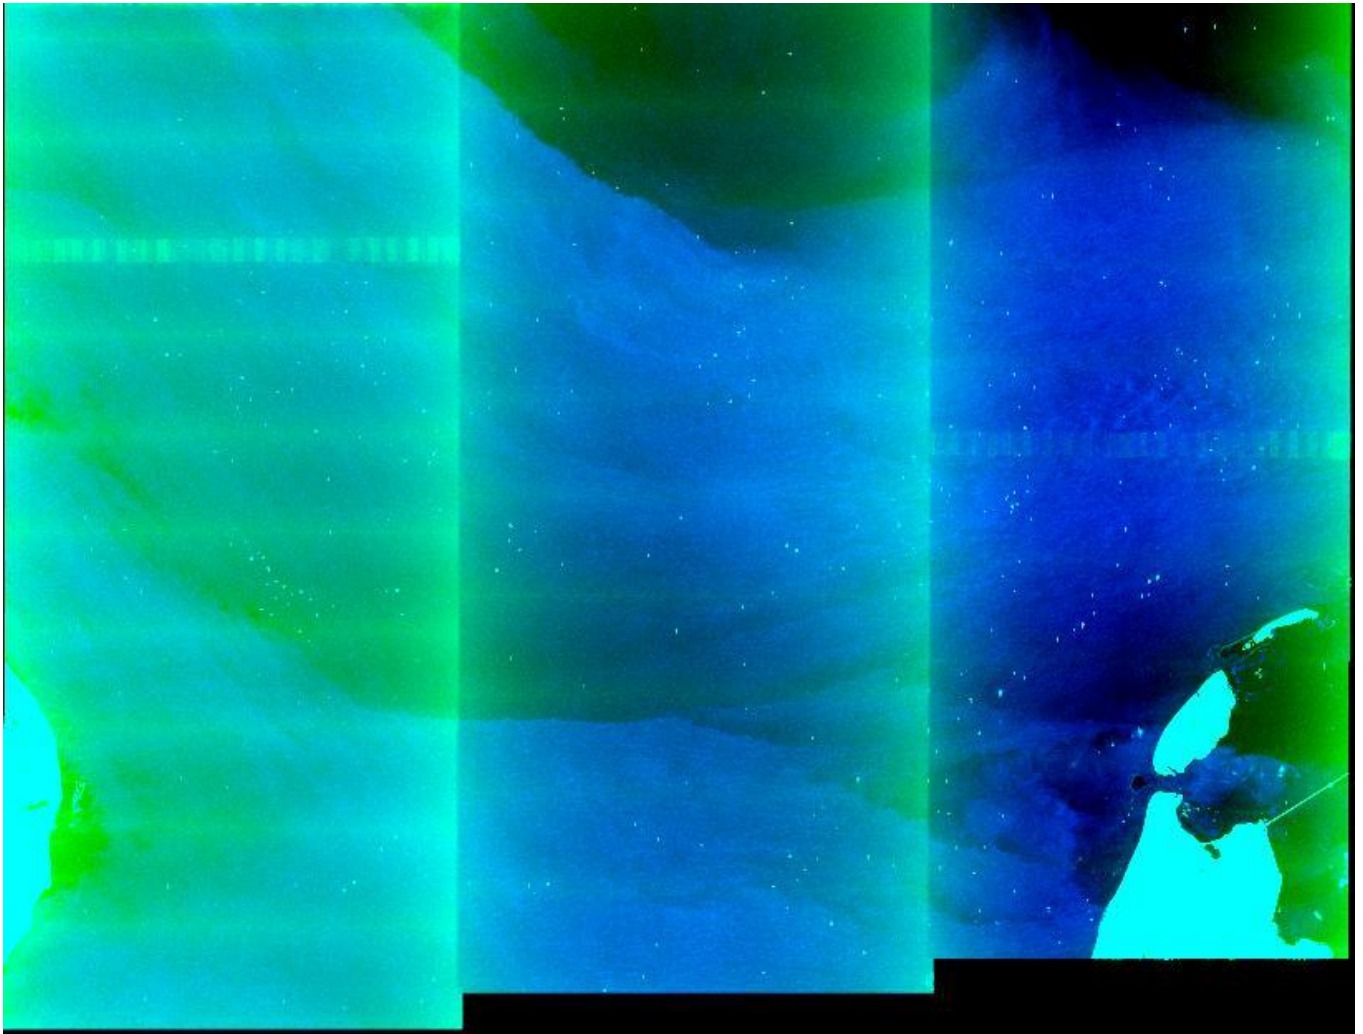*S1A_EW_GRDH_1SDH_20141010T181731_20141010T181816_002771_0031D8_192C*

Fig. 1.2. Interferometric Wide mode (IW) image, dual polarisation with VH in green and VV in blue, over the North Sea. This scene has again many ships. It is further analysed in the Section 4 on Ship detection.
S1A_IW_GRDH_1SDV_20141016T173306_20141016T173335_002858_0033AF_FA6D
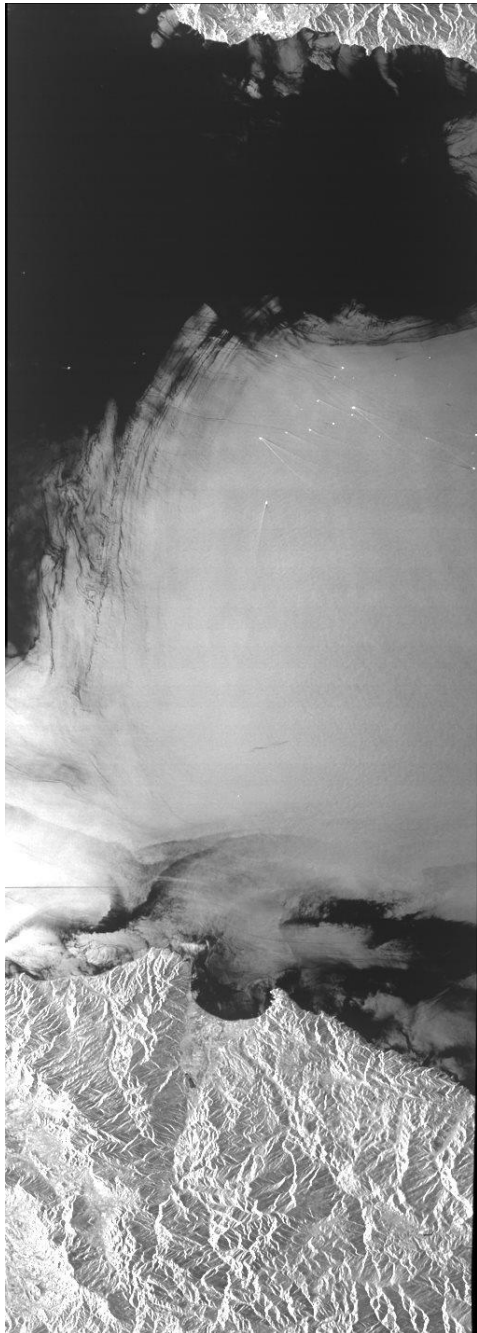
Fig. 1.3. StripMap mode (SM) image, S1 beam, VV polarisation, over the Mediterranean. Ships and V-shaped ship wakes can be seen.
S1A_S1_GRDH_1SSV_20141004T062656_20141004T062730_002676_002FC5_C589

In the dual-polarisation images, the two coloured ones above, the cross-pol channel (HV or VH) is dominated by large-scale variations that are instrumental effects. They arise because the radar backscatter from the sea surface is mostly below the noise level in cross-pol. In overview images over the sea such as the three above, the addition of this channel is not really helpful. However, for ship detection the cross-pol channel is very useful because the ship-sea contrast is highest in cross-pol. See also the cover image of the report where the ships are seen as green (VH) dots.

The remainder of the report will discuss image quality aspects, geocoding, ship detection, instrumental effects that interfere with ship detection namely ghosts and ambiguities, the use of SLC images, and scoping estimates for bulk processing.

2   Image quality

This section looks at some image quality aspects:
- Quantisation,
- Calibration,
- Antenna pattern,
- Equivalent Number of Looks,
- Spatial resolution.

These parameters are assessed by analysing the sea surface background of an image. The same example image, shown below, will be used to illustrate the analyses.
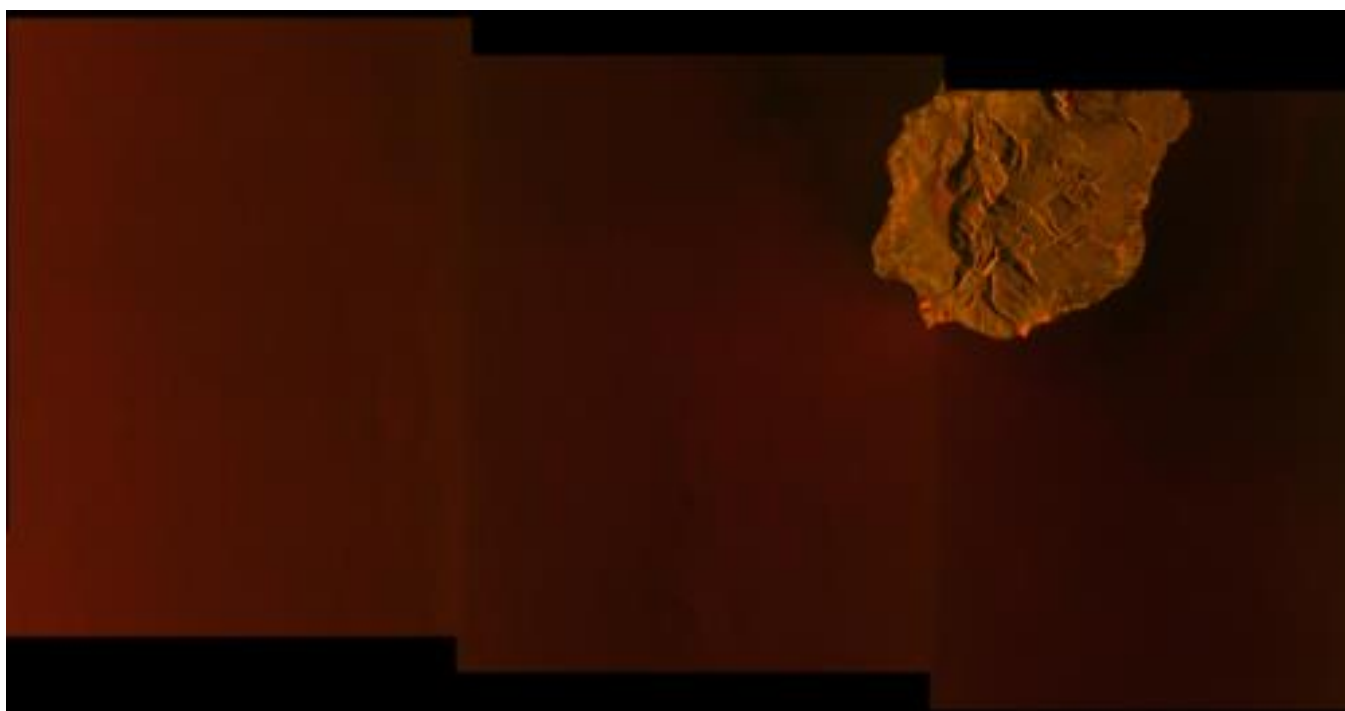

Fig. 2.1. The image used to illustrate the image quality measures. This is the quick look given with the product distribution. The island is La Réunion. Range direction is to the right, azimuth direction is vertical.
S1A_IW_GRDH_1SDV_20141003T145245_20141003T145304_002666_002F86_BFA7

In the Sections 2.1-2.3 below, results are shown for the image quality measures on this example image. The original image is an IW mode dual pol (VV+VH) GRDH product of 25220 pixels (range) by 12874 lines (azimuth). This is split up into adjacent windows of 200 pixels x 200 lines, and in the derived images below in Sections 2.1-2.3 each pixel is the average result of one 200x200 window. The land is masked out using an externally provided land mask (GSHHG [3]).

2.1   Quantisation

The quantisation looks at how the real (floating point) numbers that have come out of the SAR processing (the construction of the SAR image from the raw measured radar data) have been converted into the range of 0 to 65535 that is covered by the unsigned 2-byte integer format of the digital numbers in the image file. The dynamic range of the original real numbers may be too big to fit in the 65535 dynamic range, and then either strong targets saturate (values > 65535 are clipped) or the low levels become badly sampled (all around 0, 1, 2) eliminating details in the low-backscatter areas. In the images that were tested, however, the sea surface pixel values are all well quantised.

The example below shows that mean background values do not go much below 15, and at that level the radar backscatter from the sea surface has already dropped below the thermal noise.
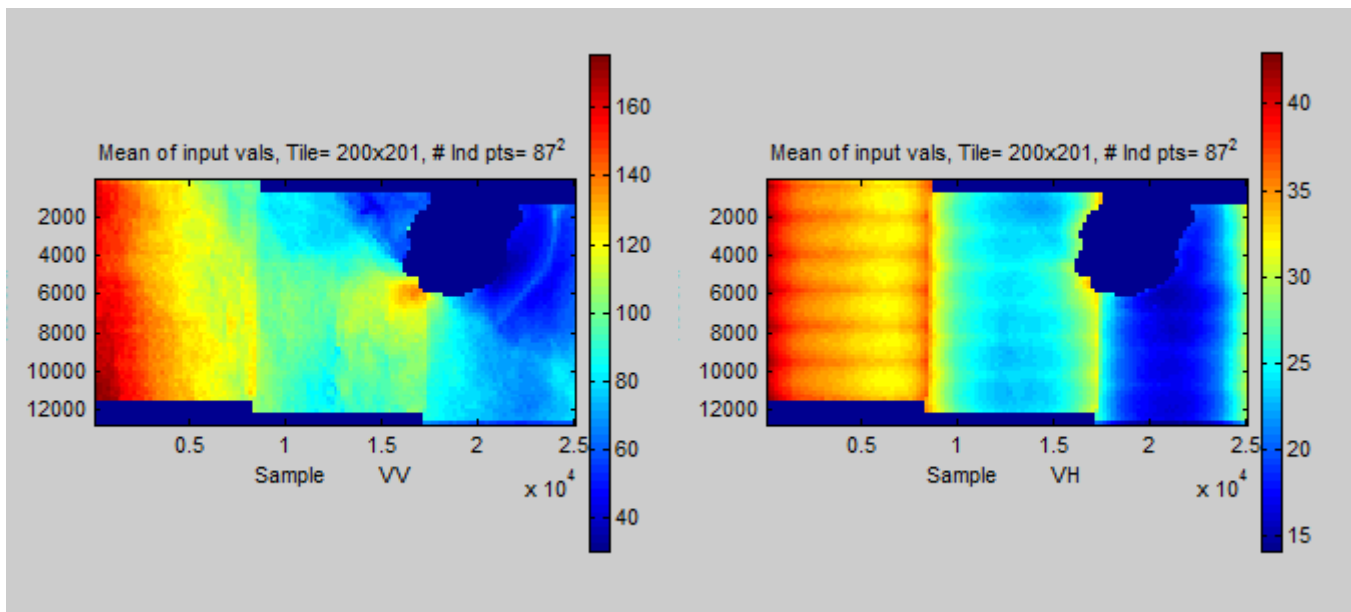

Fig. 2.2. The mean of the input pixel values (which have to be in the range 0 to 65535) in windows of 200 x 200 pixels. In this and the following figures, which all pertain to the image of Fig. 2.1, the VV polarisation is on the left and the VH polarisation on the right. The IW image is made up of three sub-swaths that are mosaicked together in the GRD products and they can be recognised as the three rectangular adjoining parts. Both VV and VH images are dominated by the large scale range drop-off in intensity (going from red to blue from left to right.) In the VH image, the antenna pattern correction is visible as a trough (bright-dark-bright in horizontal direction) in each of the three sub-swaths. Also some horizontal banding is visible in each sub-swath due to the TOPS scanning. In any case, the pixel values of the sea surface / thermal noise background are well quantised being far above the value 1.

2.2   Calibration

The calibration converts the digital number of the pixels in the image file to absolute normalised radar cross section (NRCS, also called sigma0). The S-1 product provides metadata for that conversion. As with the geo-location, the conversion factors are provided for a set of specified pixel and line values, all in an xml file, separate per polarisation. Although the (pixel, line) grid could be irregular, it is always found to be regular, greatly aiding the interpolation. The conversion factors provided are those to go from the digital pixel value to sigma0, to beta0 and to gamma, which only differ by simple geometrical factors. It is found that (up to now) the conversion factor to sigma0 is a function of range (i.e., of pixel) only (and not of azimuth / line), so the conversion to sigma0 can be described by a range profile. Moreover, the conversion factor to beta0 is always the same for one image, across the whole grid and for each polarisation. Sigma0 and beta0 are related through the cosine of the incidence angle. It seems therefore that the xml files that specify all the calibration metadata could be replaced by one single number.
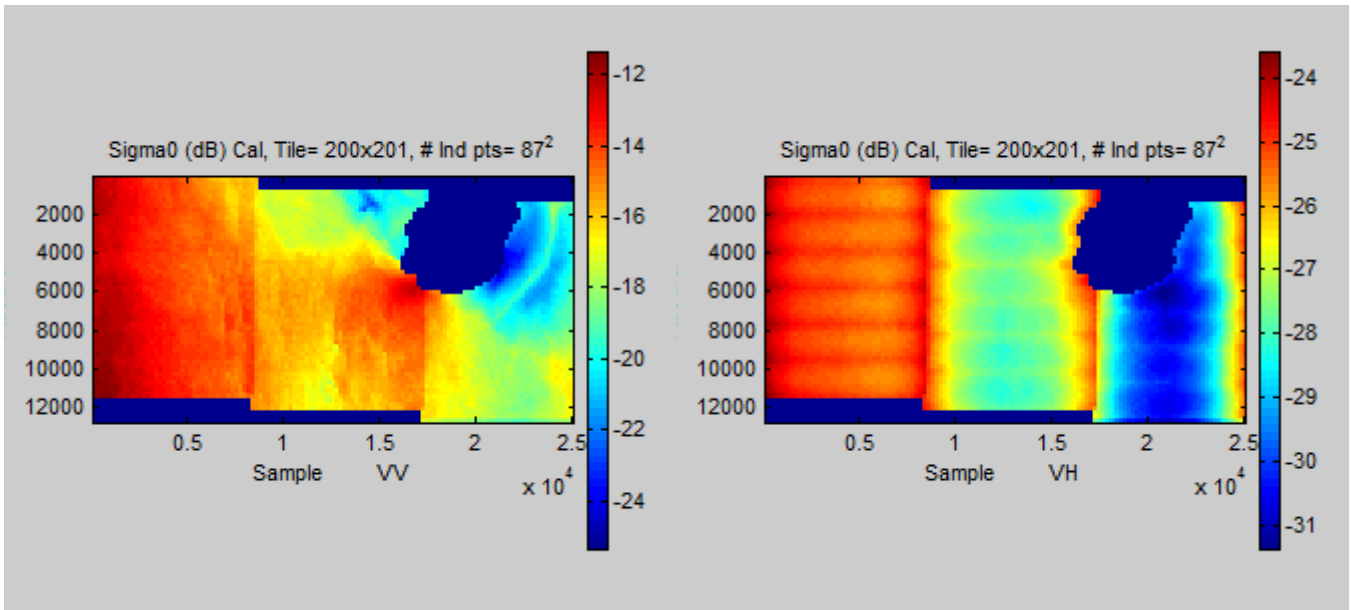
Fig. 2.3. The sea surface background values converted to absolute sigma0 (normalised radar cross section of the sea surface). Again, the drop-off in range can be seen (red to blue from left to right) which is a real physical effect of the sea surface. The VV image for the rest shows structures mostly related to the large-scale wind field. The VH image is dominated by instrument effects, because the radar backscatter from the sea surface in VH at this wind speed is below the noise level. The sigma0 values in absolute sense are reasonable – a quantitative verification would need meteorological data of wind speed, wind direction and atmospheric stability which are not available here.

The antenna pattern as a function of elevation (therefore, of range coordinate in the image) is corrected for in the calibration that was applied in making the image product. This correction flattens an image of constant normalised radar backscatter. However, when the radar backscatter drops below the sensitivity of the system, then the antenna pattern correction shows up in a varying noise profile. This is the effect in Fig. 2.3 VH (right panel). The S-1 metadata also provide noise profiles to correct for the noise level, but these have not yet been used.
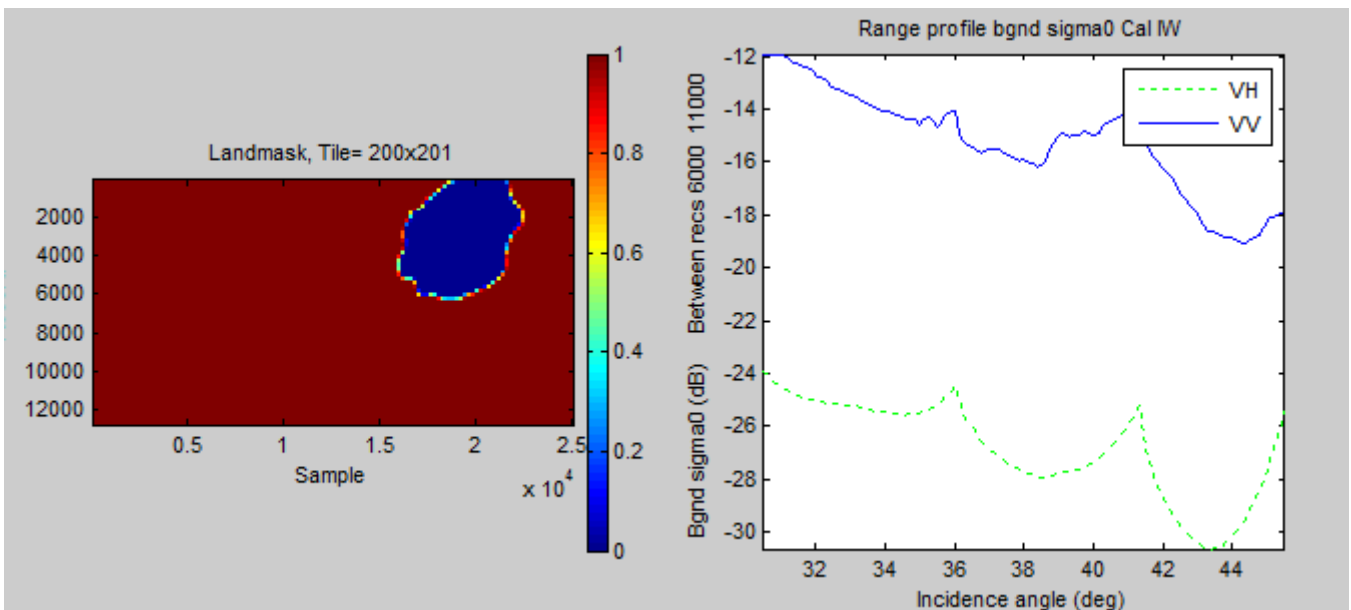

Fig. 2.4. Left, the land mask that was used. Right, the range profile of the calibrated sigma0 image background averaged between lines 6000 and 11000. The drop-off with range (converted in this graph to incidence angle) in VV (blue) is geophysical. The peaks in the VH profile (green) reflect the antenna pattern correction and mostly the noise floor.
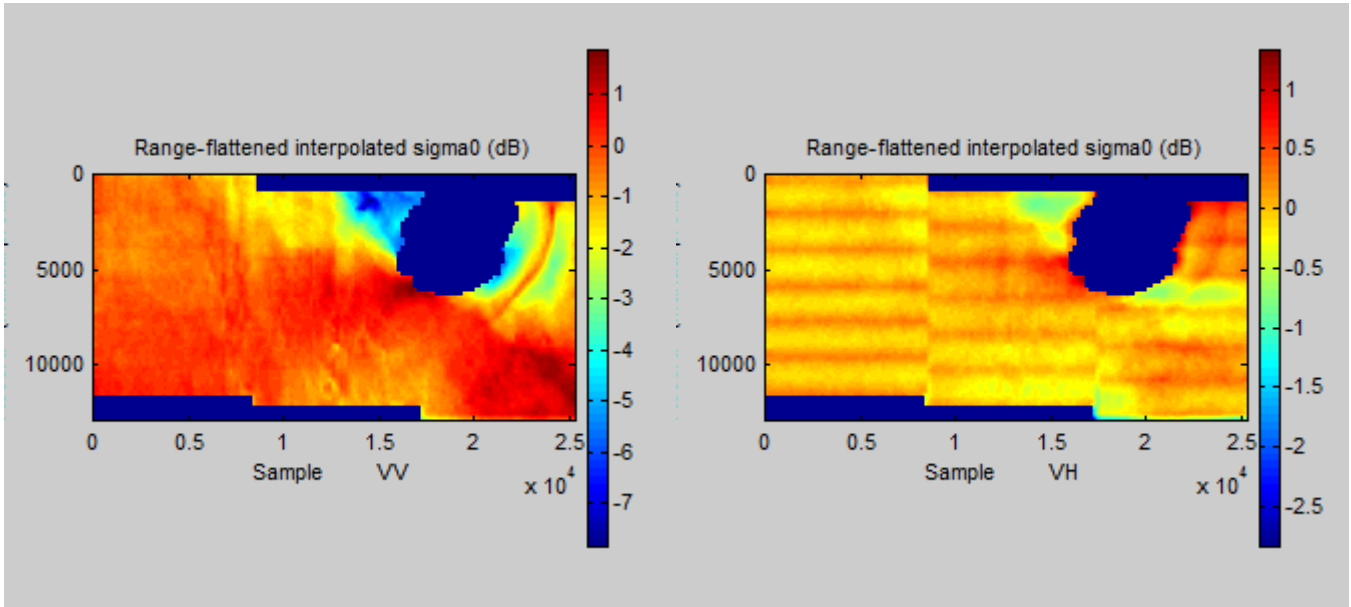
11

Fig. 2.5. The calibrated images (in dB), flattened by subtracting the mean background sigma0 range profiles of the previous figure. This removes the dominant and range-dependent trend, better revealing the real geophysical structures on the sea surface. In VV (left), a curved front can be seen to the right of the island, and what could be wind shadow on the other side. The VH image is still dominated by instrument effects although some of the geophysical variations around the island vaguely show through.

2.3   Equivalent Number of Looks

The number of looks of a SAR image is the number of independent images of the same scene that have been averaged in making the final image. The independent images of the same scene are obtained by splitting the radar bandwidth in separate intervals and constructing images separately for each interval. This can be done both in range and in azimuth direction. Each of these images shows the same scene, but with different speckle noise. In averaging these images, the speckle noise is reduced while the real scene content remains. When N such independent images are averaged, the "number of looks" is N, and the speckle noise is reduced by the square root of N. However, the images can be constructed from overlapping bandwidth intervals. This gives rise to the concept of Equivalent Number of Looks (ENL) that may be fractional. The ENL is a measure of the speckle noise level; the higher the ENL, the better. For a scene with a homogeneous background, i.e. with constant sigma0, the ENL can be estimated by measuring the noise level (standard deviation over mean, i.e. the coefficient of variation, of the pixel values). If the background in such a measurement is not homogeneous, then the ENL value obtained will be too low, because in addition to variation due to the speckle, it also contains variation due to the background. The ENL can therefore be estimated in a real image by splitting the image into many small areas, measuring the coefficient of variation in each area, converting to ENL, and selecting the highest occurring ENL values. The image should be visually assessed to make sure that the areas where the highest ENL values occur are indeed homogeneous. In calculating the mean and standard deviation of each small area, high outlier values are discarded to avoid biasing of the statistics by the presence of targets.
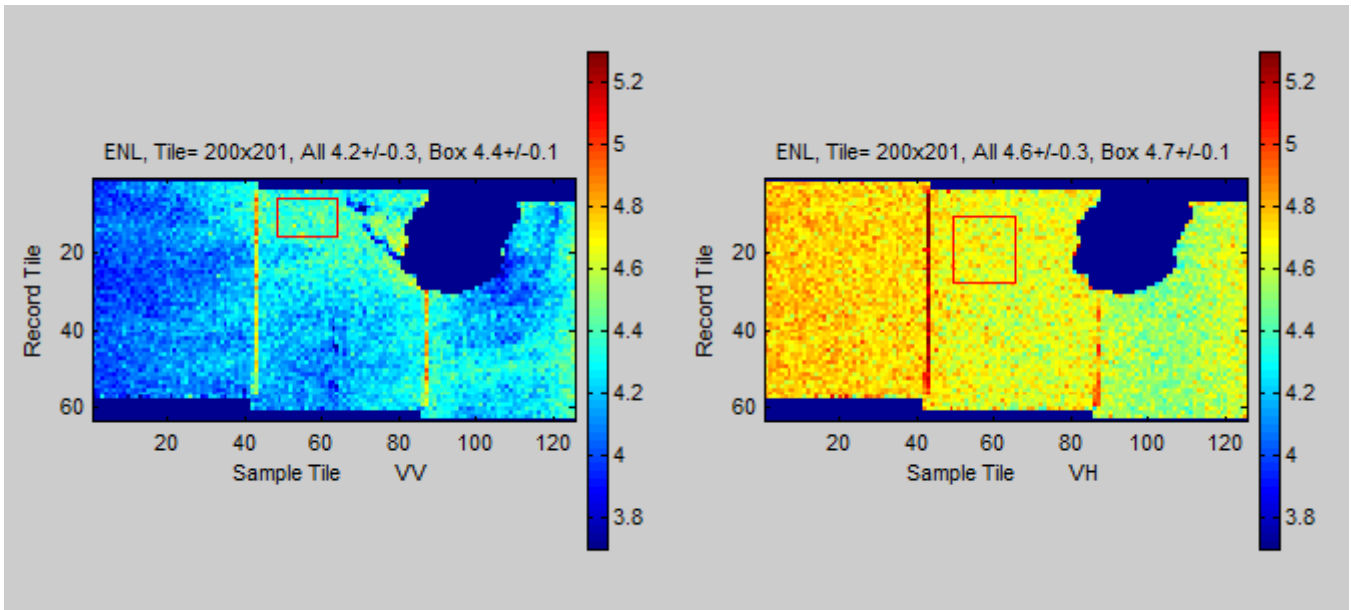
Fig. 2.6. The measured ENL values in each 200x200 pixel window as discussed above. The areas with the highest ENL values were manually selected, indicated by red boxes. The ENL value measured there is 4.4 for the VV image and 4.7 for the VH image. The latter value is the most reliable, because the VV image contains structure from the sea surface that the VH image lacks making the latter more homogeneous. Therefore, based on this and also on the histograms of the next figure, the ENL of the image is estimated as 4.7.

The ENL is an important parameter that needs to be specified for CFAR (Constant false Alarm Rate) ship detection (see Section 4), because it is needed to calculate the detection threshold. Even though the multi-beam modes of S-1 (EW, IW) have a somewhat different ENL in each beam, the beams are mosaicked (with an overlap) in the GRDH products, and for practical reasons only a single ENL is applied in the ship detection for the whole image. (It would be too cumbersome in the ship detection to change the ENL depending on the location in the image.) Therefore, the analysis in this section aims to establish a single ENL that is applied to the entire image. The small variation of the ENL per beam in the IW GRDH product is readily seen in the right panel (VH) of the previous figure. Also the beam overlap zones stand out in the ENL image, as the averaging in the overlap strips leads to higher ENL.
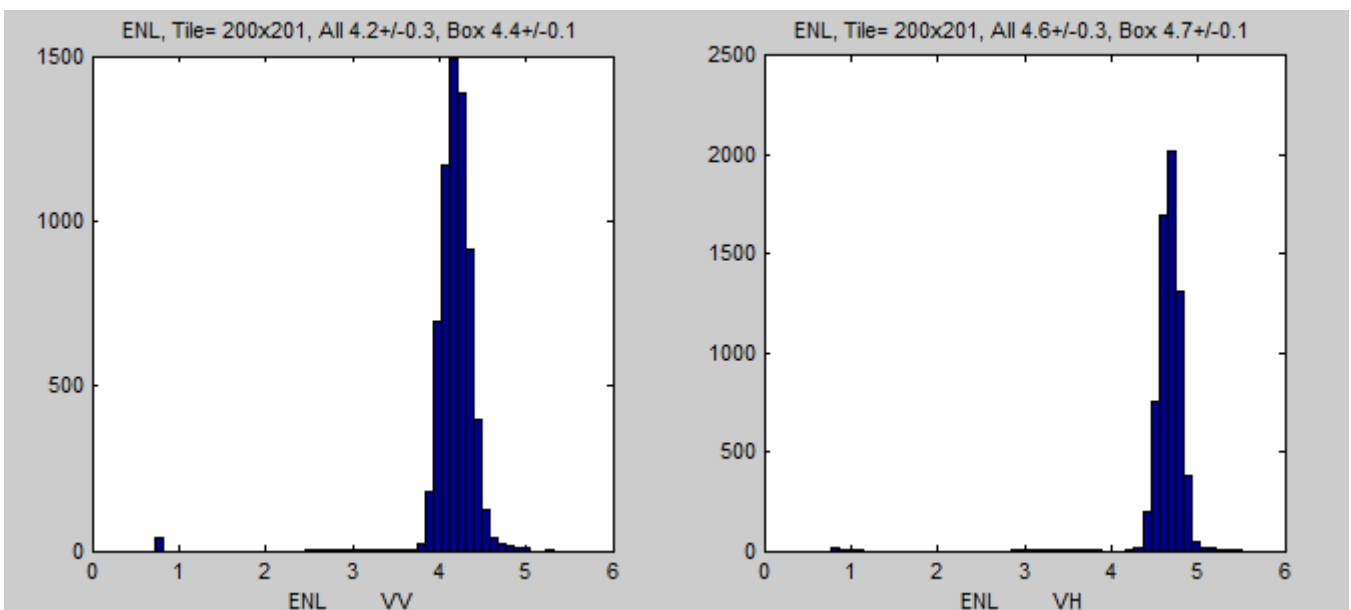

Fig. 2.7. The histograms of the ENL values from the previous figure.
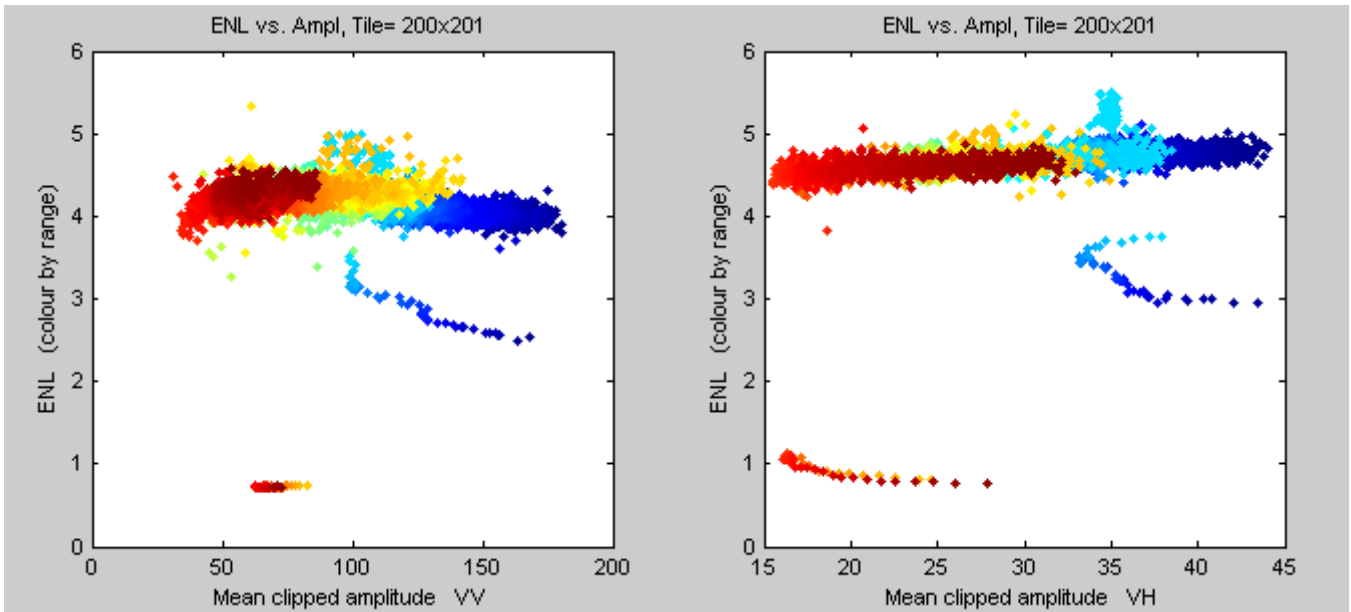
13

Fig. 2.8. These figures plot the relation between mean amplitude not counting high outlier pixel values (horizontal axis), ENL (vertical axis) and range (colour). ENL should not be a function of amplitude, nor of range; if such dependencies are found, there is likely something wrong with the image. In this case, it can be seen that ENL is mostly always around 4.3 (VV) and 4.7 (VH), as already seen in the histograms of the previous figure, independent of amplitude (left-right) or range (colour).

The ENL measurements of a series of images are listed in three tables below, separately for EW, IW and SM modes. Also the specification values are listed. For EW and IW modes, measured values are well in accordance with the specifications, considering the uncertainty in the measurements and the small discrepancies in the specifications in different reference documents. For SM mode, the measured ENL values are somewhat lower than expected. This could be due to the higher resolution in the SM mode, which shows more structure on the sea surface and therefore would lead to a lower ENL estimate.

Image	EW1	EW2	EW3	EW4	EW5
520E	2.5*/2.9/-	2.5*/2.9/-	2.6/2.8/-	2.6/2.7/-	2.7/2.7/-
E33F	2.6/2.9/-	2.7/2.9/-	2.7/2.8/-	2.7/2.8/-	2.7/2.8/-
4FDE	2.6*/2.9/-	2.7/2.9/-	2.7/2.9/-	2.8/2.8/-	2.8/2.8/-
F8D5	-/2.9/2.9	-/2.9/2.9	-/2.9/2.9	-/2.8/?	-/?/?
22A7	-/2.9/2.9	-/2.9/2.9	-/2.9/2.9	-/2.8/2.9	-/2.8/2.9
Overall	2.9	2.9	2.9	2.8	2.8
Specs [2]	2.8	2.7	2.7	2.7	2.7
Specs [1]	2.9				

Table 2.1. ENL measurements in 5 different EW GRDH images. Values are quoted for each of the 5 EW beams (EW1 – EW5) as well as for each polarisation in the image in the order HH/HV-or-VH/VV. An asterisk means that the ENL could not be well estimated due to lack of homogeneous areas. The line 'Overall' makes the best estimate. The lines 'Specs' give the nominal values according to the specifications; slightly different values are given in the Technical Guide (per beam) and in the User Guide (one number for all beams). The left column identifies the image by its 4-character code at the end of the full image name:
S1A_EW_GRDH_1SDH_20141006T165117_20141006T165221_002711_00308A_520E
S1A_EW_GRDH_1SDH_20141008T163548_20141008T163652_002740_003132_E33F
S1A_EW_GRDH_1SDH_20141015T162811_20141015T162915_002842_00334C_4FDE
S1A_EW_GRDH_1SDV_20141020T055155_20141020T055259_002909_0034C1_F8D5
S1A_EW_GRDH_1SDV_20141020T055259_20141020T055359_002909_0034C1_22A7

Image	IW1	IW2	IW3
322E	3.3/4.8/-	3.4/4.7/-	3.5/4.6/-
BFA7	-/4.8/4.3	-/4.7/4.3	-/4.6/4.4
E8A2	4.8/4.8/-	4.3*/4.7/-	4.4*/4.6/-
7860	-/4.7/4.7	-/4.7/4.5	-/4.6/4.5
9A42	-/-/4.3	-/-/4.4	-/-/4.5
E77C	3.0*/4.7/-	3.3*/4.6/-	3.6*/4.5/-
Overall	4.8	4.7	4.6
Specs [2]	4.4	4.4	4.3
Specs [1]	4.9		

Table 2.2. ENL measurements in 6 different IW GRDH images. Values are quoted for each of the 3 IW beams (IW1 – IW3) as well as for each polarisation in the image in the order HH/HV-or-VH/VV. The line 'Overall' makes the best estimate. The lines 'Specs' give the nominal values as in the previous table. The full image names are:
S1A_IW_GRDH_1SDH_20141101T163542_20141101T163611_003090_0038A9_322E
S1A_IW_GRDH_1SDV_20141003T145245_20141003T145304_002666_002F86_BFA7
S1A_IW_GRDH_1SDH_20141029T181208_20141029T181228_003048_0037BA_E8A2
S1A_IW_GRDH_1SDV_20141030T084345_20141030T084410_003056_0037EC_7860
S1A_IW_GRDH_1SSV_20141019T200626_20141019T200655_002903_0034A0_9A42
S1A_IW_GRDH_1SDH_20141108T072001_20141108T072020_003187_003ABA_E77C

	S1	S2	S3	S4	S5	S6
6C0E	-/-/8.8					
C589	-/-/22.5					
E338	-/23.3/5.2					
52DB	-/23.2/6.0					
3A2D		-/-/7.7				
0861			-/-/20.1			
2BC9				-/23.5/13.4		
961F					22.9/19.0/-	
E886						21.4/-/-
Overall	23					
Specs [2]	26.8	26.3	29.7	26.8	29.7	26.8
Specs [1]	34.4					

Table 2.3. ENL measurements in 9 different SM GRDH images. Values are quoted for each of the 6 SM beams (S1 – S6) as well as for each polarisation in the image in the order HH/HV-or-VH/VV. The line 'Overall' makes the best estimate (only for S1 because the other beams have too few measurements). The lines 'Specs' give the nominal values as in the table for the EW images. The full image names are:
S1A_S1_GRDH_1SSV_20141003T162847_20141003T162910_002667_002F8A_6C0E
S1A_S1_GRDH_1SSV_20141004T062656_20141004T062730_002676_002FC5_C589
S1A_S1_GRDH_1SDV_20141020T020949_20141020T021014_002907_0034B3_E338
S1A_S1_GRDH_1SDV_20141020T021039_20141020T021104_002907_0034B3_52DB
S1A_S2_GRDH_1SSV_20141003T183332_20141003T183402_002669_002F92_3A2D
S1A_S3_GRDH_1SSV_20141003T140020_20141003T140048_002666_002F84_0861
S1A_S4_GRDH_1SDV_20141017T181251_20141017T181315_002873_0033FE_2BC9
S1A_S5_GRDH_1SDH_20141017T102409_20141017T102424_002868_0033E7_961F
S1A_S6_GRDH_1SSH_20141020T054944_20141020T055013_002909_0034BE_E886

2.4 Resolution

The resolution is estimated in a few images, and compared with the specified resolution. Resolution is best measured by analysing a point target, but in images over sea it is difficult to find confirmed point targets, and therefore the sea surface background itself is used instead. The smallest structures in the background are at the size of the resolution, and this scale is measured by the central width of the autocovariance function. The autocovariance function is computed via the Fourier Transform of a small square window of sea surface (512 x 512 or 1024 x 1024 pixels). The window should be located in a homogenous area. The figures below show measured autocovariance functions from such a window. The autocovariance functions are extracted separately for range and azimuth direction, and for the two polarisations (HH and HV in the image below). The resolution is estimated from the width of the autocovariance function by dividing by $\sqrt{2}$ [12] and listed in the table following the figures. Only GRDH images have been analysed.

(Fig. 2.9 on next page, caption:) Top row: EW image, containing a coast (the upper half is sea), in two polarisations. The left column is HH, the right column HV, for all rows. Middle row: The window indicated in the red square in the top-left image. Bottom row: Four plots of autocovariance functions of the window. Only the right half of the (symmetric) autocovariance functions is shown. The top two are in range (i.e. samples or pixels) direction; the bottom two in azimuth (i.e. records or lines) direction. The green plots use amplitude values; the red plots powers (= square of amplitudes); both give nearly identical results. The numbers are measured half-width at half maximum (HWHM). S1A_EW_GRDH_1SDH_20140902T221101_20140902T221211_002219_00243F_757B

HH   window= 1065x1255

HV   window= 1065x1255

Intp'd avg autocovariance, centre HH

Intp'd avg autocovariance, centre HV
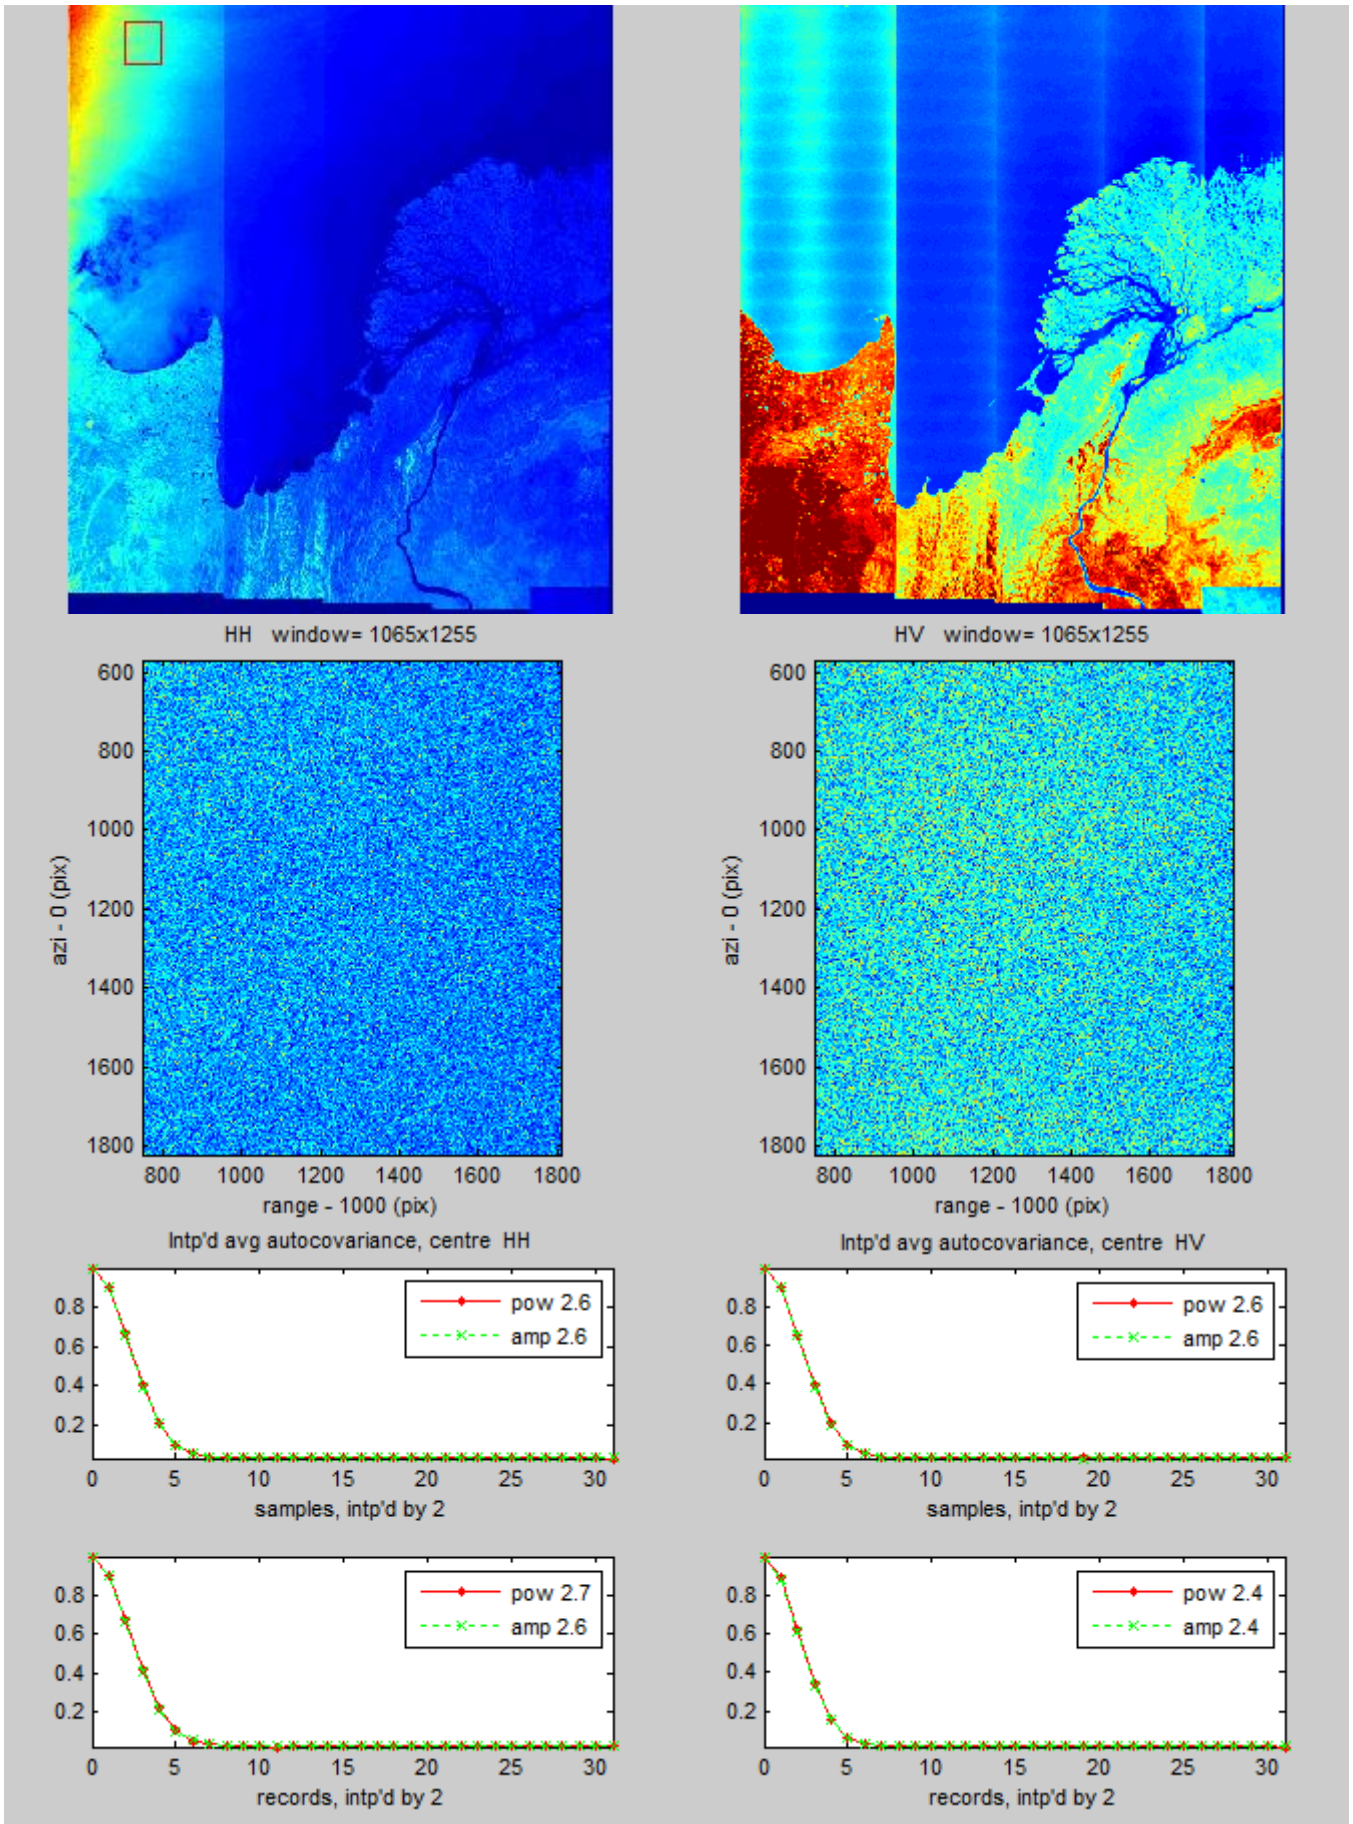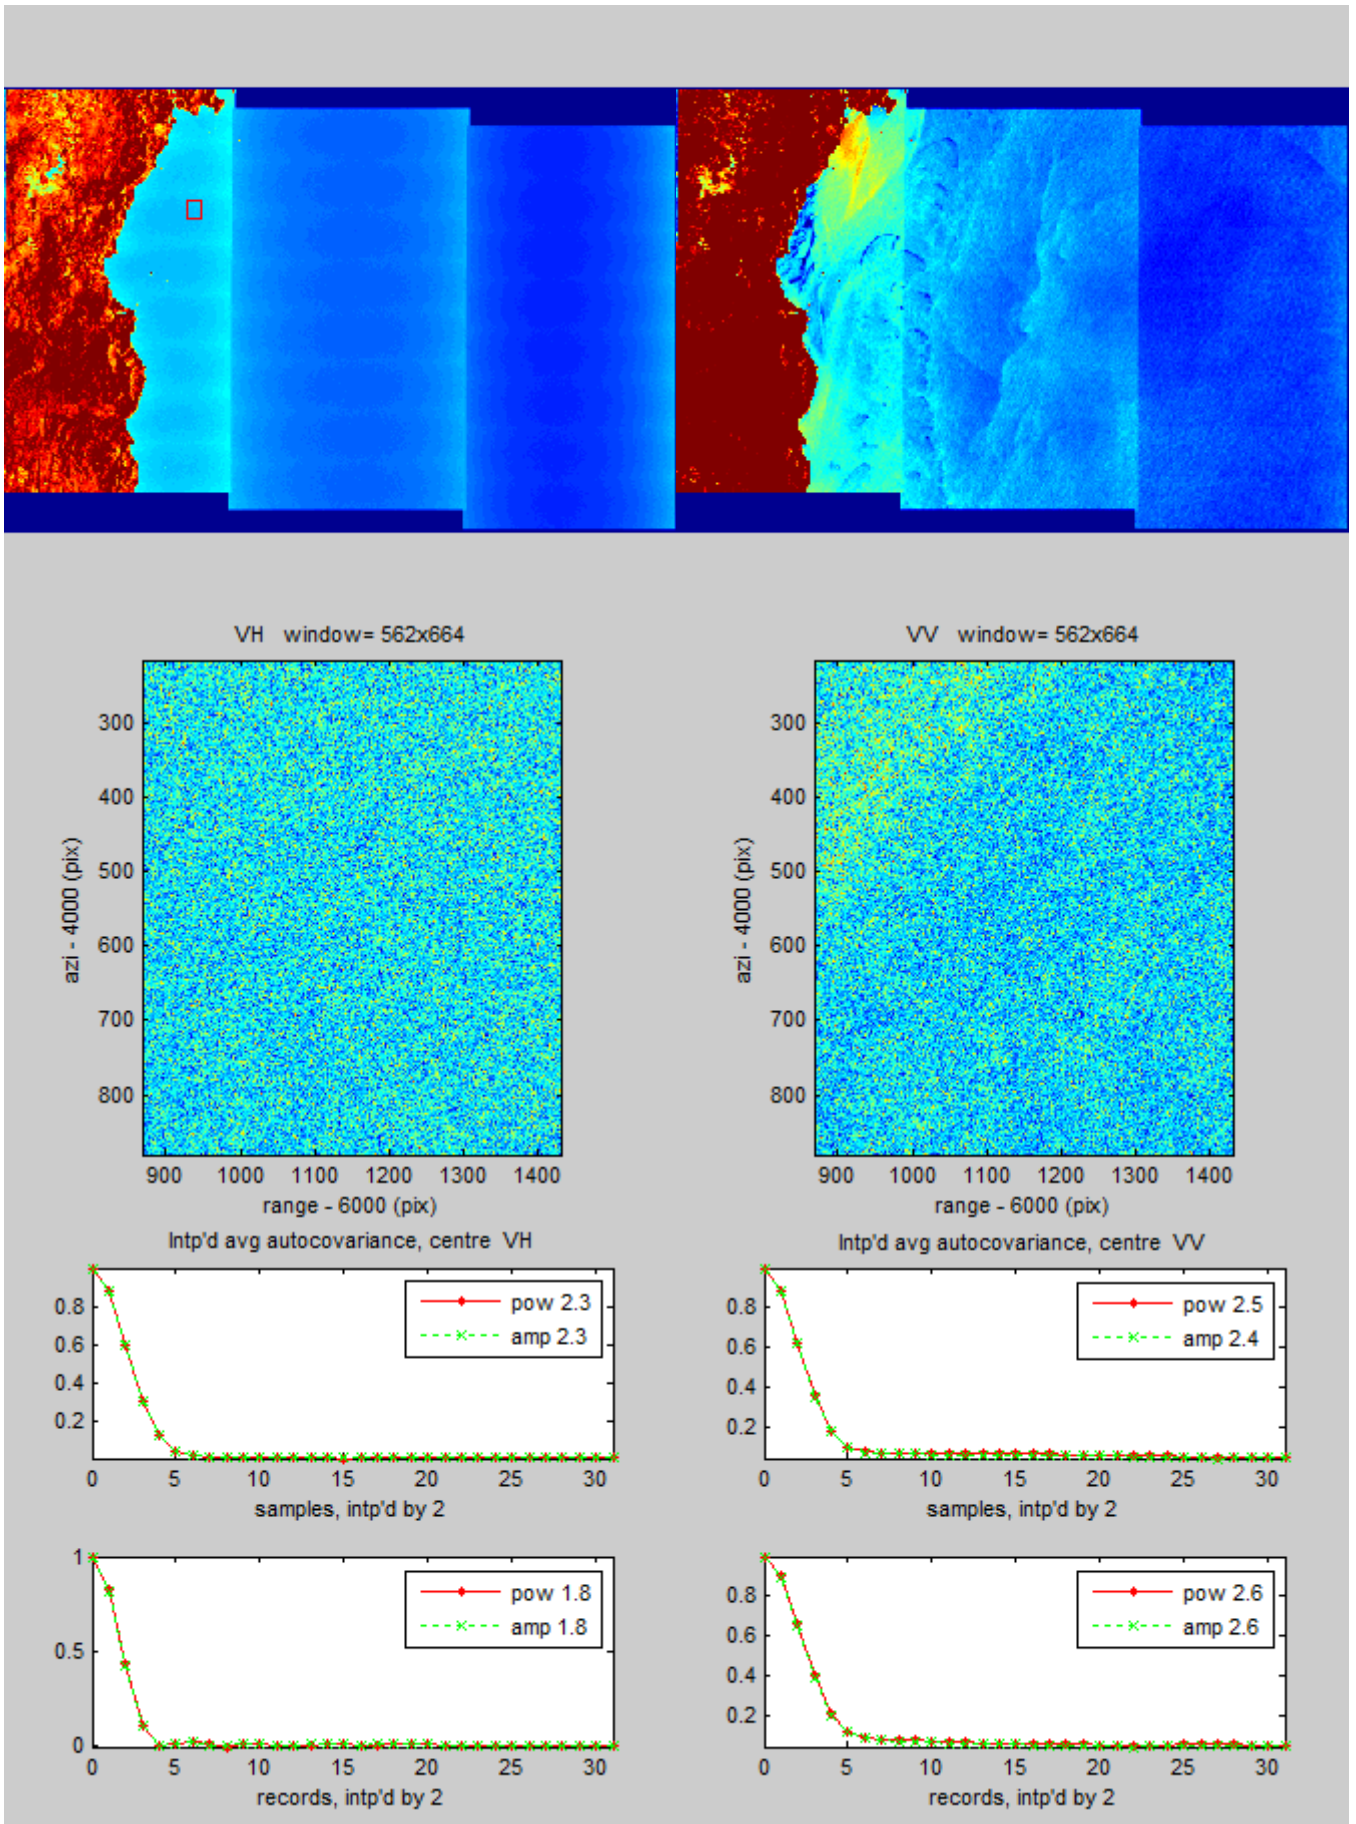
Fig. 2.9. Caption on previous page.

Fig. 2.10. As the previous figure, but now for an IW image. Polarisations VH and VV.
S1A_IW_GRDH_1SDV_20141030T084345_20141030T084410_003056_0037EC_7860

VH   window= 647x609   VV   window= 647x609

Intp'd avg autocovariance, centre VH   Intp'd avg autocovariance, centre VV
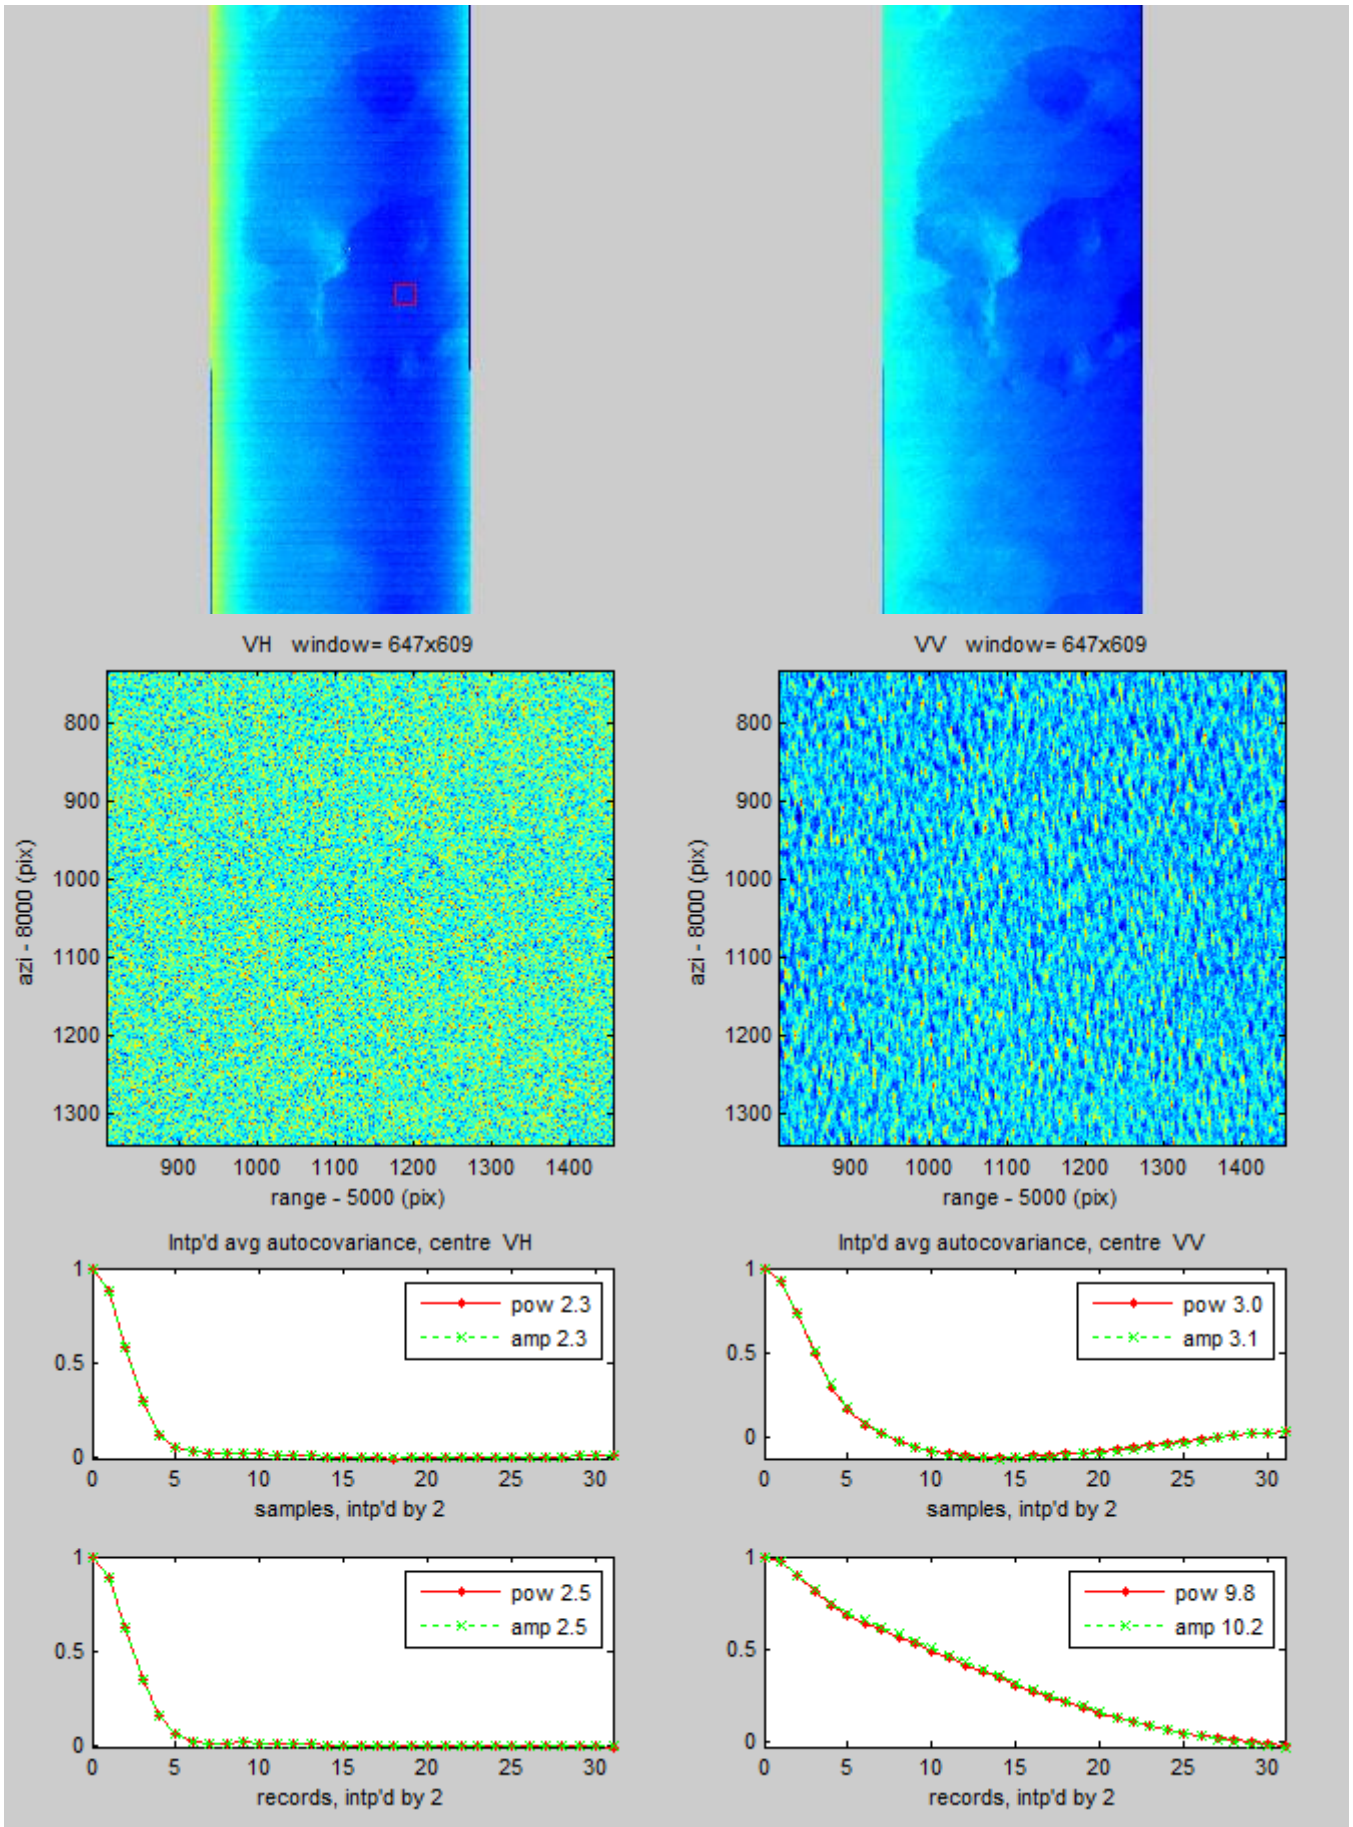
Fig. 2.11. Caption on next page.

(Fig. 2.11 caption:) As the previous two figures, but now for an SM (S1) image. Polarisations VH and VV. In this case, there is too much structure on the sea surface in VV (due to waves), so that the autocovariance function in VV is broadened and cannot be used to estimate the resolution. S1A_S1_GRDH_1SDV_20141018T022725_20141018T022753_002878_003419_F05C

Mode	Pixel size (m)	Direction	HH (pixels)	HV (pixels)	VV (pixels)	Best estimate (m)	Nominal (m)
EW	25	Range	1.9	1.8		46.0	50
	25	Azimuth	1.9	1.7		45.0	50
IW	10	Range		1.6	1.7	17.0	20
	10	Azimuth		1.3	1.9	15.7	22
S1	10	Range		1.6	(2.1)	16.2	21
	10	Azimuth		1.7	(6.9)	17.5	21

Table 2.4. Measured resolutions derived from the autocovariance function widths shown in the previous figures, compared to nominal values (from [1, 2]).

3   Geo-location

Geo-location is the process of converting the (pixel, line) coordinate in the image to geographical (longitude, latitude); and vice-versa. There are two ways to do this. First is using the geo-location grid that is provided in the metadata. The grid is a list of pixel, line numbers distributed over the image with their lon, lat coordinates. The second way is by using the satellite orbit data and the slant range to the point.

3.1   Geo-location with the provided geo-location grid

Although in the metadata xml file the geo-location grid is specified as a list of grid pixels for a certain grid line, and that for a set of grid lines, it is found that the list of grid pixels is identical for each grid line, thereby providing a regular grid. This greatly aids interpolation (although it leads to serious redundancy in the metadata.)

In addition, each grid line also specifies a height value ("geo-location grid height"), which impacts on the geo-location and which is found to be indeed different for each line in the grid. Height values are also provided in another place in the metadata, under "Reference height"; however, not on the same lines grid as the geo-location grid. The two sets of height values are consistent, seeming to derive from the same continuous function, however sampled at different (1D) grids. By comparing the two, for the cases inspected it is found that the geo-location grid heights can be interpolated from the reference heights; but the reference heights cannot be interpolated from the geo-location grid heights. See e.g. the figure below, where the red circles can be derived from the blue dots by interpolation, but not the other way around.
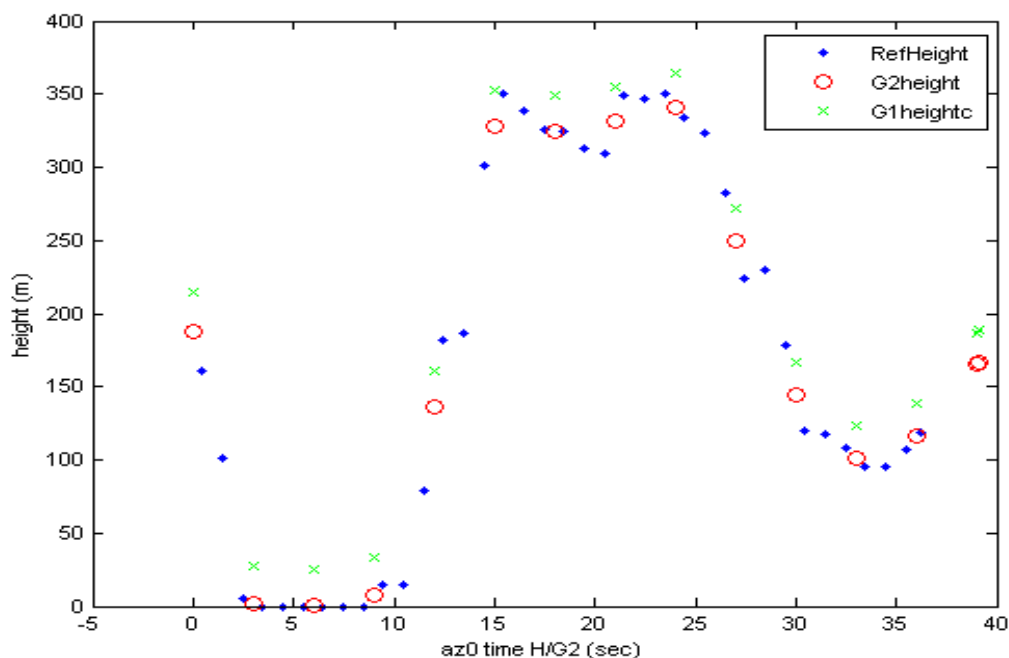

Fig. 3.1. Height values from the metadata of an S1 image. The red circles are the heights specified in the geo-location grid. The blue dots are the heights specified under the "Reference height" xml keyword. The two sets are consistent, but sampled differently, and the blue points cannot be derived from the red ones by interpolation while the inverse is probably true. The green crosses are heights calculated based on satellite orbit data provided in the metadata and a reference ellipsoid; they show a constant deviation to the specified blue and red heights.
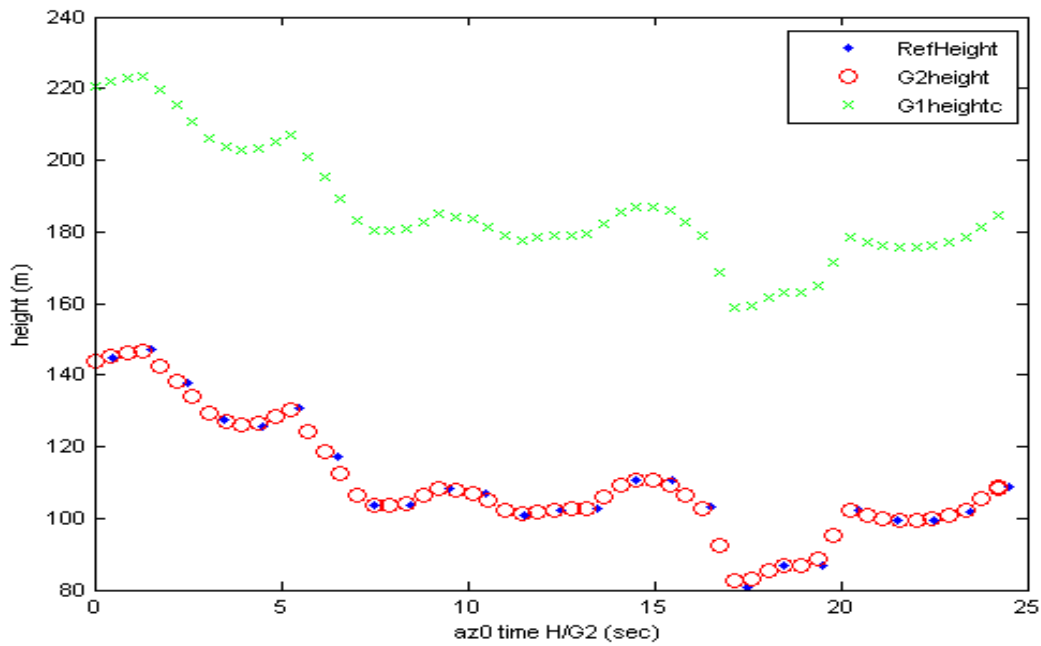*S1A_EW_GRDM_1SDH_20140812T222712_20140812T222751_001913_001D13_C802*

Fig. 3.2. Same as previous figure, but for another image. Now, the geo-location grid heights (red) are more densely sampled than the "Reference height" heights (blue). The two sets are consistent, and possibly the red points can be correctly derived from the blue points by interpolation.
S1A_S1_SLC__1SDV_20140607T172812_20140607T172836_000947_000EBD_4DB2

The (lon, lat) of an arbitrary (pixel, line) in the image must be found through interpolation of the geo-location grid. The interpolation can be linear, or e.g. spline for more accuracy, or a single polynomial surface can be fit to the entire grid. Any interpolation method needs a certain degree of smoothness of the grid, but this criterion can usually be met because the image is naturally smoothly draped over the Earth's (longitude, latitude) grid. However, for S-1 the geo-location is also determined by the height values that have been used and that change from line to line. It is found that this introduces a jumpiness that hampers accurate interpolation of the geo-location grid. The jumpiness is visualised in the figures below. The first figure shows that lon, lat depend on pixel (i.e. range direction) in the geo-location grid in a smooth way. The second figure below, however, shows that the lon, lat dependence on line (i.e. azimuth direction) is not smooth. The irregularities seem to be introduced by the varying reference heights used. The example shows one particular data set, but the same is found in all images of any mode. In addition, as a separate issue, it is often found that the last value of lon, lat as a function of line is very irregular; an error in the values is suspected.
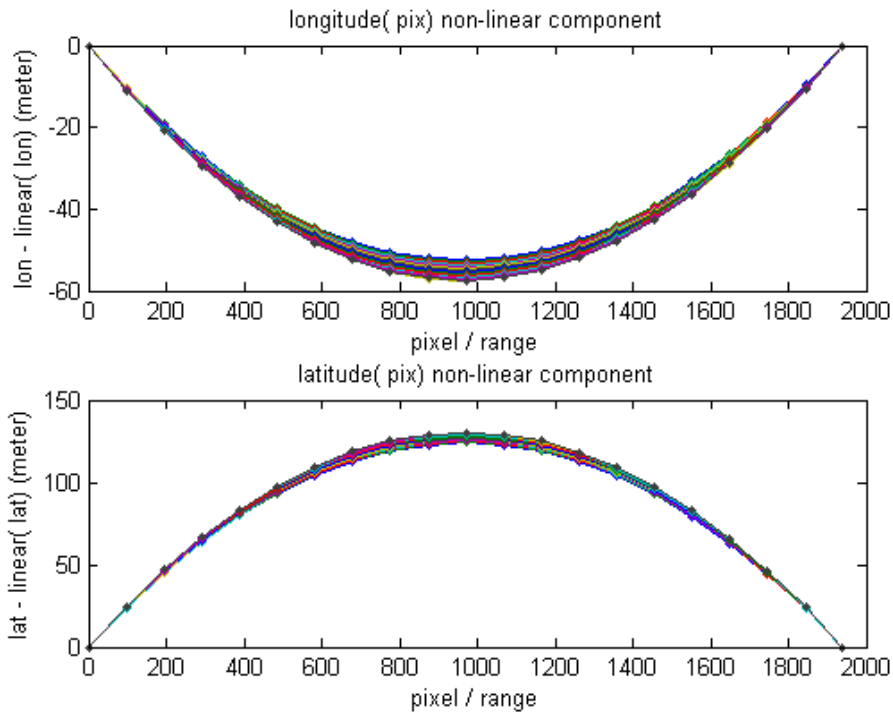
Fig. 3.3. Longitudes (top) and latitudes (bottom) specified in the geo-location grid as a function of pixel (which is range direction), with the linear component removed. Each of the two plot frames contains the graphs (as a function of pixel) for each grid line in different colours. The graphs are smooth, indicating that they can be interpolated, or modelled by a (at least 2^{nd} degree) polynomial. S1A_S1_GRDM_1SDV_20140607T172812_20140607T172836_000947_000EBD_7543
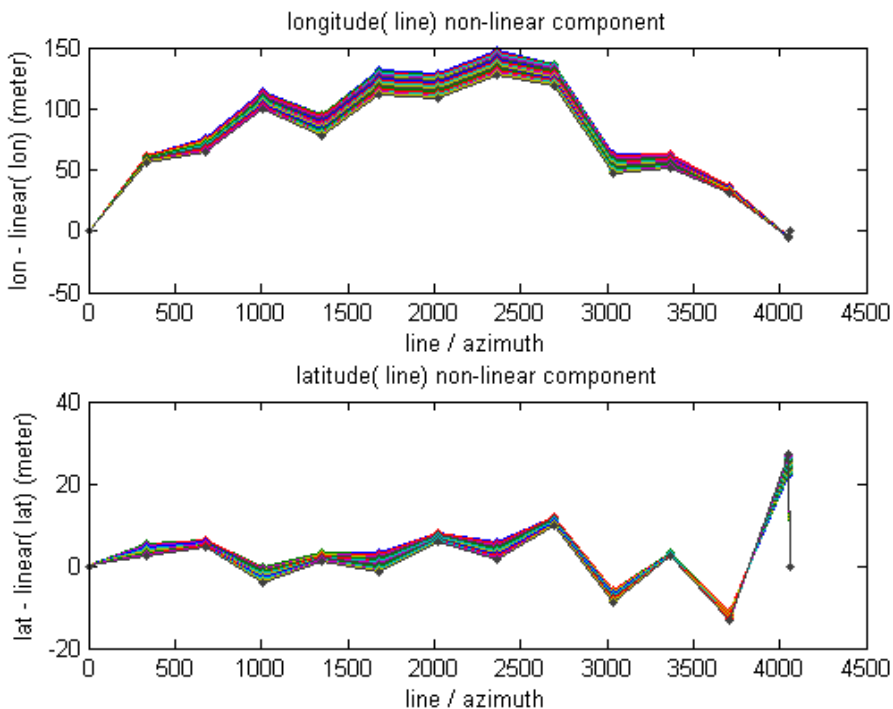

Fig. 3.4. Longitudes (top) and latitudes (bottom) specified in the geo-location grid as a function of line (which is azimuth direction), with the linear component removed. Each of the two plot frames contains the graphs (as a function of line) for each grid pixel in different colours. Contrary to the previous figure, these graphs are not smooth, indicating that they cannot be accurately interpolated, or modelled by a polynomial. Same image as previous figure.

23

The transformation from (pixel, line) to (lon, lat) can be done in a straightforward way by interpolation of the geo-location grid values, because the provided grid is regular in (pixel, line). In order to go the other way, the inverse transformation from (lon, lat) to (pixel, line), in an efficient manner, first a grid that is regular in (lon, lat) must be constructed. This is done by Delauney triangulation which uses linear interpolation. To have enough accuracy, the (pixel, line) grid that is being inverted must be sufficiently densely sampled. Therefore, the (pixel, line) grid is resampled to a denser grid before doing the inversion. How dense this resampling must be is assessed by doing an accuracy check in the following way. A dense grid of (pixel, line) test points is defined across the image. These are transformed to (lon, lat) coordinates, and those are in turn transformed back to (pixel, line) values; then, the difference between the initial (pixel, line) and the re-obtained (pixel, line) values is computed. The maximum occurring difference is not allowed to be more than 0.5 pixel or line. If it is more, then the geo-location grid is interpolated (using splines) to a yet denser grid before the inversion is made, and the process is repeated. The maximum number of iterations is set to 20, and the maximum number of interpolated grid points is set to 120. (The original geo-location grid is typically provided on 20 x 20 points order of magnitude depending on the size of the image.) Using this procedure, it is found that typically in the order of 5 iterations are needed to an interpolated grid size of the order of 80 x 80 in order to reach the 0.5 pixel / line accuracy criterion. However, in some cases this accuracy is never reached under the conditions mentioned. The figure below shows the remaining deviations at the test pixel grid for one example image.
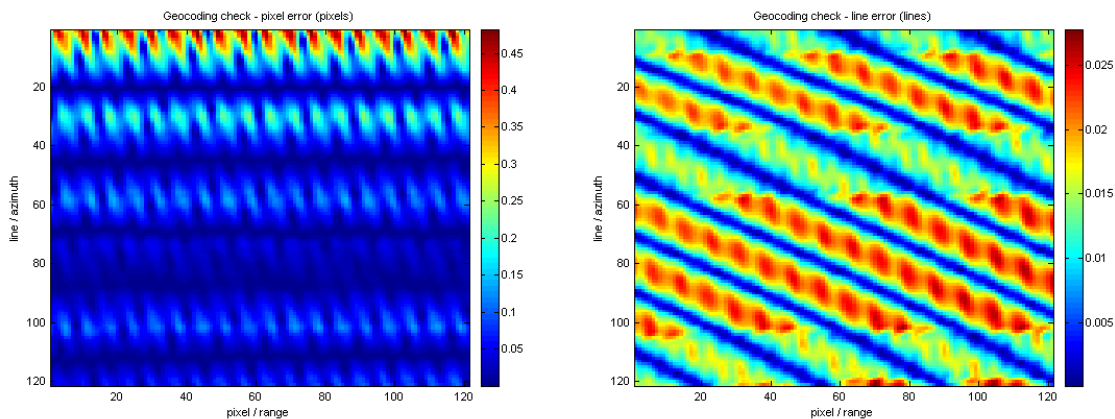

Fig. 3.5. Residual geo-coding errors measured from a test grid of 121 x 121 points after iteratively re-interpolating the provided geo-location grid until the residual errors are less than 0.5 pixel / line. Left, residual errors in pixels. Right, in lines.
S1A_IW_GRDH_1SSV_20141006T152706_20141006T152731_002711_003086_06DE.SAFE

Also the accuracy of doing the geo-coding using a fitted 2D polynomial is measured. A 3[rd] degree polynomial is fitted in the least-squared sense to the original geo-location grid, that can be used to go from (pixel, line) to (lon, lat). Another 3[rd] degree polynomial is fitted to go from (lon, lat) to (pixel, line). As before, the accuracy is measured by taking a test grid in (pixel, line), transforming with the first polynomial to (lon, lat), transforming back again with the second polynomial to (pixel, line), and measuring the deviations. This procedure, however, leads to very large deviations, due to the use of varying reference heights as discussed earlier. However, when the pixel locations of the original (pixel, line) geo-coding grid are corrected for the used reference heights, then the deviations become much smaller. The correction subtracts the reference height divided by the cosine of the incidence angle from the pixel location. Even so, the remaining deviations are unacceptably large (>> 1 pixel or line). The figure below visualises the deviations.
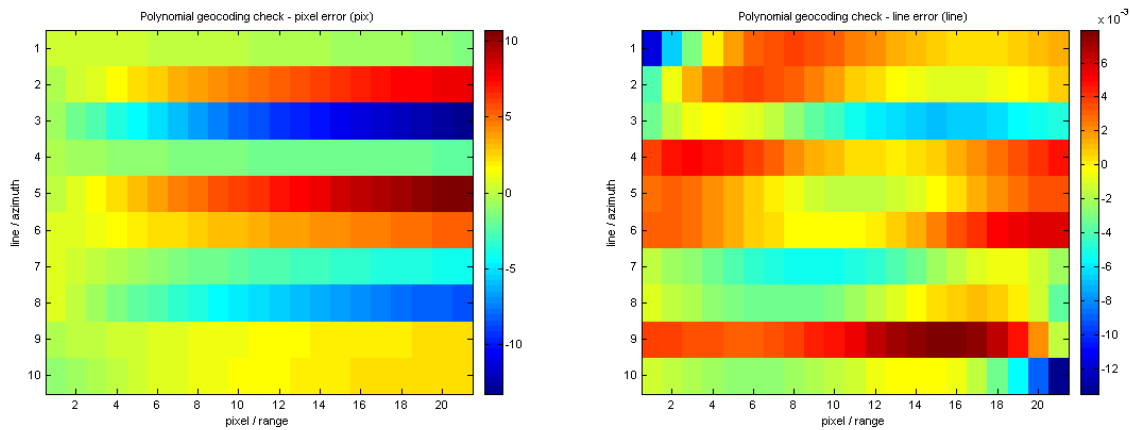
Fig. 3.6. Residual geo-coding errors measured from a test grid of 121 x 121 points using a 3rd degree polynomial model fitted to the geo-location grid. Left, residual errors in pixels. Right, in lines. Same image as previous figure.

From the above, it is concluded that the use of a polynomial to do geo-location on the S-1 images is not accurate enough. Interpolating with a tie points grid can often be accurate enough but only after re-interpolation of the provided tie-points grid; but in some cases, also that is not accurate enough.

Geo-coding of SLC images with a tie points grid is further discussed under the Section 6.

3.2   Geo-location in SUMO

SUMO uses an approximate method for geo-location, because for ship detection, the exact location of a ship is in many cases not so important: the ship is moving, so in a short time its position will have changed anyway. More importance is given to processing speed instead. With S-1 images, accuracies are found to be not better than a few hundred meters. Absolute position accuracy was measured by inspecting easily recognisable features in the landscape, by comparing the (lon, lat) that SUMO gives on that feature to the coordinates of the same feature as seen in Google Earth. In addition, the coordinates given by NEST, the toolbox provided by ESA, were included in the comparison. An example is shown in the figure below.
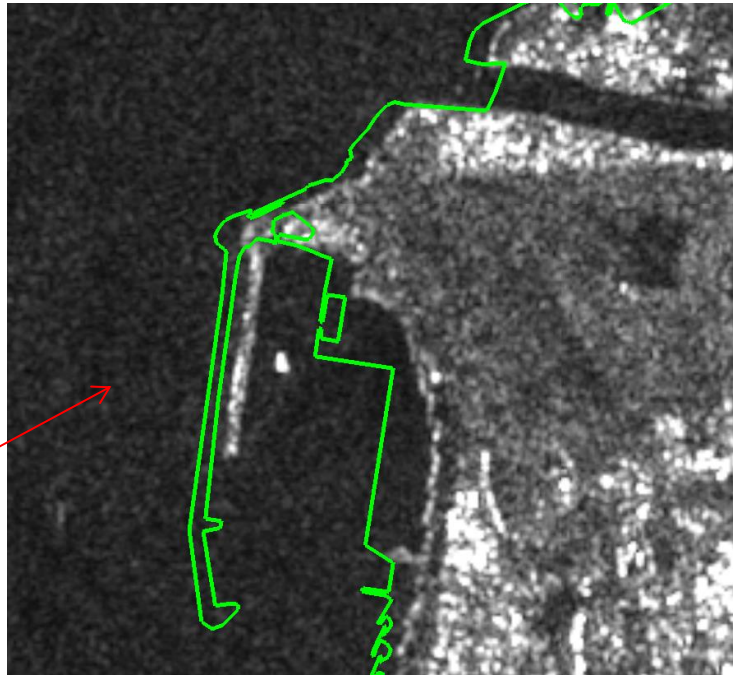
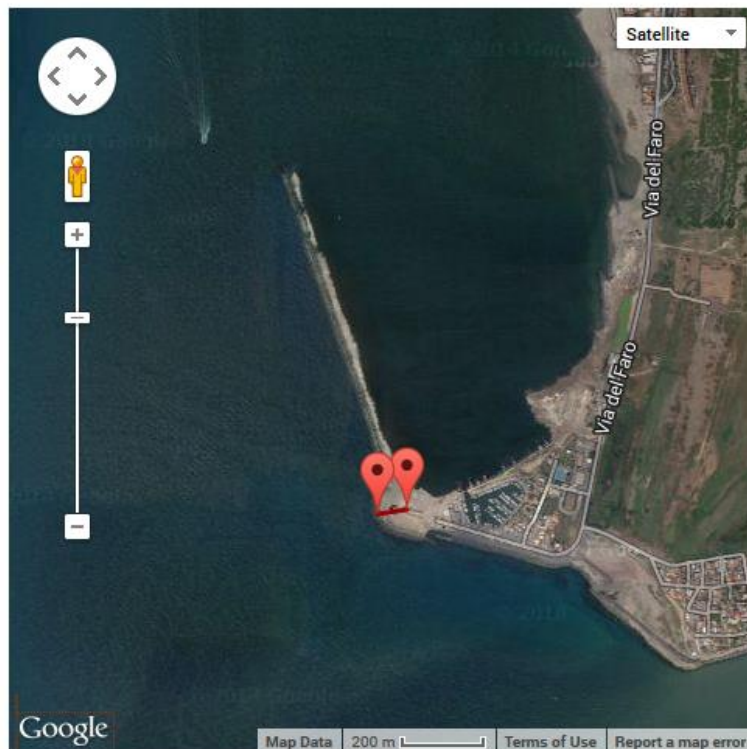*Fig. 3.7. A harbour protection pier near Ostia, Italy. Top: seen in an S-1 image viewed in SUMO, with the OpenStreetMap coastline overlaid in green. Bottom: seen in Google Earth (GE) and indicating the coordinates of the feature as seen in SUMO and as seen in GE. In this case, their distance is 68 meter. Screenshot from http://www.movable-type.co.uk*
S1A_IW_GRDH_1SDV_20140716T170245_20140716T170310_001516_00176D_FCF8

This type of measurement was done on a number of landscape features in various S-1 images. It is found that the NEST coordinates agree very well with GE coordinates. The SUMO coordinates, however, may deviate up to several hundreds of meters. Although this is within the original SUMO specs, that sacrificed accuracy over speed, it should be considered to improve SUMO's geo-location accuracy, as processing speed is becoming less of a bottleneck due to increased availability. An improved SUMO geo-location algorithm must include the correct application of the terrain reference height used in the processing.

3.3   Geo-location using orbit data

This geo-location method makes use of the orbit state vectors (position and velocity) and timing references in the image metadata to solve the range and Doppler equations. In the case of reverse geocoding (i.e. finding the image coordinates for a given geographic coordinates), the steps to geo-locate a point P with geographic coordinates (lat, lon and height) are [4, 5]:

- Interpolate the orbit state vectors. In the case of S-1, the metadata presents one state vector every second. This sampling rate has to be increased to a rate that is similar to the image sampling rate (PRF). The interpolation method that has been implemented is based on Hermite polynomials [6]. For reference, MDA suggests a cubic spline interpolation for Radarsat-2 images [4], and for Terrasar-X images the recommended interpolation method (generally Chebyshev) is stated in the metadata (<recProcessingTechnique> field).
- Convert the point P's geographic coordinates to Earth-Centred Earth-Fixed coordinates.
- Compute the Euclidean distance between P and each position in the interpolated orbit (slant range distances). This is the "range history".
- Compute the position and value of the minimum distance in the range history. This is the zero Doppler position (position of closest approach).
- The line (or record) image coordinate is calculated from the position of the minimum distance. For S-1, TerraSAR-X, Cosmo-SkyMed and Radarsat-2 the metadata contains the times of the orbit state vector and of the first and last line of the image in the zero Doppler reference frame, so this step can be done by simple interpolation.
- The pixel (or sample) image coordinate is calculated from the value of the minimum distance, using the near range distance and the range sampling distance (both obtained from the image metadata).
- If the image is Ground Range Detected (GRD) an additional intermediate step is needed to convert the minimum slant range distance to ground range distance. This should be done using the Slant to Ground polynomials provided in the metadata. However those polynomials have not proved to be very accurate. As a consequence, the Ground to Slant polynomials (also from the metadata) are inverted and used instead. S-1 gives a set of polynomial coefficients every second in azimuth time; to achieve good geolocation accuracy these coefficients have to be linearly interpolated at the point's zero Doppler time.

The forward geocoding (i.e. finding the geographic coordinates for a given image coordinate) is slightly more difficult and includes some additional geometric calculations [4].

Both the forward and the reverse geocoding using orbit state vectors have been implemented in MATLAB and have proved more accurate at geo-locating features than the method based on tie points. The evaluation is presented next. The steps in the evaluation have been:

- Identify the same salient coastal features on Google Earth and on SAR images. 62 such features from 38 different S-1 images have been identified. Most of the images were of the IW mode, with a few ones being EW and SM. Only images acquired after the entry into operational phase of S-1 (3 October 2014) have been used in this evaluation. All the products are Ground Range Detected (GRD).
- For each feature, extract the geographic coordinate from Google Earth and the image

coordinate (i.e. line and pixel) from SUMO. Use the coordinate from Google Earth as ground truth.

- In the case of reverse geocoding, convert Google Earth's geographic coordinate of the feature (lat and lon, height is obtained from a geoid model) into image coordinates (line, pixel) using the geocoding based on orbit data, and compare the computed image coordinates with the ones observed in SUMO.
- In the case of forward geocoding, convert SUMO's image coordinates of the feature into geographic coordinates using orbit data and tie points, and compare the computed geographic coordinates with the ones observed in Google Earth (the ground truth).

Fig 3.8 presents a scatterplot with the reverse geocoding error (in meters) of the 62 features using the method based on orbit data. It shows most of the points clustered closely around the origin. 84% of the points are inside the bounding box with corners (-25,-25), (-25,25), (25,25) and (25,-25) m. The geolocation accuracy (in meters) in EW images tends to be lower than in SM images, due to the lower resolution of the EW mode. No noticeable difference in the geolocation accuracy is observed depending on the orbit direction (ascending or descending); this is not represented in the scatterplot.

Figs 3.9-3.10 present two different zoom levels of the scatterplot with the forward geocoding error (in meters) of the 62 features using the orbit (in blue) and tie points (in red) methods. It is clear from the plots that the errors using the tie points methods are more spread out than those using the orbit method. In fact, some red points are outside the limits shown in the figure (the largest such error is more than 8000 meters in longitude). In this case, around 79% of the blue points but only 10% of the red ones are inside the 50x50 meters box centred at (0,0). It has to be noted as well that a significant amount of red points have geolocation errors in the order of several hundreds of meters (especially in longitude).

It is important to realise that the exact geolocation of the features used in this evaluation is not known (i.e. they are not known corner reflectors or transponders), hence the errors represented in the scatterplots can be caused by different sources, one of them being inaccuracies in the geocoding processes. Other possible sources of error are human error in visually identifying the exact same feature on Google Earth (optical image) and on the SAR image, changes in the scene between the acquisition times of the Google Earth image and the SAR image (e.g., tides), or inaccuracies in the image metadata provided. Additionally, no attempt has been made to try to find the features' location on the SAR image with sub-pixel accuracy, as that level of accuracy is deemed unnecessary for maritime surveillance. The locations are therefore always taken at the centres of pixels. It is therefore not possible to isolate the part of the error related solely to inaccuracies in the geocoding process, but it is probably safe to assume that this component of the error is smaller than the overall error seen in the scatterplots.
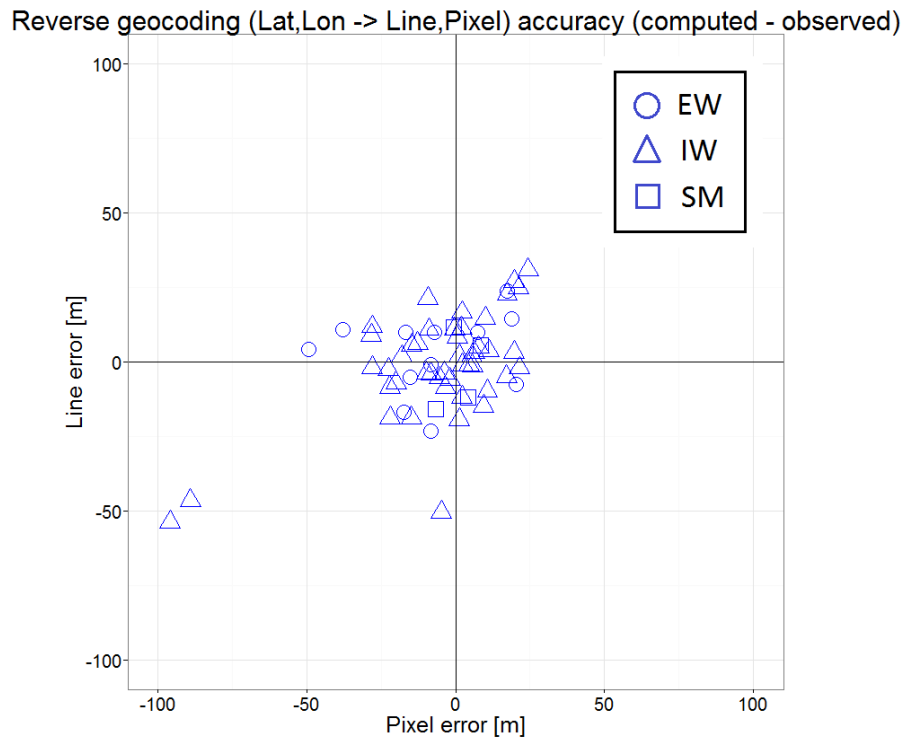
Fig 3.8. Reverse geocoding error using orbit state vectors. 62 points in 38 S-1 GRD images.
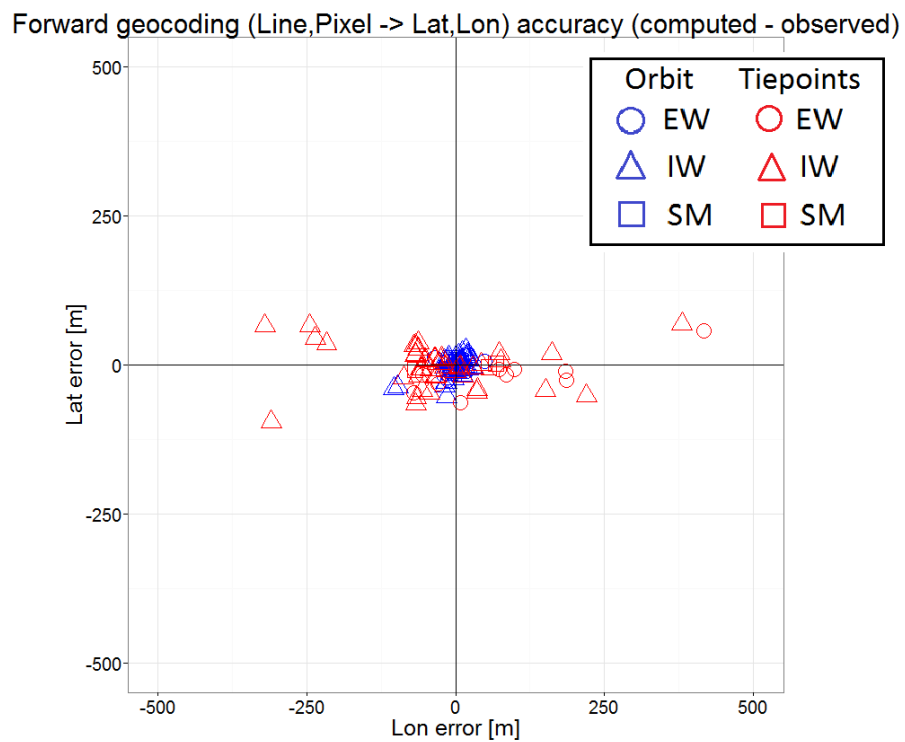

Fig 3.9. Forward geocoding error using orbit state vectors (in blue) and tie points (in red). 62 points in 38 S-1 GRD images. Some red points are outside the scatterplot's limits.
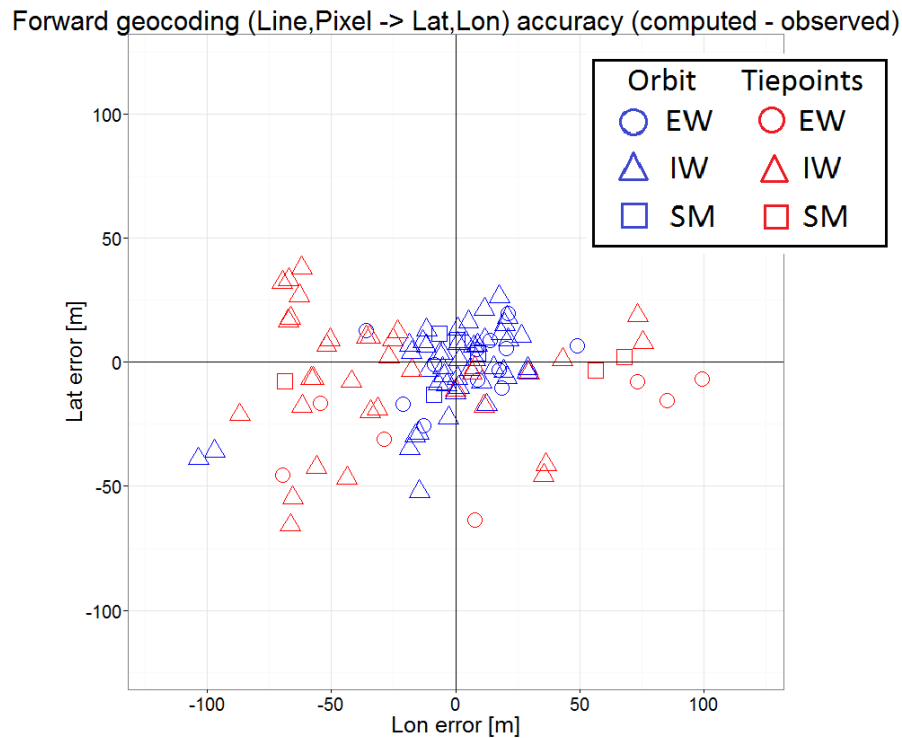
Fig 3.10. Zoomed in on previous figure.

Outside this evaluation on GRDH products, one feature in a IW Single Look Complex (SLC) image has been geocoded, yielding very accurate results. Geocoding of GRD products is considered as slightly more challenging than geocoding of SLC products, as it includes the extra step of converting a distance from slant range to ground range (or vice versa). In the geocoding of S-1 SLC TOPS (IW and EW) images, the metadata bursts timing parameters (explained in Section 6) need to be taken into account.

The recommended way to use the geolocation based on orbit data for ship detection (SUMO) is:
- Reverse geocode only the points of the coastline shape file. The points are therefore converted to image coordinates, which SUMO uses to mask out the land.
- VDS analyse the image as usual by SUMO.
- Forward geocode the detected VDS targets to find their geographic coordinates.

The first step is expected to give very precise land masking results as long as the coastline shape file is accurate enough. It should also have a low computational cost because only a small number of points need to be geocoded (the most complex and detailed shape file is likely to have fewer than 50,000 points within a SAR image boundary). It will however not give precise geocoding in the rare instances where coastal features are affected by layover. Examples of such features are the rock of Gibraltar and the cliffs on the west part of Lampedusa. If deemed necessary, a few points of a DEM around the coastline should be geocoded in order to increase the land masking accuracy in those sections of coastline.

4 Ship detection

4.1 Expected performance

Prior to the launch of Sentinel-1, studies had been made to estimate its ship detection performance, also in comparison to existing satellite SARs used for that purpose such as Radarsat-2 and ENVISAT-ASAR. In particular [7] from 2008 presented a very detailed analysis with an emphasis on the IW mode, also reported in [8]. The main results are reported in the form of minimum detectable ship length as a function of incidence angle for the various Sentinel-1 modes and polarisations. These predictions have not yet been quantitatively verified – indeed, such a verification is very hard to do due to the high level of randomness of ship radar backscatter and the amount of (system, target and ambient) parameters that influence the detection performance.

It can be noted, however, that Sentinel-1's system parameters as they have turned out are quite similar to what was used in the quoted study. Today's specified NESZ [2] stands at < –26.3 dB which is somewhat better than the –24 to –23 dB used in [7] for the IW mode. Today's quoted IW SLC resolution of 2.7 to 3.5 m in slant range (corresponding to 5 m ground range) by 22 m in azimuth [2] on the other hand is slightly worse than the 2.4 to 3.2 m slant (corresponding to 4.5 m ground) by 19.5 m used in [7]. The lower resolution should lower the detection performance somewhat, but the difference is marginal. The better NESZ should lead to somewhat improved detection in cross-polarisation where the detection is noise-limited; in co-pol, where the detection is limited by the sea clutter, the impact should be less.

More relevant than these small differences, however, is the fact that the 2008 study used the single-look parameters for its assessment. While by itself correct, most IW products over the sea are made available not as SLC but as GRDH, so it is important to take that into account. The 2008 report provides a tool to do that as it includes a sensitivity analysis on exchanging resolution for ENL, as is done in going from SLC to GRDH. Its Table 16 shows that whereas a 4.4 m ground range / ENL 1 product (corresponding to IW SLC as used in the 2008 study) has a minimum detectable ship length of 31.4 m according to the simulations, a 22.0 m ground range resolution / ENL 5 product made from the same raw data (approximately corresponding to today's IW GRDH) has a 41.9 m minimum detectable ship length. This figure presents still a significant improvement over previous systems.

4.2 Automatic ship detection

For ship detection, the JRC's SUMO software is used. SUMO can ingest, display and analyse SAR images of the most common SAR satellites: Radarsat-1, Radarsat-2, ENVISAT-ASAR, Cosmo-SkyMed, TerraSAR-X, and now also Sentinel-1. It takes considerable effort to adapt SUMO to each new satellite, because the data of each satellite comes in a different format. The adaptations for Sentinel-1 were done in 2014. In order to do the ship detection, SUMO overlays a land mask, divides the image in tiles, and for each tile computes the sea surface statistics (mean and standard deviation) and from those – through the use of a look-up table – the CFAR (Constant False Alarm Rate) threshold above which pixels are deemed to be part of a target. The threshold can be raised or lowered to make the detection process less or more sensitive. Then it clusters all neighbouring target pixels into detected ships. In the case of dual-pol images, it does so separately for each polarisation, and then combines the results into a single set of detections. For each detected ship, the geographical position (lon, lat), size (length, width), orientation, target strength and detection reliability are computed. The detected ship positions are displayed, and the results can be output in different formats (csv, xml, kml). This is done fully automatically. An example of a SUMO result is shown in the figure below.
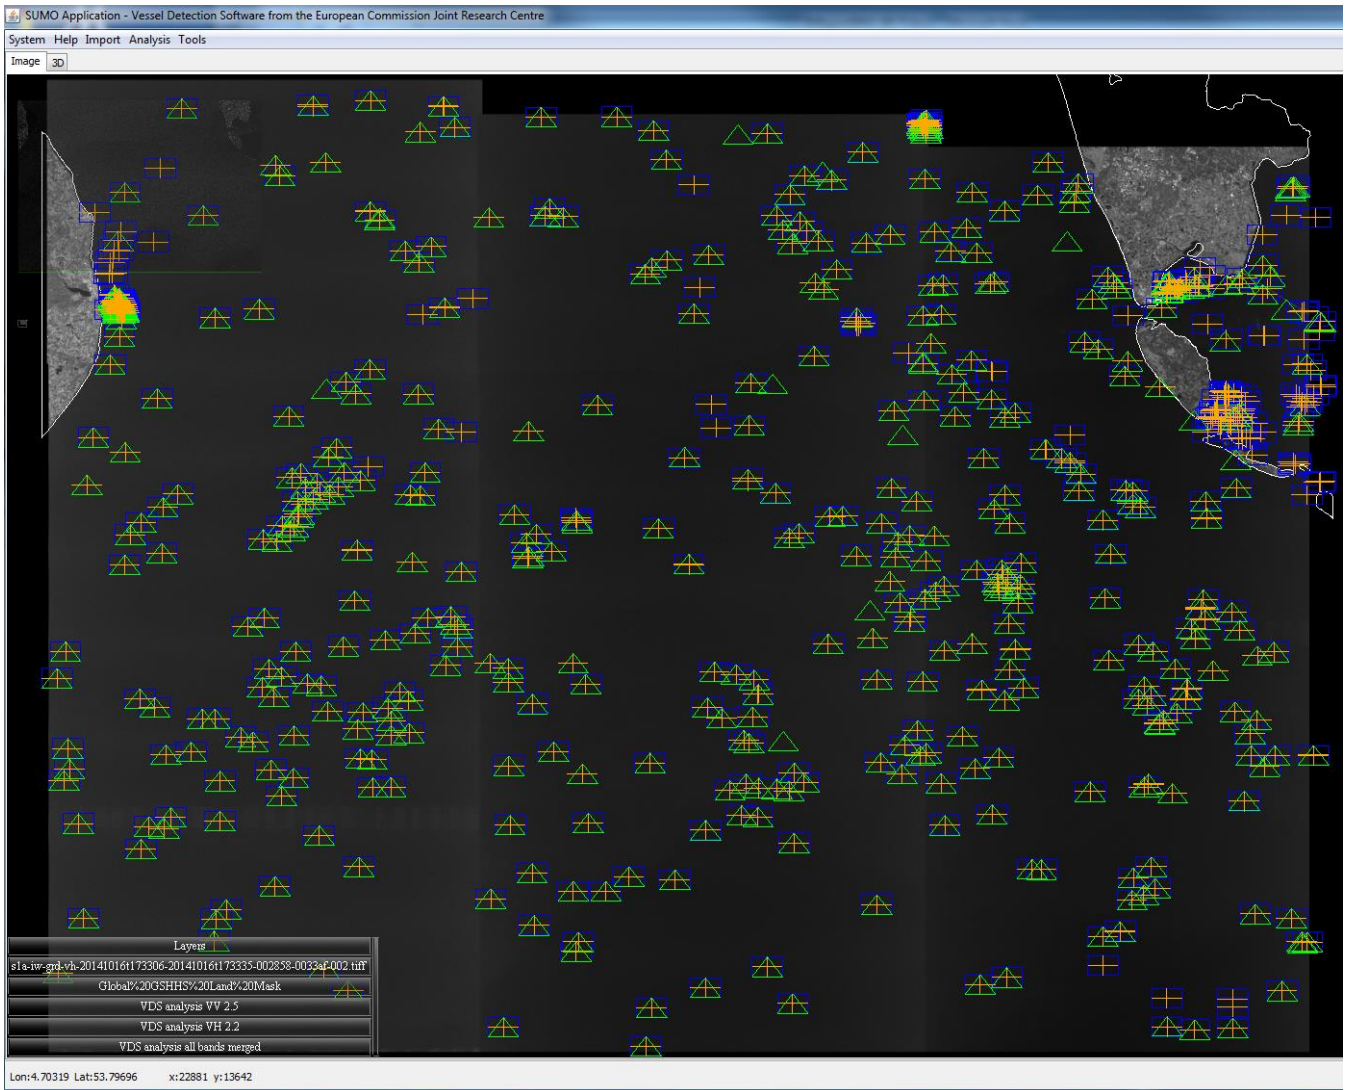
Fig. 4.1 SUMO automatic ship detection result over an IW image of the North Sea. The image is in natural coordinates which makes it appear, in this case, upside-down in comparison to a normal map. The green triangles are targets detected in the VH channel. The blue rectangles are detections in the VV channel. The orange crosses are merged from the two, and thereby form the final result. This is a fully automatic result, and the total number of detected targets is 614, which is very high for one image.
S1A_IW_GRDH_1SDV_20141016T173306_20141016T173335_002858_0033AF_FA6D

4.3 Manual verification of detected ships

The results of the fully automatic detection can be inspected manually in order to discard some false alarms that still occur. Today, in most types of satellite SAR imagery including S-1, the false alarms are mostly due to: coastal features not properly indicated on the coastline map (such as shallows, reefs, piers); spurious detections on ship wakes, especially in cross-pol; unrecognised range or azimuth ambiguities (see Section 5); strong local oceanic or atmospheric effects such as rain cells or fronts; and edges of very low wind areas. A human operator is able to recognise many of these effects and flag them as false alarms, whereas the automatic detection algorithm is not yet reliable enough for these tasks. The SUMO output after operator validation of the previous figure is shown in the figure below. The figures after that show some examples of ships and false alarms in close up.
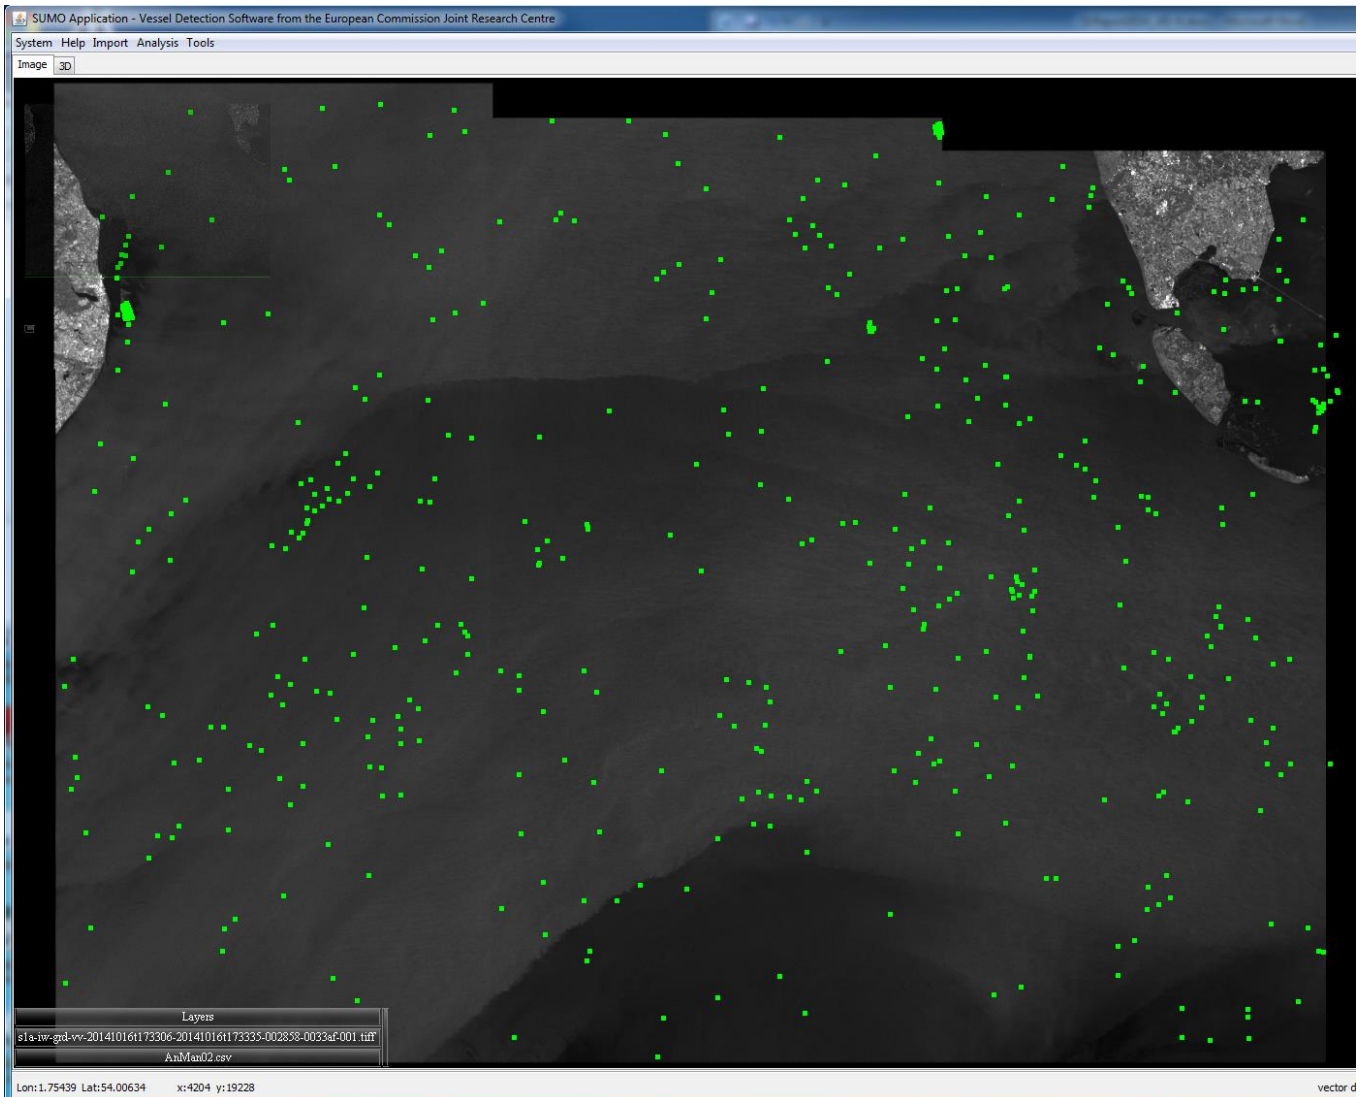
Fig. 4.2. The ship detections from the previous figure after manual rejection of false alarms based on visually inspecting each target. The remaining number of validated detections is 507. Many false alarms occurred on the tidal flats off the North of Holland (top right). Same image as previous figure.
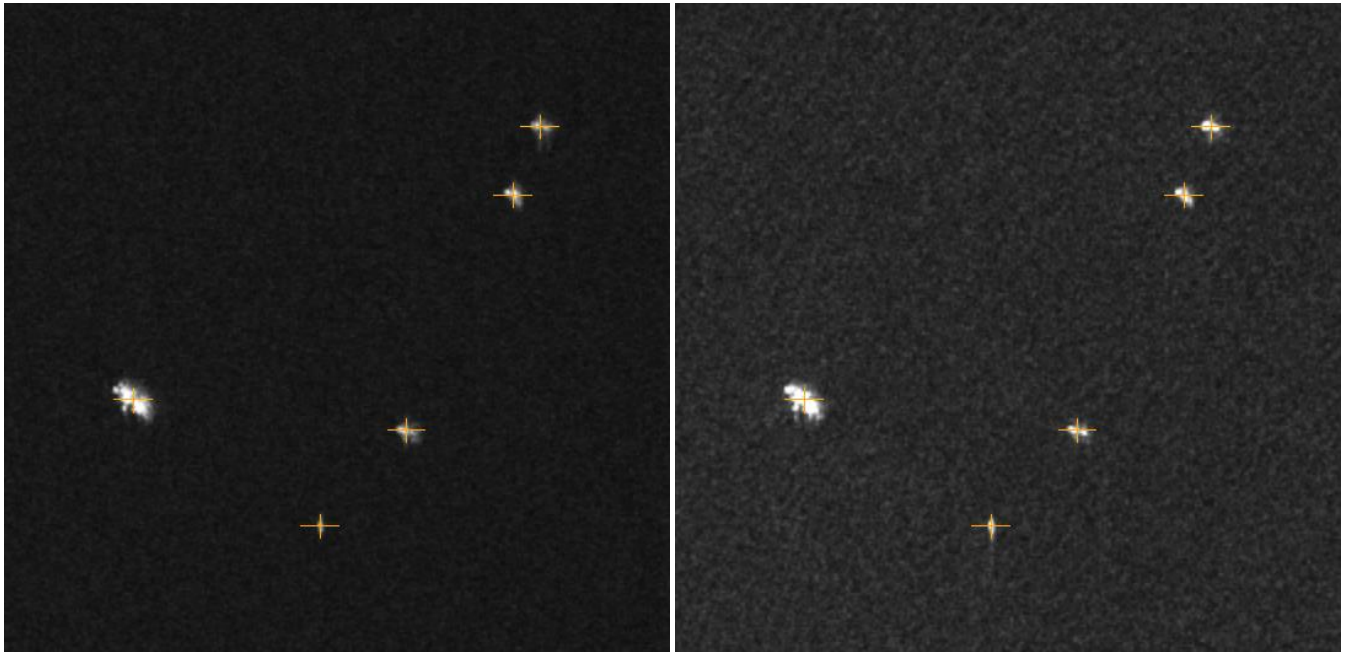
Fig. 4.3. Zoom-in at full resolution showing five nearby ships of different sizes against the sea surface background. Left in VH; right in VV. The crosses are the SUMO automatic ship detections merged form the VH and VV channels. Same image as previous figures.
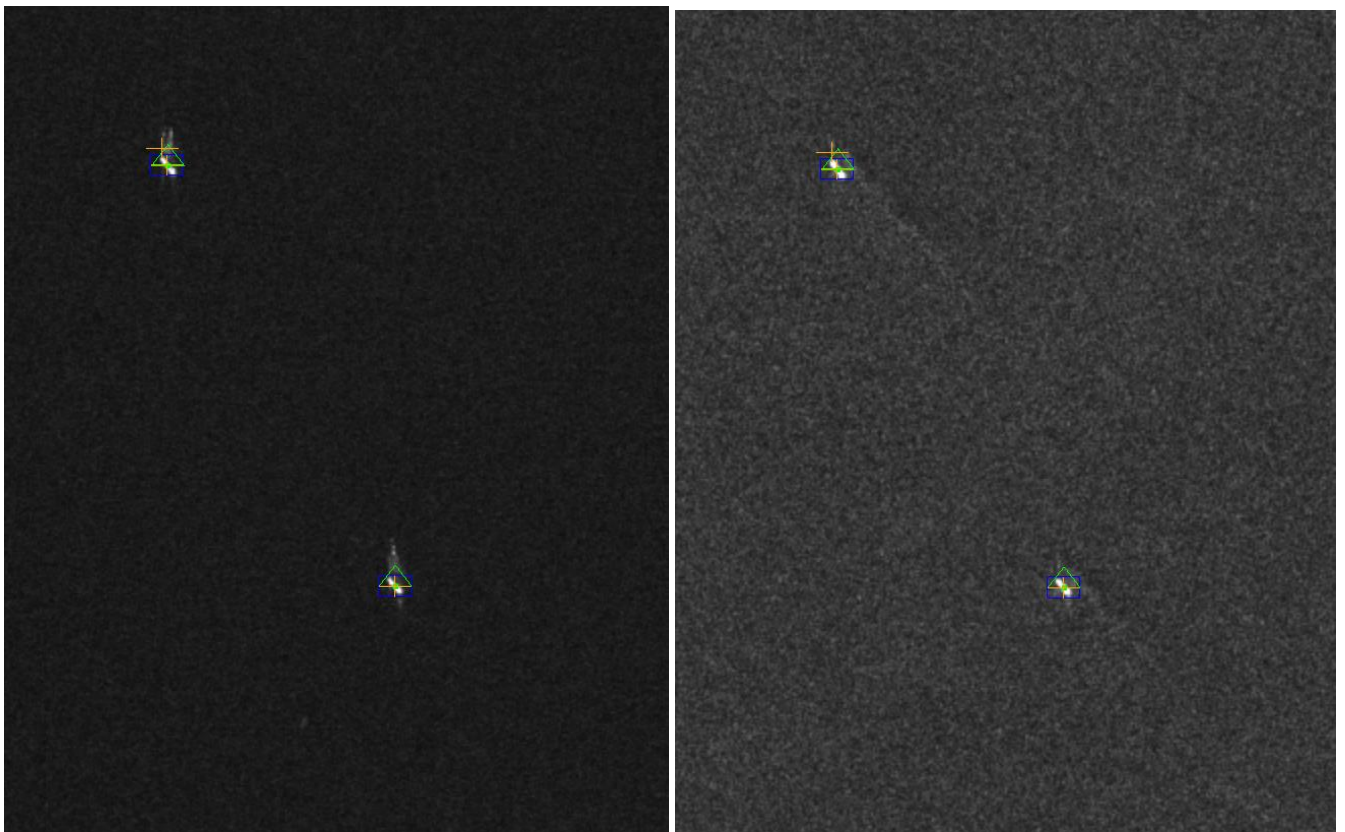

Fig. 4.4. Zoom-in at full resolution showing two ships. Left in VH; right in VV. The green triangles are the SUMO automatic detections in VH; the blue rectangles in VV; and the orange crosses the merged result from the two. The ship wake signatures show up in VH as bright smudges above the actual targets; in VV they are fuzzy lines extending from the targets to the bottom right indicating the ship's track. For the upper ship, a spurious detection on the VH wake created two ship detections (two orange crosses) instead of one; one of these is a false alarm and is removed in the manual verification. For the lower ship, the automatic result was correct. Same image as previous figures.
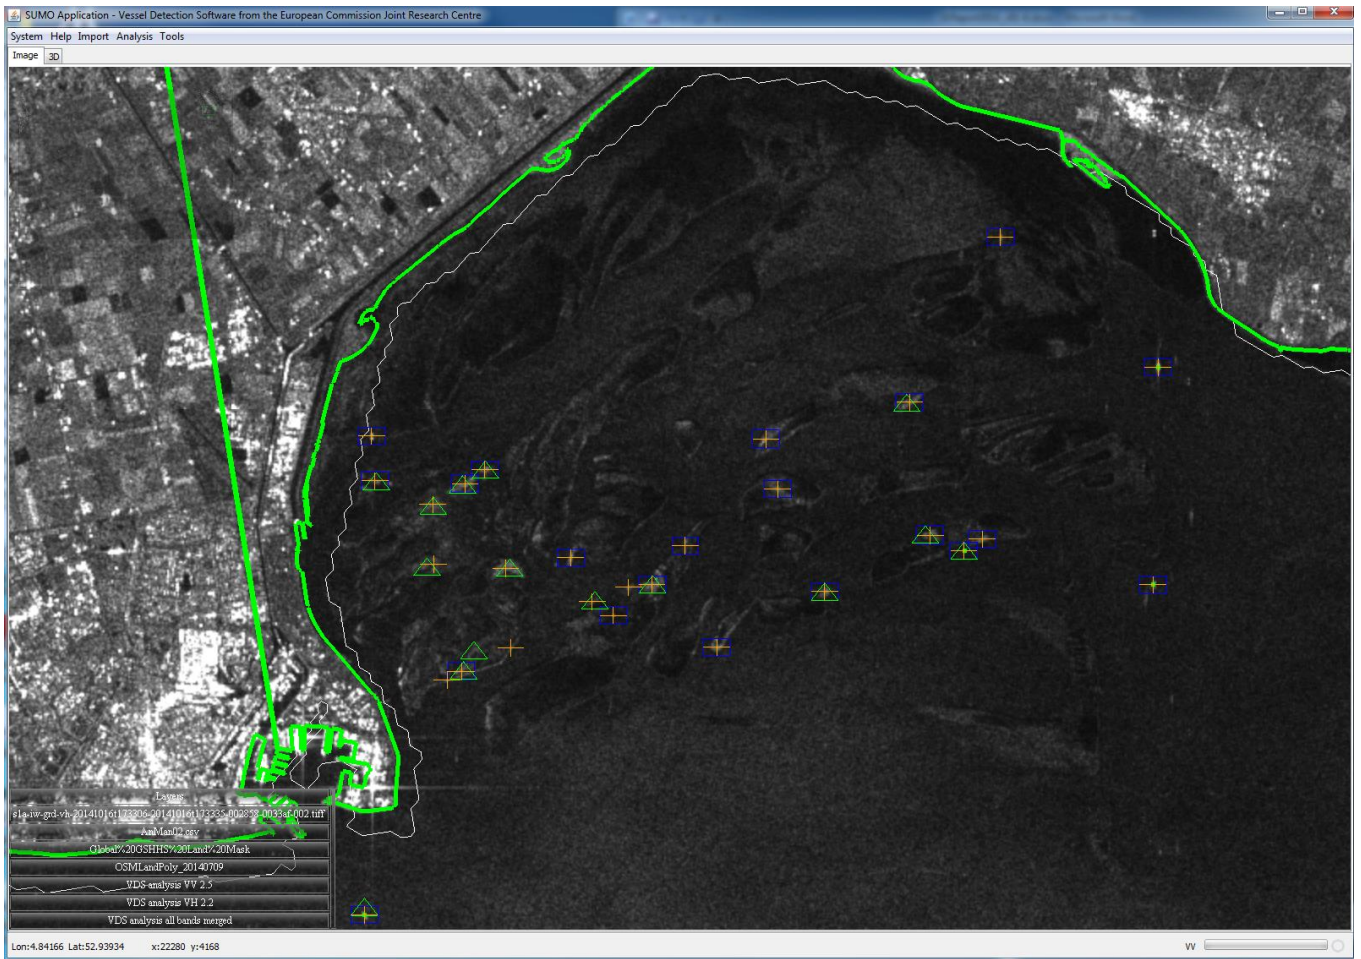
Fig. 4.5. Example of false alarms on tidal flats. The VV channel is shown. This is next to Den Helder, The Netherlands. Same image as previous figures.

5 Ambiguities and ghosts

5.1 Azimuth ambiguities and ghosts

SAR imagery always suffers from ghosts, which are repeated occurrences of bright targets, displaced from the target location. The displacement direction is often nearly exactly in range or azimuth. Most common ghosts are "range ambiguities" and "azimuth ambiguities". The former are weakened and distorted repeats of a bright target displaced in range direction, at a distance of one to several hundred kilometres. The source of the range ambiguity that is seen in an image is therefore often outside the image. Azimuth ambiguities are similar ghosts but displaced in azimuth at a distance of some 4-8 km. Often, the source of the azimuth ambiguity can be readily found inside the image. It can be a bright ship, but also a bright reflector on land. Both range and azimuth ambiguities occur at known distances, that depend on parameters such as the radar PRF, platform velocity and target range distance. They can occur repeatedly at integer multiples of the primary distance.

In S-1 images, range and azimuth ambiguities are found for very bright targets, as expected according to the system specs. However, also another type of ghost is found, which is similar to the azimuth ambiguity as it is displaced and repeated in azimuth, but occurs in addition to it and at another distance. The figures below show examples.
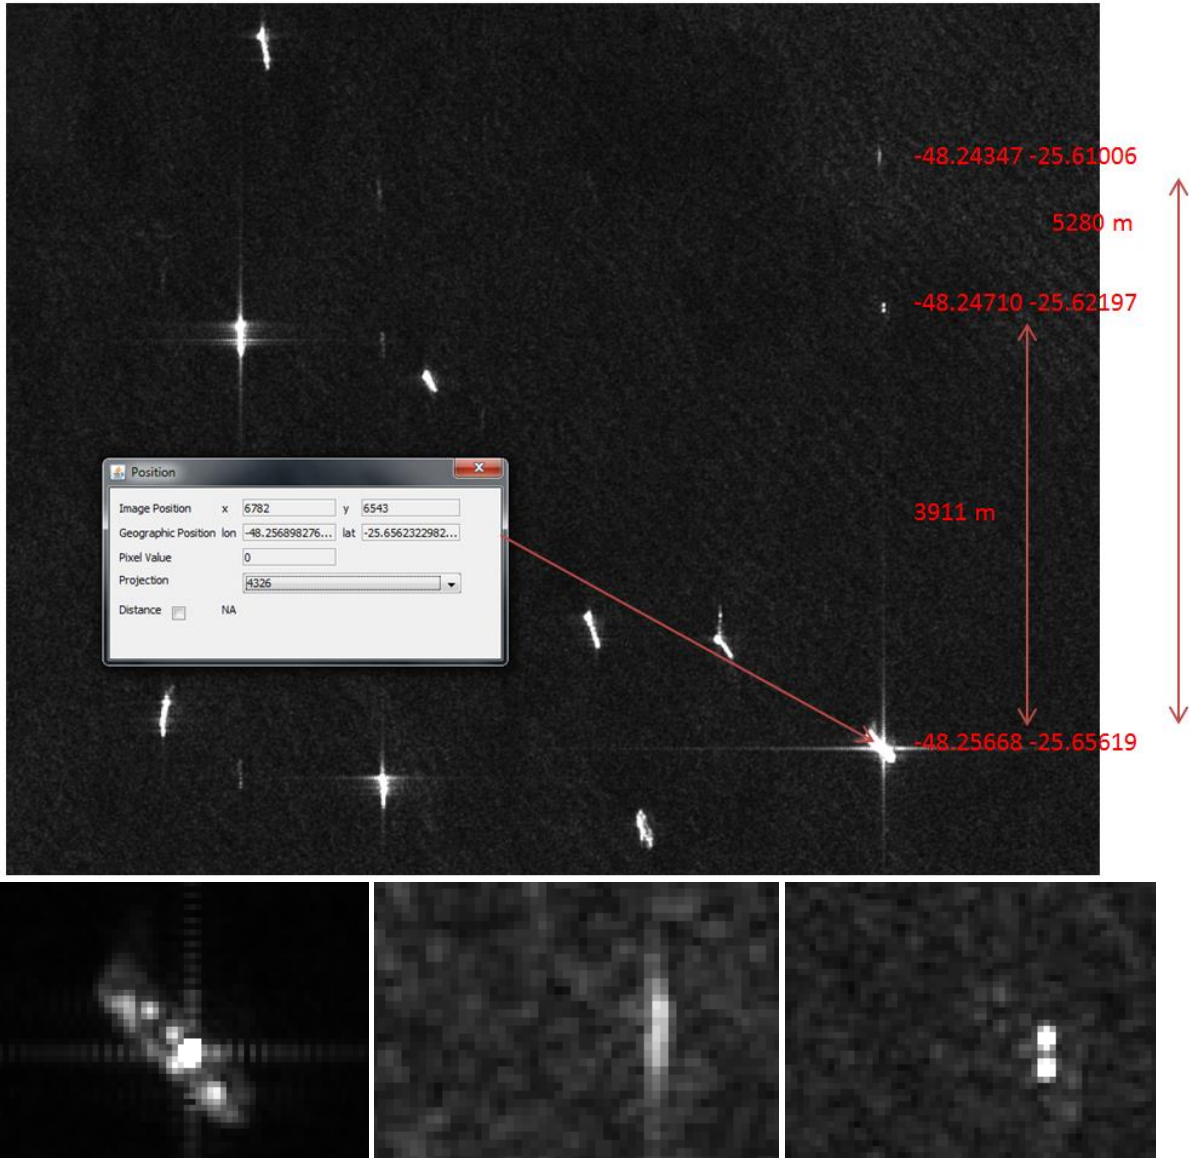
Fig. 5.1. Example of azimuth ambiguities and ghosts. Top: A part of an S-1 image with several strong targets. The main target, indicated by a red arrow, is a bright ship; it is enlarged on bottom left. At 5,280 meters offset in azimuth (vertically), an azimuth ambiguity is visible; it is enlarged at bottom centre. At 3,911 m offset in azimuth, another ghost is visible, enlarged on bottom right. That last ghost is characterised by a double peak.
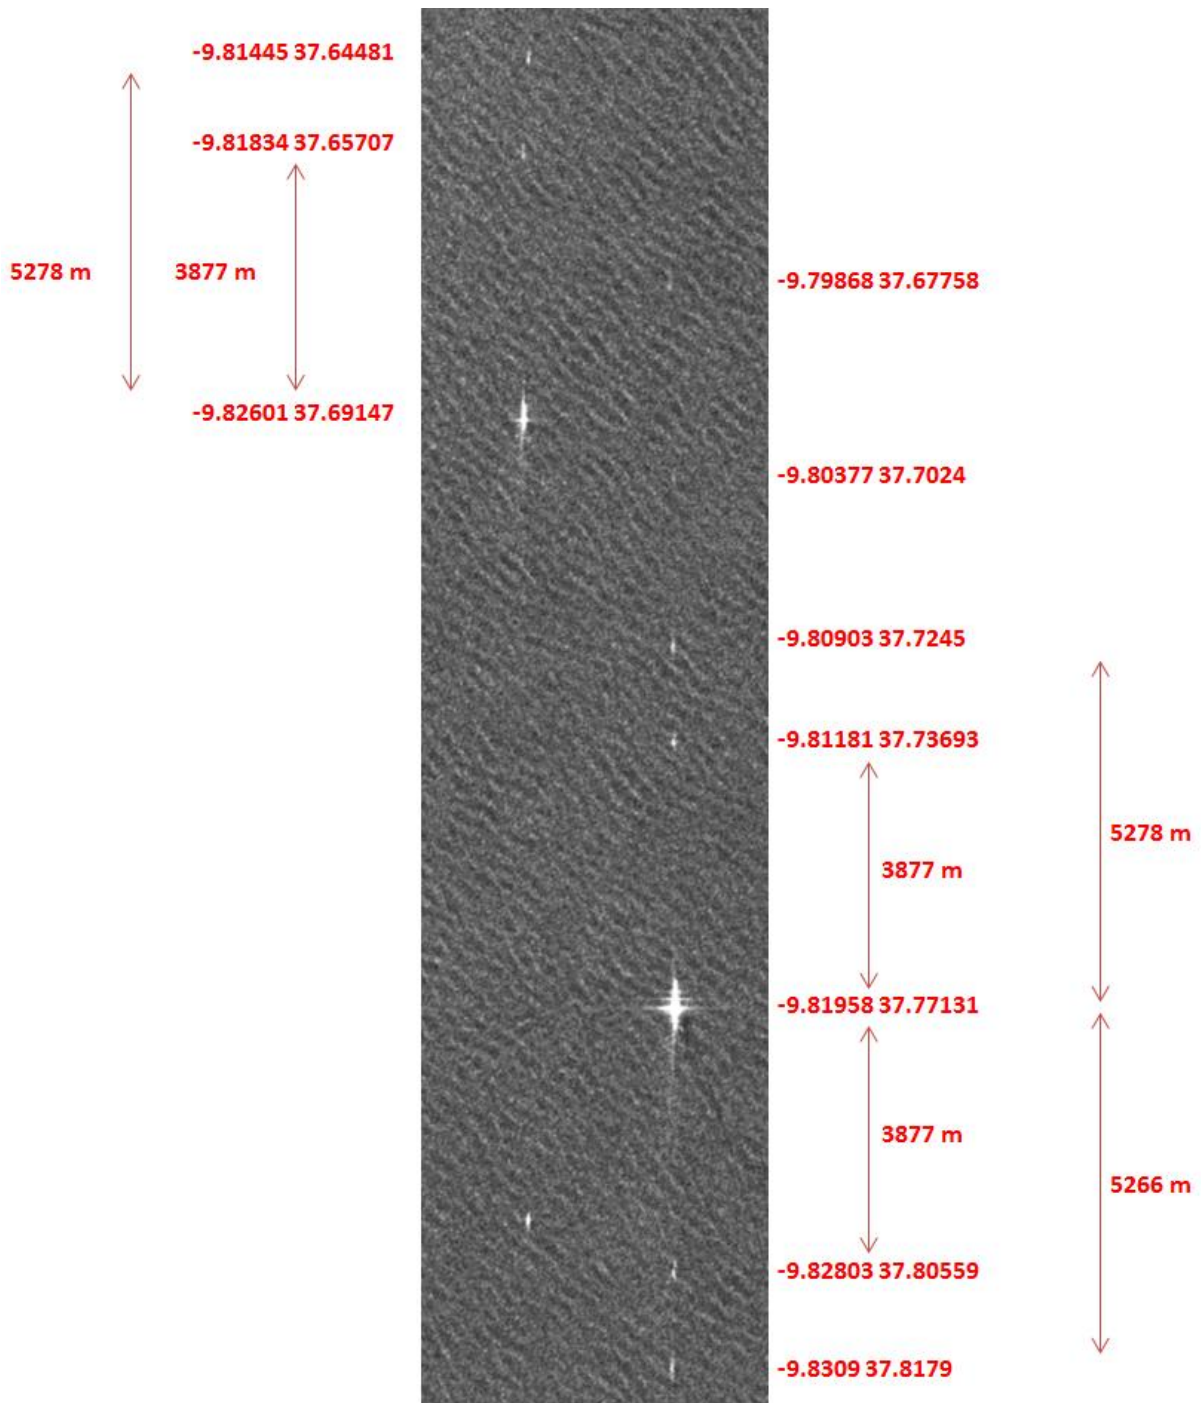*S1A_IW_GRDH_1SDH_20140620T083950_20140620T084014_001131_00119F_2B95*

Fig. 5.2. Another example of azimuth ambiguities (at 5.27 km distance from the bright target) and the other type of azimuth ghost (3.877 km distance).
S1A_IW_GRDH_1SDV_20140902T183442_20140902T183510_002217_00242A_6C68

Inspection of a number of targets in various images shows that azimuth ambiguities occur at theoretically expected distances, and that the other type of azimuth ghost is also often present. See the table below. In the table, all the images are from the pre-operational phase (i.e. before 3 October 2014). No detailed analysis has been done of images in the operational phase, although by simple visual inspection it has been verified that the "other ghosts" still appear in IW and EW images. These ghosts have not been seen in SM images.

Azimuth ambiguity distance, theoretical (m)	Azimuth ambiguity distance, measured (m)	"Other ghost" distance, measured (m)	Beam	Pol	Acquisition date
5634.67	5710	4252	IW3	VV	20140812
5720.16	5678	4183	IW3	VV	20140716
4505.56	4503	NA	IW2	VV	20140812
4684.43	4701	4418	IW2	VV	20140806
~4662	4749	4376	IW2	VV	20140726
5133.34	5280	3911	IW1	HH	20140620
5201.78	5327	3906	IW1	VV	20140716
~5200	5187	3862	IW1	VV	20140720
~5200	5273	3885	IW1	VV	20140726
5238	5278	3877	IW1	VV	20140902

Table 5.1. Azimuth ambiguity and other ghost distances as seen for a number of targets and images.

In conclusion, it can be said that:
- The azimuth ambiguity distance on the image closely matches the expected theoretical value.
- The measured azimuth ambiguity distance generally overestimates the theoretical azimuth ambiguity distance. This is probably because the theoretical distance is only the shift in the azimuth direction, whereas the distance measured on the image is from target to ambiguity, which is normally shifted by a small range component as well.
- The origin of the "other ghost" is unknown. Its distance from the target is shorter than the azimuth ambiguity one, and, similarly to the azimuth ambiguities, it can repeat at constant intervals in the azimuth direction.
- The value of the "other ghost" distance varies (as a function of beam, maybe of incidence angle as well) and follows an unknown formula.
- The "other ghosts" are only visible in the co-pol channel of IW and EW images.

5.2   Azimuth ambiguities and ghosts in repeating images

S-1 has a repeat cycle of 12 days, meaning that a scene on the earth can be re-imaged in exactly the same configuration every 12 days. A collection of acquisitions of the same scene, imaged in the same mode, multiples of 12 days apart can be called a stack, and image pixels in the stack can be made to exactly coincide enabling pixel-based change detection, also coherently (using the phase – see the section on SLC). A detailed comparison of a few stacks that could be found in the S-1 image archive so far shows that the ambiguities and ghosts in azimuth direction that are caused by bright targets on land, are closely repeated over time in the stack. See the figure below for an example.
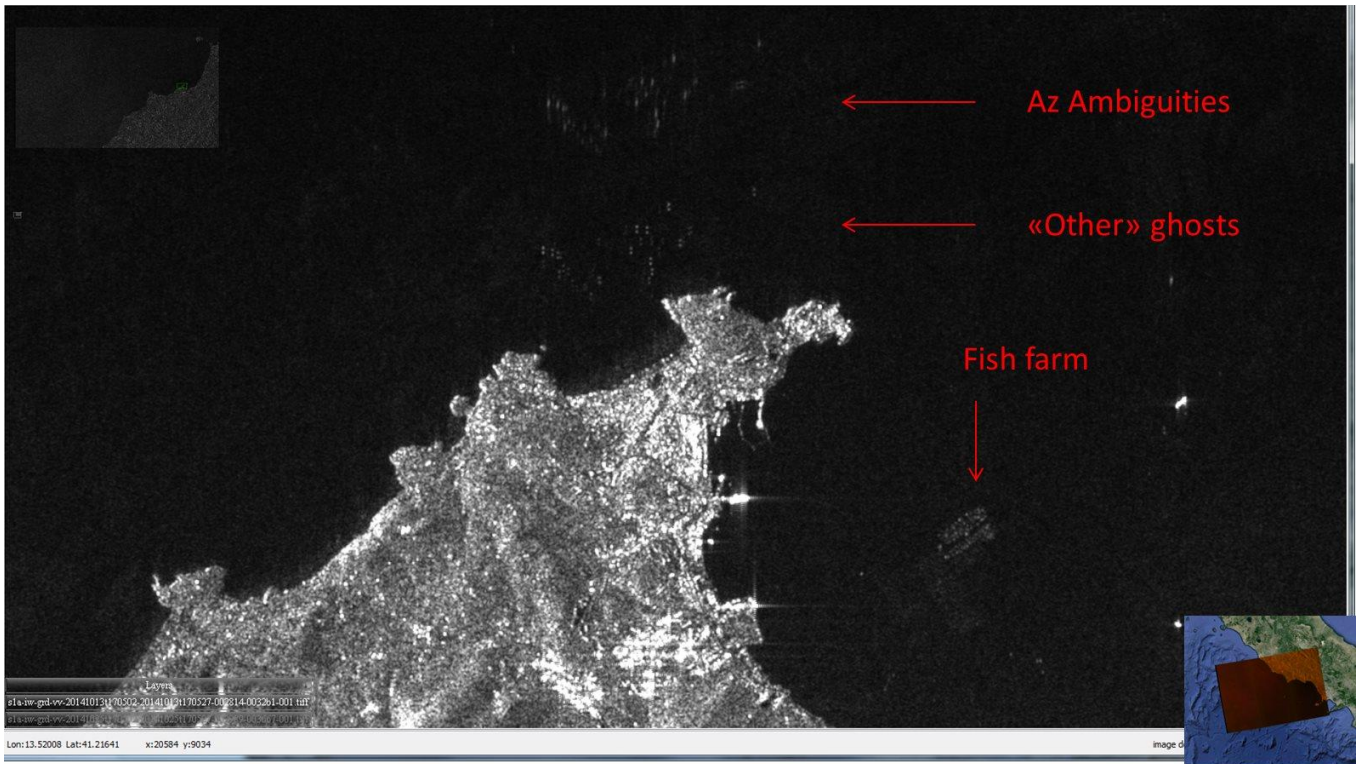
Fig. 5.3. Top: A scene on the Italian coast imaged in IW mode (beam IW3). The upper cluster of dots on the water in the top centre of the image are ambiguities from the bright targets on land at bottom centre (at about 6 km separation). The similar cluster of dots just below the upper one (but still above the coast) are the other type of azimuth ghosts discussed earlier. The object at centre right is a fish farm (as verified from Google Earth).

Bottom: The same scene but imaged 12 days later in exactly the same mode. Most of the ambiguities and ghosts appear in the same places, and at the same intensities.

Top: S1A_IW_GRDH_1SDV_20141013T170502_20141013T170527_002814_0032B1_525A
Bottom: S1A_IW_GRDH_1SDV_20141025T170502_20141025T170527_002989_003667_0B06

Of course, if the structures on land that cause the ghosts change between the 12-day acquisitions, also their ghosts won't be constant. See the example below.
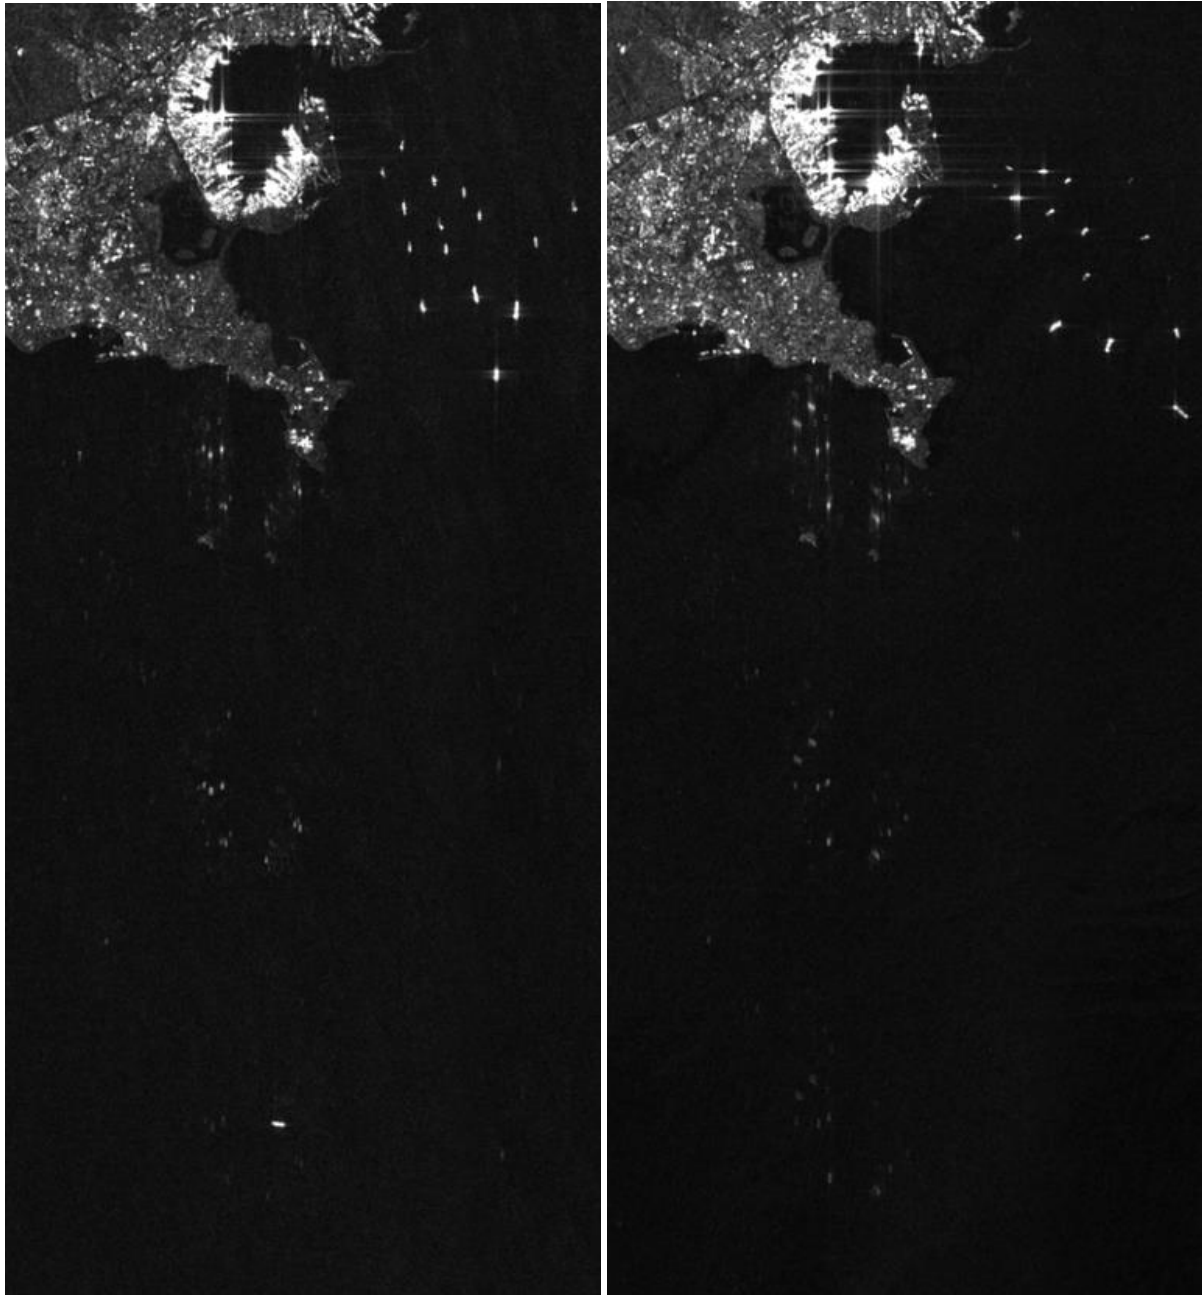

Fig. 5.4. Azimuth ambiguities and ghosts from a container port. In this case, many are not constant from one acquisition to another 12 days later, because the containers and moored ships causing the ambiguities have moved in the mean time.
Left: S1A_IW_GRDH_1SDV_20141015T040550_20141015T040619_002835_003325_239D
Right: S1A_IW_GRDH_1SDV_20141027T040550_20141027T040619_003010_0036E1_DD09

Section 7.3 comes back to the concept of image stacks.

5.3   Range ambiguities

The temporal consistence of range ambiguities in images within the same stack has also been verified, as in the example below. This figure shows an area of the Tyrrhenian Sea (identified with a red dot in the lower left corner) imaged by S-1 with a difference of 24 days (i.e. 2 repeat cycles

apart). In addition to two ships that are traversing the area at the image acquisition times (some azimuth ambiguities of the ships are visible), the two images show common artefacts that are likely range ambiguities originating from strong scatterers in the city of Rome, which is around 145 km towards far range (this distance coincides with the theoretical range ambiguity distance for that area of the image).
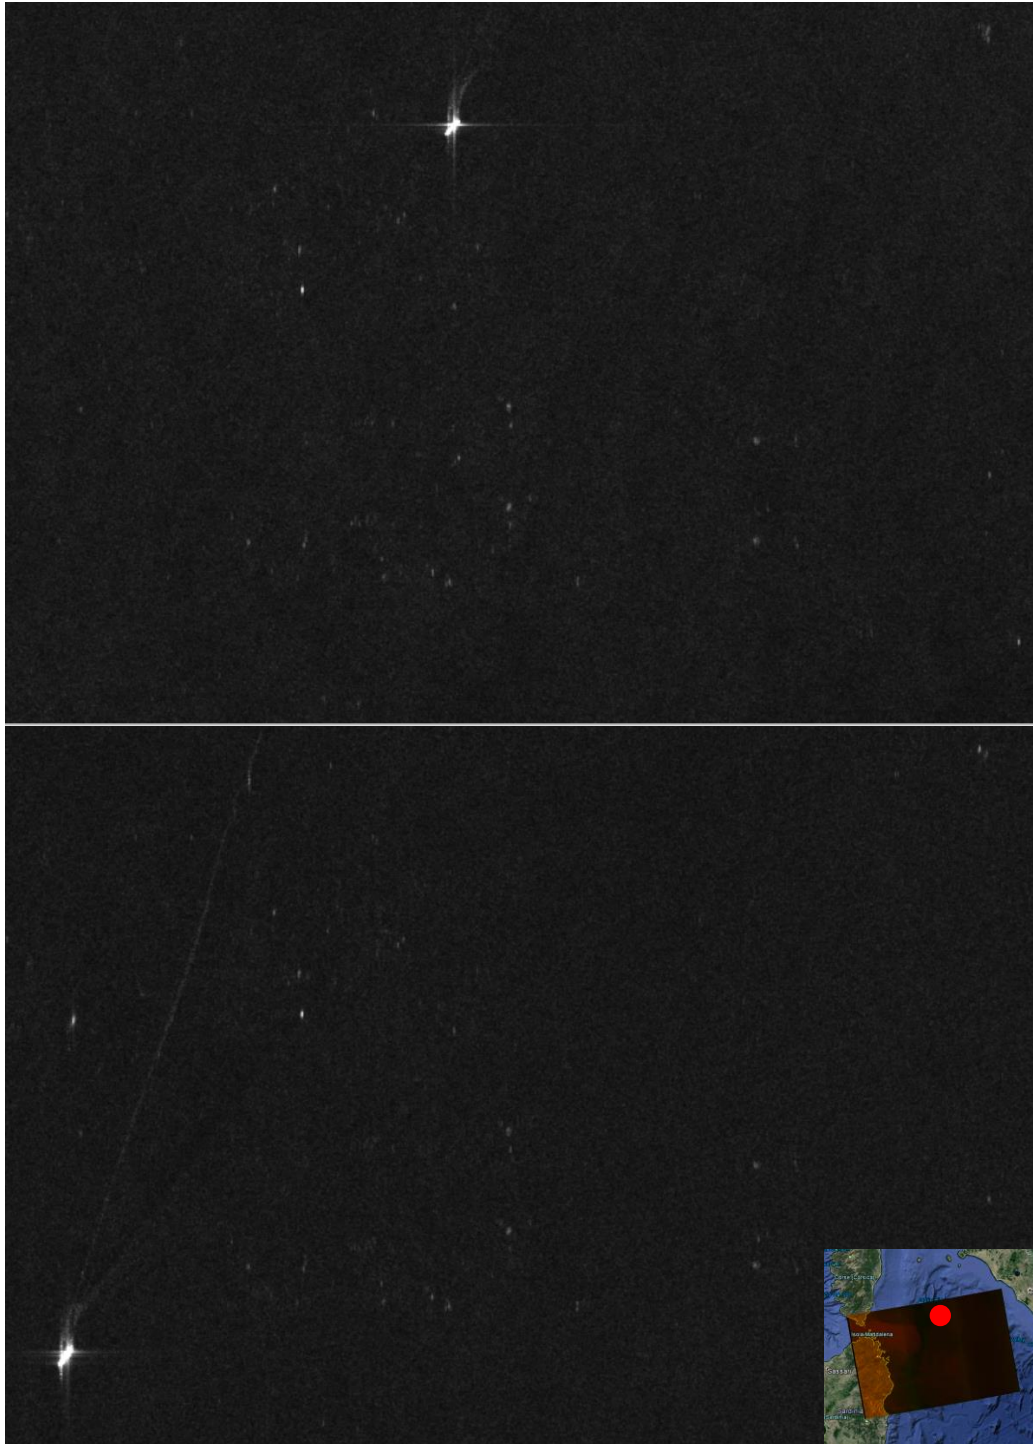

Fig 5.5. Range ambiguities likely originating from the city of Rome, 145 km towards far range. The images also show two different ships traversing the area and some azimuth ambiguities of the ships. S1A_IW_GRDH_1SDV_20141006T171315_20141006T171340_002712_00308D_06C0 S1A_IW_GRDH_1SDV_20141030T171314_20141030T171339_003062_003809_27C3

6 SLC images

The Single-Look Complex (SLC) is a form of the image where the pixels are complex values, i.e., having a real and imaginary part, or equivalently, an amplitude and phase. The more usual GRD (Ground Range Detected) image product retains only the amplitude.

6.1 Possible uses of SLC images

The phase enables, in some cases, the extraction of additional information. The SLC image can be used to construct a GRD image that has different values of resolution and number of looks than the standard GRD product. A GRD image is usually multi-looked, meaning it is the average of a set of sub-apertures. These sub-apertures can be extracted separately from the SLC. Alternatively, the phase can be used for interferometry in combining two images of the same area taken from different viewing angles or at different times, to measure heights or speeds. The phase can also have an application in polarimetry.

In the context of maritime surveillance, four uses of SLC data could potentially be considered. First, for ship detection, the use of sub-aperture cross-correlation has been proposed as an alternative way to doing detection on GRD imagery. However, its superiority over the normal use of the GRD product is not clearly confirmed.

Secondly, Sentinel-1 uses the TOPS mode of scanning to produce its multi-beam (wide swath, IW or EW) images. This entails the use of bursts, which are strips of imaged terrain sequenced in azimuth. The beginning of a burst overlaps with the end of the previous burst. Therefore, a target that is in the overlap zone is seen twice with a separation of the order of 2 seconds. If the target is moving, its movement could possibly be detected and its speed and course measured. The overlap zone is relatively narrow, so only a fraction of the targets in the image could possibly make use of this effect.

Thirdly, ship detection needs a land mask to prevent false detections on land. For that, an external coastline map is used. However, these maps have limited accuracy, and actual detection of the coastline would be preferable. This is difficult in GRD images (which can only use image intensity that can be similar for sea and land depending on polarisation, incidence angle and wind speed), but with SLC images, the complex coherency can be used, which is expected to be quite different for land and sea.

In the fourth place, the phase difference between the co-pol and the cross-pol channels of dual-polarisation images might possibly help to distinguish real ships from false alarms or give information about the target. This is however something that needs additional research. The following three figures of a ship seen in a dual-polarisation IW SLC image give an example, in the form of the complex coherence between the two polarisation channels. With A_i the complex pixels of the first channel and B_i those of the second, writing . as product, * as complex conjugate, < > as convolution with a boxcar function i.e. spatial averaging and | | as absolute value, the complex coherence is defined in the usual way as

$$cc = <A_i \cdot B_i^*> / sqrt(<A_i \cdot A_i^*><B_i \cdot B_i^*>)$$

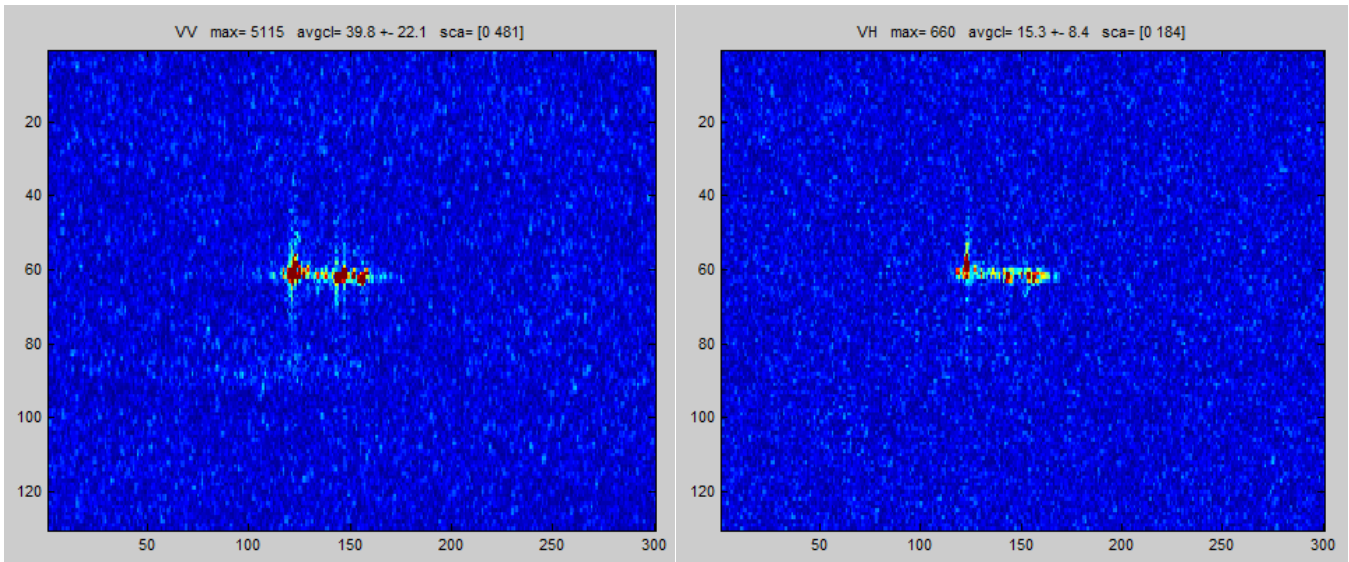(note that $<A_i \cdot A_i^*> = <|A_i|^2>$ is real, and likewise for B).

Fig. 6.1. Small area around a ship in an IW SLC image. Left VV, right VH polarisation. SLC pixel amplitude is plotted as false colour from blue to red.
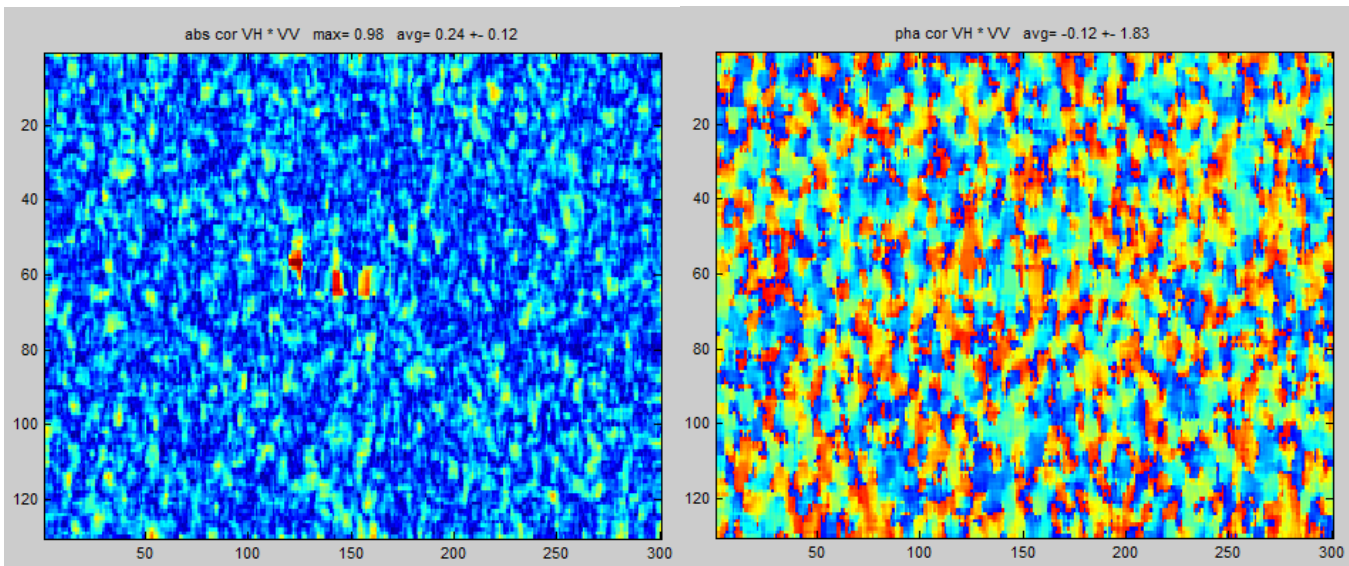*S1A_IW_SLC__1SDV_20141202T040708_20141202T040735_003535_004293_C60C*

Fig. 6.2. The complex coherence between the VV and the VH channel from the previous picture, computed over a 5 x 5 pixel extent. Left shows the amplitude of the complex coherence, scaled in false colour between 0 and 1. Right shows the phase, scaled between –pi and +pi.
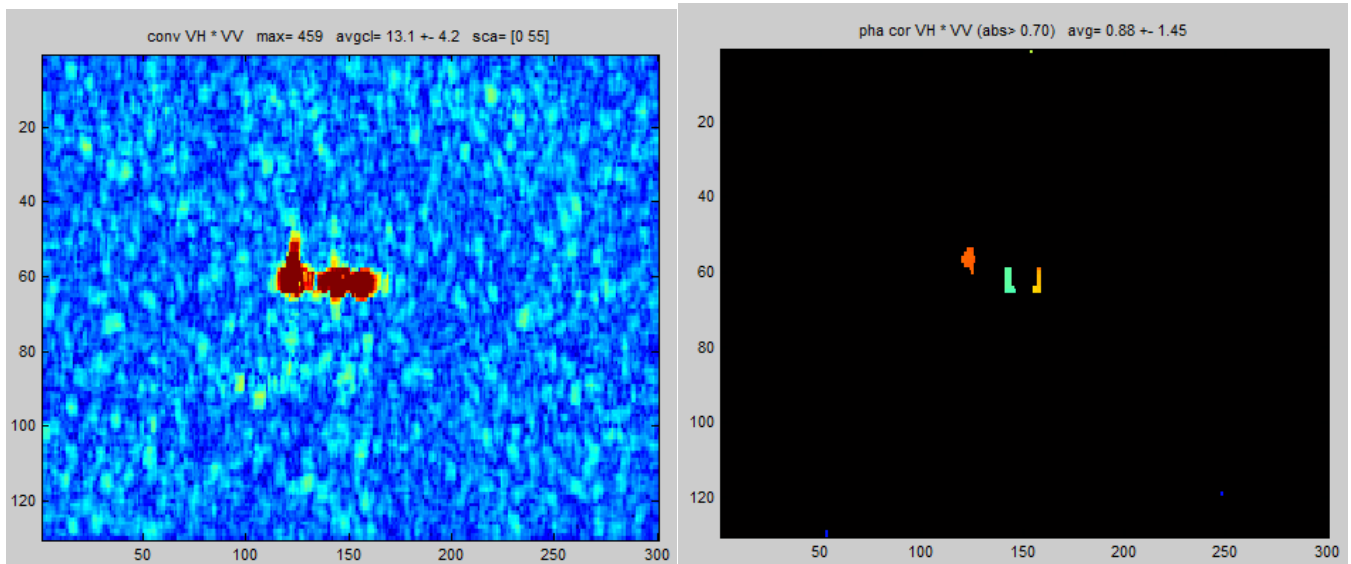*Same image as previous figure.*

Fig. 6.3. Left: The square root of the absolute value of the convolved product of the complex VV and VH channels, i.e. sqrt(|<A_i . B_i*>|) *from the text above. Right: The phase of the complex coherence from Fig. 6.2 right, but only plotted where the absolute value of the complex coherence is larger than 0.7.*
Same image as previous figures.

The figures show that the polarimetric phase difference of the strong scatterers on the ship is well defined. Fig. 6.3 left shows where both polarisation channels are strong, so their phase is not dominated by noise, and Fig. 6.3 right shows the polarimetric phase difference in those places.

6.2   Burst storage

For an SLC product, all the bursts are stored consecutively in one single tiff image, separated by bands of zeros. Therefore, one SLC image is not really an 'image' because the empty bands and overlap zones in there do not correspond (uniquely) to geographical locations on the ground. The bursts have to be extracted based on their start and end locations that are specified in the metadata xml files provided with the product. The relevant metadata to handle bursts is in the <swathTiming> and <geolocationGrid> sections of the metadata files. With these fields it is possible to correctly geo-locate and timestamp each line of the tiff files. Relevant fields include:

Field	Where	Explanation
azimuthTimeInterval	<imageAnnotation> <imageInformation> <azimuthTimeInterval>	Time spacing between lines of the output image [s]
linesPerBurst	<swathTiming> <linesPerBurst>	Number of range lines within each burst (constant for all the bursts within the swath)
burstFirstLineAzimuthTime[]	<swathTiming> <burstList count="nn"> <burst> <azimuthTime>	Zero Doppler azimuth time of the first line of this burst [UTC]
nBursts	<swathTiming> <burstList count="nn">	Number of bursts records within the list
firstValidSample	<swathTiming> <burstList count="nn"> <burst> <firstValidSample>	An array of integers indicating the offset of the first valid image sample within each range line. This array contains count attribute integers, equal to the linesPerBurst field (i.e. one value per range line within the burst), separated by spaces. If a range line does not contain any valid image samples, the integer is set to -1.

Table 6.1. Relevant metadata to extract bursts.

The following diagram illustrates the layout of one of the beams in IW and EW SLC images.
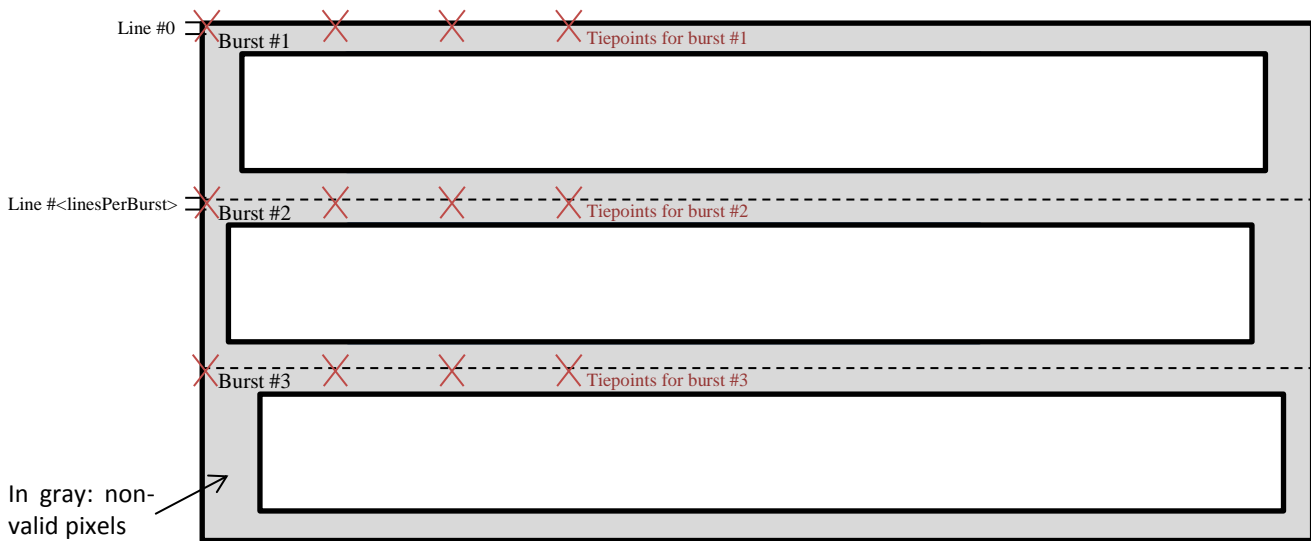

Fig. 6.4. Layout of IW and EW SLC data (tiff files). Valid pixels are in white, non-valid pixels in grey, the dashed lines indicate the separation between bursts, and the crosses indicate the tie points provided in the metadata.

6.3 Geo-coding in SLC images

As explained in Section 3, two different methods to geo-locate the Sentinel-1 images have been explored. The first way uses tie points, the second uses orbit data. The process to geo-locate multi-beam SLC images using either of these methods is slightly different to the process to geo-locate GRD images, and the relevant burst timing information (Table 6.1) needs to be taken into account.

For the tie points method , according to the documentation [9, Table 6-85] a new set of tie points is

provided every 1 second of azimuth time. However, this does not appear to be true in the sample IW SLC images: a set of tie points is provided for the first line of each burst, i.e., a new set of tie points every ~2.75 s.

The information in a set of tie points (<line>, <latitude>, <longitude>, …) corresponds to the pixels in the first line of the burst. This means that without de-bursting, these tie points cannot be used to geo-locate the pixels of the previous burst by simple interpolation. In order to interpolate correctly, the coordinates (<line>) of the tie points need first to be de-bursted. To perform this de-bursting, the metadata fields in Table 6.1 are needed. This is graphically seen in the next figure.
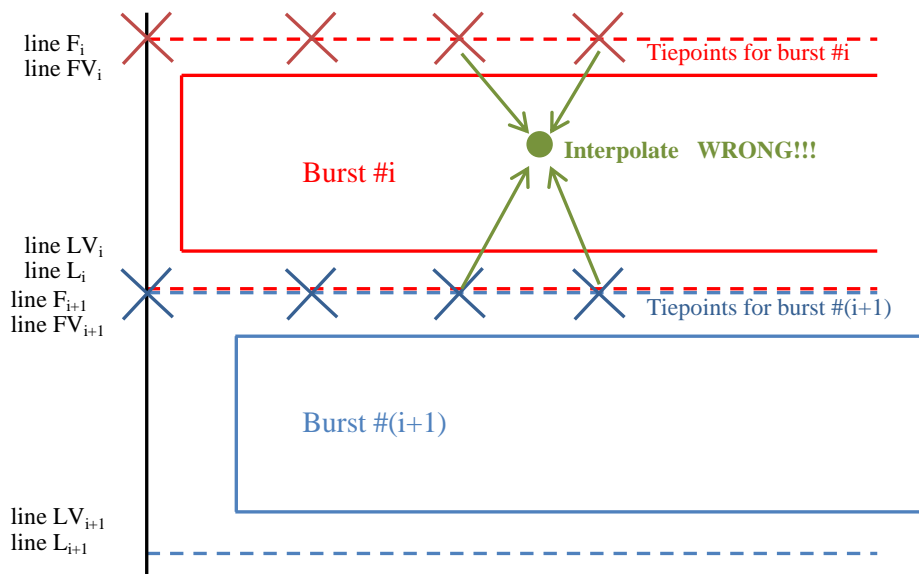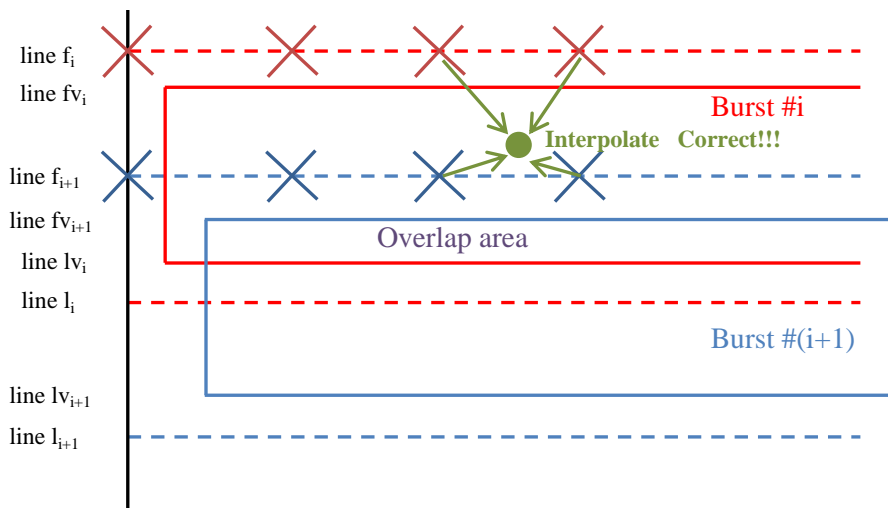
Fig. 6.5. Graphical representation of the geocoding process using tie points. Before de-bursting (top half of the figure) the interpolation of the tie points coordinates would give wrong results. Images need to be de-bursted (bottom half) in order to use simple interpolation correctly.

In the geo-locating method that uses orbit data, the metadata fields in Table 6.1 are also needed, in this case to accurately determine the zero Doppler time of each image line.

7 Scoping for bulk processing

Sentinel-1 routinely acquires images over maritime areas. As mentioned in the introduction, that is over the European seas, the North Atlantic, the Arctic and the world's coasts. Given that these data are there, it should be considered to analyse them for maritime surveillance. Routine ship detection on these data would build up a history of ship traffic patterns, which can be useful for policy decisions or economic applications. At the moment, efforts are ongoing to implement a batch mode for SUMO so that it can automatically analyse all S-1 images as they come in. Manual verification is not possible because it would take too much time. Therefore, the detection threshold would need to be set higher than usual, to avoid too many false alarms on weak targets, and the resulting detections would represent the medium and bigger ships, not the smaller ones. It is expected that the AIS- and LRIT-carrying ships, the ones of 300 GT which is roughly comparable to 45 m length and up, should mostly be detected in this way. Still, the SUMO detection algorithm should be improved further for better automatic performance. In particular over Arctic areas, many false alarms may be expected due to sea ice; at present, SUMO has no facility to deal with that.

7.1 Data volume

Today, the only way to access S-1 imagery is through ESA's rolling archive, that holds 2 months of the most recent data. No clear mechanism is proposed to easily access S-1 data older than that. The need to access such data is expected, however, as new information demands arise (e.g., questions like, how did the ship activity develop in the Gulf of Guinea during the last 3 years?) or as improved processing algorithms become available. In such cases, even if all S-1 data have been analysed as they came in, one needs to re-analyse the data archives to perform manual analysis or to apply the newly available algorithms. Therefore, a need is felt to store an easily accessible archive of S-1 maritime images.

An estimate for the storage space needed for such an archive is as follows.

S-1 parameters:
* Duty cycle: 25 min / orbit [10]
* Repeat period: 12 days = 175 orbits [11]
* Image modes: EW, IW, SM [1]
* Polarisations: single, dual [1]
* Products: SLC, GRDF (SM only), GRDH, GRDM [1]

Main products for bulk maritime surveillance:
* EW-GRDH, IW-GRDH; dual pol and single pol [10, and own choice]

Measured from downloaded S-1 images:
* EW-GRDH dual pol: 1 min acquisition gives 1.0 GB data set
* IW-GRDH dual pol: 25 sec acquisition gives 1.6 GB data set
* single pol is half the quantity

Observation schedule: [10]
* EW-GRDH dual pol, over open sea incl. arctic seas: 5 % of the time
* IW-GRDH dual pol, over open sea incl. arctic seas: 5 % of the time
* IW-GRDH dual pol, over coasts: 20 % of the time
See also Fig. 7.1 below.

Derived:
- Acquisition time per day (max) = 175 orbits / 12 days * 25 min / orbit = 365 min / day
- EW-GRDH dual pol = 1.0 GB / min
- IW-GRDH dual pol = (60 / 25) * 1.6 = 3.8 GB / min
- Open sea EW-GRDH dual pol: 5 % * 365 min / day * 1.0 GB / min = 18.3 GB / day
- Open sea IW-GRDH dual pol: 5 % * 365 min / day * 3.8 GB / min = 69.4 GB / day
- Coastal IW-GRDH dual pol: 20 % * 365 min / day * 3.8 GB / min = 277.4 GB / day

Conclusion:
- Total: 18.3 + 69.4 + 277.4 GB / day = 365.1 GB / day = 133 TB / year
- Collecting all open sea images, but being more restrictive with coastal images, could half the total to: 18.3 + 69.4 + 33 % * 277.4 GB / day = 180.2 GB / day = 66 TB / year

During the first time of its operations (that started 3 Oct 2014), S-1 will not be at full capacity. Also the full 25 % duty cycle will probably not be reached in practice. Also some scenes may be acquired in single pol, leading to half the volume per scene. These factors lead to lower volume.
On the other hand, the calculation excluded SLC products which can be up to 8x as big. However, in the context of ship detection, these are more of experimental nature, so would not be used for bulk processing. Nor are they much produced over the sea by ESA judging from the rolling archive so far.

The uncertainties in this calculation of what S-1 will acquire stem from:
- The estimates of 5 %, 5 %, 20 % above;
- Which products will actually be acquired (EW/IW, single/dual pol);
- Actual use of the duty cycle.
These could lead to maybe 50 % deviations in the total.

From the total of what S-1 will acquire, it is JRC's choice of what fraction of the maritime data actually to collect (download) and process.

Therefore, an estimated need of 60 TB storage space for the coming year could be realistic.
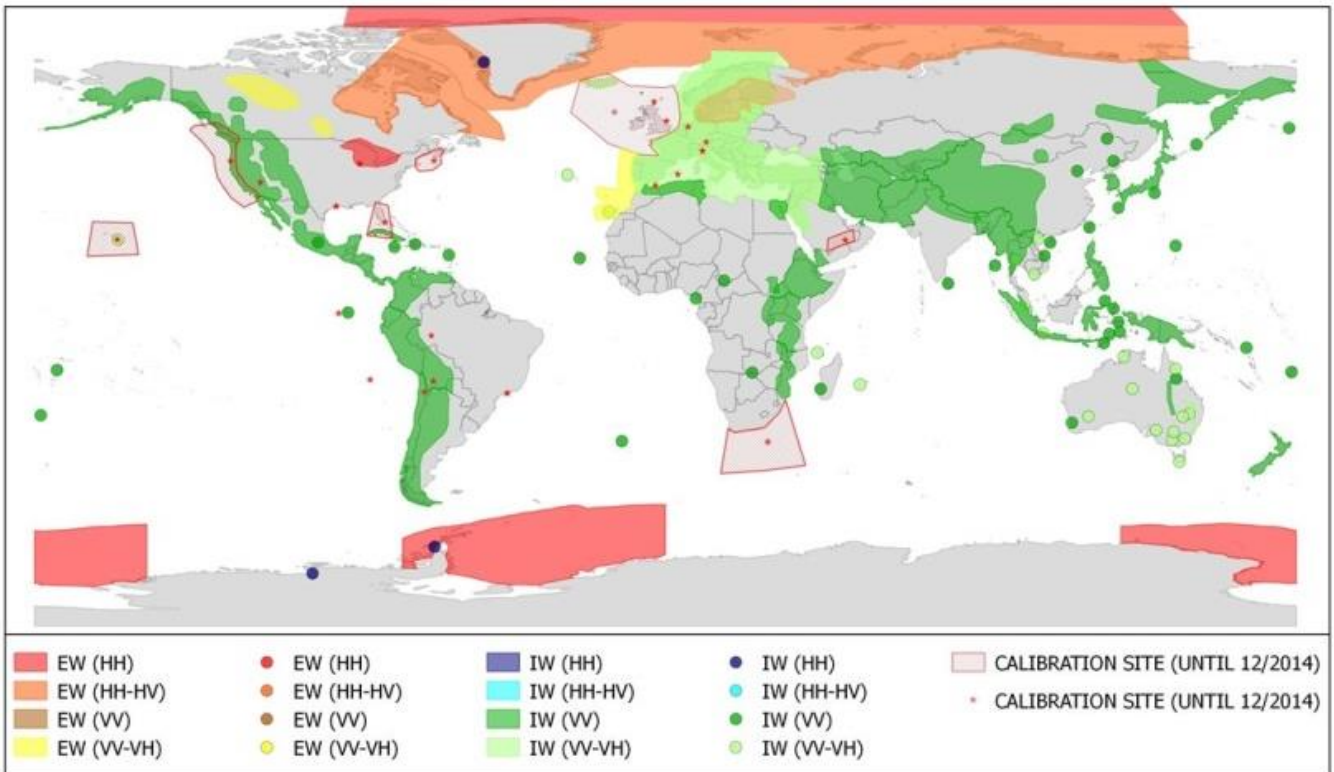
Fig. 7.1. Current S-1 acquisition plan, still during the ramp-up to full operations [10].

7.2   Processing time

The time to analyse a Sentinel-1 image with the SUMO ship detector was measured. Having the image on the local PC's disk was compared with having the image on a network shared disk (which could be a way that a large image archive is accessible). The image used is a typical one: GRDH product, IW mode, dual pol (1.86 GB); it is the one treated in the chapter on Ship detection. The PC that was used is Windows64, 2 core, 2.83 GHz, 8 GB RAM.

The results are shown in the table below.

SUMO action	What it does	Local drive	Network drive	Factor
Loading image first time	Reading xml header (3 MB), reading first corner of image (some MB) for display	12 s	45 s	4
Getting entire overview image in first polarisation	Reading 953 MB tiff file	40 s	13 m 20 s	20
… in second polarisation	Reading 953 MB tiff file	50 s	13 m 30 s	20
Ship detection, - first polarisation	Analysing 953 MB tiff file, and going back to the tiff file to analyse the local surroundings of each detected target in detail	1 m	approx. 20 m	20
- second polarisation	idem	1 m 10 s	approx. 25 min	21
- merging two polarisations	Analysing the local surroundings of each detected target in detail. There were 613 targets detected	7 m 0 s	not finished after 1 hour, stopped	> 9

Table 7.1. Timing the automatic ship detection analysis of an S-1 image with SUMO.

The conclusion is that accessing the Sentinel-1 data with SUMO on the network drive is 20x slower than on the local drive. The SUMO processing time is dominated by I/O, not by actual computations. So fast access to the data is essential.

Batch processing of such an image would take 12 s + 1 m + 1 m 10 s + 7 m = 9.4 min (the visualisation steps are not needed) on the local drive. In the previous section, it was estimated that 180 GB/day of Sentinel-1 data might need to be analysed. So that would mean 180 / 1.86 * 9.4 = 15.2 hours / day. That is a feasible amount, considering this is on just one normal PC. With that speed, it is possible to keep up with the incoming data. On the other hand, if one would want to do the re-analysis of one year of Sentinel-1 data, faster speeds would be needed.

The times above exclude manual verification, which – if done – would add maybe of the order of 5-10 seconds per target for an experienced operator working on a fast platform. So for the present example image, that had 614 targets, it would mean another 1–1.5 hour. Most images would have far fewer targets; but still, routine manual analysis of all maritime images would be too much work.

7.3 Stacks of images

The figure below presents the S-1 images of the Mediterranean Sea that have been downloaded from the ESA rolling archive from the start of the operational phase (3 October 2014) until mid-November. There are 131 images, and they only represent some of the images acquired over the Med and available on ESA's archive. The overall size on disk of these images is 207 GB. The green frames indicate stacks with more than one image from 12-day repeat cycles in exactly the same mode and same starting time in the day. One example of a 4-image stack located east of Crete is:

S1A_IW_GRDH_1SDV_20141003T040734_20141003T040754_002660_002F64_288E.SAFE
S1A_IW_GRDH_1SDV_20141015T040734_20141015T040754_002835_003325_D502.SAFE
S1A_IW_GRDH_1SDV_20141027T040734_20141027T040754_003010_0036E1_830B.SAFE
S1A_IW_GRDH_1SDV_20141108T040734_20141108T040754_003185_003AB0_2C2A.SAFE

Such stacks can be used to identify coastlines and fixed false alarm causes on the sea such as ambiguities caused by scatterers on land (ref. Section 5.2).
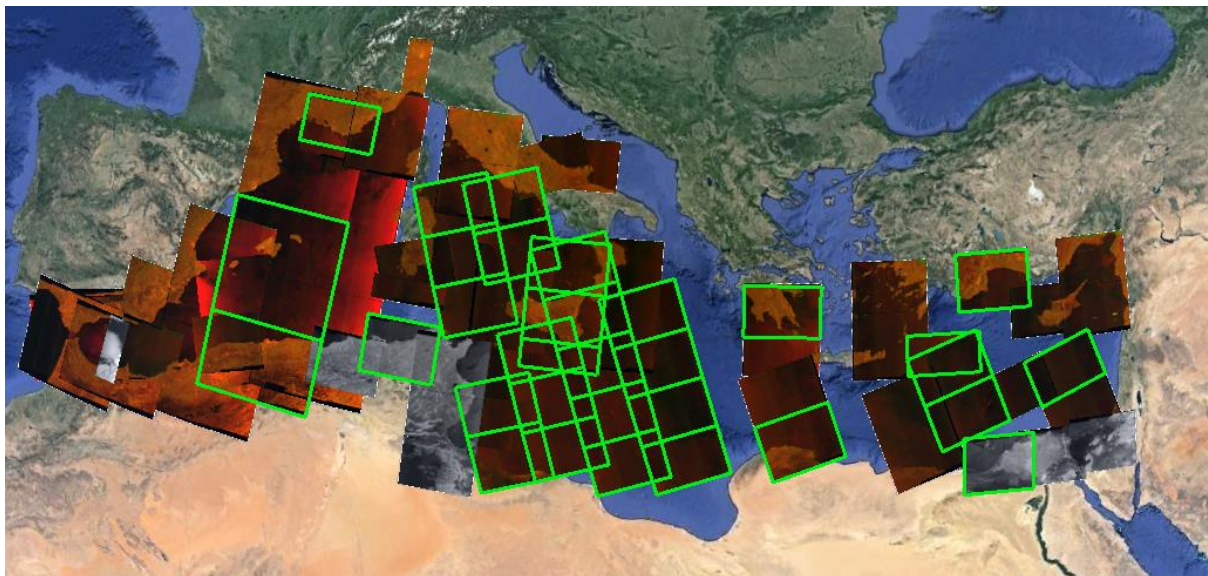

Fig. 7.2. S-1 images of the Med downloaded from 3 October 2014 until mid-November. Most images are in IW mode, some are EW (mainly in the western Med) and two are SM (near Genoa and near Gibraltar). Reddish images are dual pol (HH+HV or VH+VV), grey images are single pol (HH or VV). The green frames indicate that there is more than one image in a stack. Google Earth background.

8 Conclusions

The main conclusions are included in the Executive Summary. Sentinel-1 image quality as assessed here lives up to its expectations, and the system is very well suited for maritime surveillance. Although a detailed quantitative assessment of ship detection performance is still to be done, the findings so far are positive and in line with predictions.

One unexpected instrumental effect was found in the form of ghost repetitions of bright targets, positioned with an azimuth offset. This is in addition to the normal azimuth ambiguities, and the cause is not known. These ghosts can lead to false alarms in the ship detection. Although not a major effect, hopefully further tuning of the instrument or the processing can alleviate this phenomenon. Otherwise its precise behaviour (distance from the source target as a function of image parameters) must be further analysed so that it can be corrected for in the detection process.

The best possible performance for bulk ship detection may not be reached in practice because the routine products (GRDH) made from the acquired raw data are not the optimum ones (SLC) for ship detection, but that is also a matter of weighing a large increase in processing time against a small increase in performance. Furthermore, the fixed choice of using the wide modes of IW over coastal waters and EW further from the coast means that very little imagery will be acquired over seas in SM mode – which would have shown smaller ships at the expense of losing significant coverage. Similarly, the fixed geographic acquisition schedule means that some areas will be frequently monitored at the expense of never looking at many others.

The routine 12-day repeat imaging of the same areas opens up the possibility to make a catalogue of constant image features caused by structures on land (bright reflectors, fixed infrastructures) that traditionally lead to high false alarms in the littoral zone. This could be very helpful to obtain better results especially in bulk automatic processing for ship detection.

9   References

[1]   SENTINEL-1 SAR User Guide, ESA. https://sentinel.esa.int/web/sentinel/user-guides/sentinel-1-sar

[2]   SENTINEL-1 Technical Guide, ESA. https://sentinel.esa.int/web/sentinel/sentinel-1-sar-wiki

[3]   GSHHG. A Global Self-consistent, Hierarchical, High-resolution Geography Database. http://www.soest.hawaii.edu/pwessel/gshhg/

[4]   "Geolocation of RADARSAT-2 Georeferenced Products", RADARSAT-2 Applications Technical Note, RN-TN-53-0076, Issue 1/2, January 14, 2014

[5]   D. Small et al., "Preliminary Validation of ASAR Geometric Accuracy", Proc. of ENVISAT Validation Workshop, Frascati, Italy, 9-13 December 2002 (ESA SP-531, August 2003)

[6]   P.D. Beaulne and I. Sikaneta, "A Simple and Precise Approach to Position and Velocity Estimation of Low Earth Orbit Satellites", Defence R&D Canada, Ottawa Technical Memorandum, DRDC Ottawa TM 2005-250, December 2005

[7]   P.W. Vachon and J. Wolfe, "GMES Sentinel-1 Analysis of Marine Applications Potential (AMAP)", Defence R&D Canada – Ottawa, ECR 2008-218, October 2008

[8]   Paris W. Vachon, John Wolfe and Harm Greidanus, "Analysis of Sentinel-1 Marine Applications Potential", IEEE International Geoscience and Remote Sensing Symposium, Munich, Germany, 22-27 July 2012

[9]   S1-RS-MDA-52-7441 Sentinel-1 Product Specification

[10]   SENTINEL-1 Observation Scenario, ESA. https://sentinel.esa.int/web/sentinel/missions/sentinel-1/observation-scenario

[11]   SENTINEL-1, Satellite description, Orbit, ESA. https://sentinel.esa.int/web/sentinel/missions/sentinel-1/satellite-description/orbit

[12]   Michele Vespe and Harm Greidanus, "SAR Image Quality Assessment and Indicators for Vessel and Oil Spill Detection", IEEE Transactions on Geoscience and Remote Sensing, Vol. 50, No. 11, Nov 2012, pp. 4726-4734

Acknowledgments

Europe Direct is a service to help you find answers to your questions about the European Union
Freephone number (*): 00 800 6 7 8 9 10 11
(*) Certain mobile telephone operators do not allow access to 00 800 numbers or these calls may be billed.

A great deal of additional information on the European Union is available on the Internet.
It can be accessed through the Europa server http://europa.eu.

How to obtain EU publications

Our publications are available from EU Bookshop (http://bookshop.europa.eu),
where you can place an order with the sales agent of your choice.

The Publications Office has a worldwide network of sales agents.
You can obtain their contact details by sending a fax to (352) 29 29-42758.

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