

CP40

SARin for Coastal Ocean

WP4000 ATBD contribution

isardSAT Reference: ISARD_ESA_CPO_SCO_ATBD_225
Issue: 1.a

Prepared by: isardSAT CP40 team

18 December 2013
Activity: CryoSat+ for Ocean

This page has been intentionally left blank

Change Record

Date	Issue	Section	Page	Comment
18 December 2013	1.a	all	all	Initial Issue

Control Document

Process	Name	Date
Written by:	Pablo García	18 December 2013
Checked by:	Cristina Martín-Puig	18 December 2013
Approved by:	Mònica Roca	18 December 2013

Distribution List

Company	Name
ESA	Jerome Benveniste
SATOC	David Cotton
isardSAT	Pablo García Mònica Roca Cristina Martín-Puig

Table of Contents

1	INTRODUCTION	5
1.1	SCOPE	5
1.2	ACRONYMS	5
1.3	REFERENCES	5
2	OVERVIEW	6
3	WORK DONE	9
3.1	INPUT DATA SET	9
3.2	AREA OF INTEREST	10
3.3	ALGORITHMS	12
3.4	RESULTS	15
4	CONCLUSIONS	23
5	WORK TO BE DONE	24

1 Introduction

1.1 Scope

The scope of this document is to describe the investigation carried out by the isardSAT team within the CP4O project. It is about a small scientific study of the capability of CryoSat-2 Synthetic Aperture Radar interferometry (SARin) mode to mitigate the echoes contamination on coastal ocean areas.

1.2 Acronyms

AoA	Angle of Arrival
CS2	CryoSat-2
CP4O	CryoSat Plus For Ocean
DS	Data Set
L2I	Level 2 Intermediate/Interim (ESA product)
SARin	Synthetic Aperture Radar interferometry
SSH	Sea Surface Height
TBD	To be defined

1.3 References

- AD-1 CryoSat +: Ocean Theme. CP4O-Cryosat Plus 4 Oceans. Technical Proposal. November 2011. SATOC, DTU Space, isardSAT, NOC, Noveltis, STARLAB, TUDelft, University of Porto and CLS.

2 Overview

The investigation described in this document is based in a uniqueness of CryoSat-2 (CS2): its capability to operate in Synthetic Aperture Radar interferometry (SARin) mode; possible thanks to carrying onboard a fully duplicated altimeter system for the processing of the echoes received by two antennas simultaneously.

This mode allows for deriving the angle of arrival (AoA) of the echoes in the same direction of the antennas baseline (across-track direction as shown in Figure 1). This capability adds to the SAR enhanced along-track spatial resolution (see central footprint on Figure 1), the ability to determine the origin of the echo across-track, specifically the angle of inclination from nadir (see right hand side footprint on Figure 1).

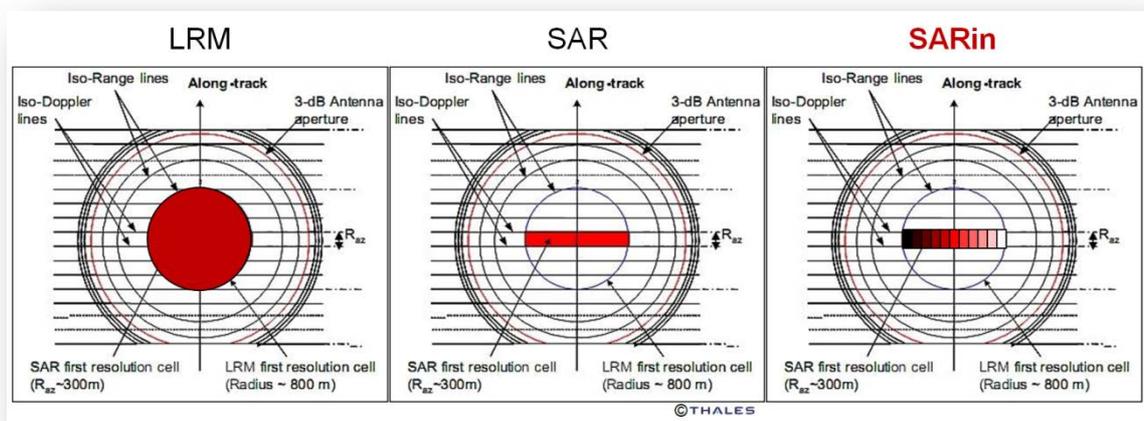


Figure 1. The three CS2 possible modes: LRM, SAR & SARin (courtesy of Thales Alenia Space).

The AoA is derived from the phase information of the signals received by the two antennas. The processing of the phase difference between the two antennas, knowing the characteristics of the pulse emitted and the dimensions of the onboard antennas baseline, gives us the waveform figure in terms of AoA. For equal phase between echoes (zero difference between them), the backscatters are classified as nadir. Otherwise, it shall be proved that the echo comes from one side or the other of the sub-satellite track in the across-track direction.

In Figure 2 we can observe, in a case of a coastal ocean scenario, the well known power waveform on top of the figure, and the less known AoA waveform (intrinsic to the SARin mode) on the bottom. We can see how the two clear peaks of this contaminated power waveform are in line with two correspondent coherent (less noisy) AoA zones: the peak on the left, very specular, shows a AoA far from zero, the peak on the right, ocean like, is in line with the red line reference (AoA = 0) crossing.

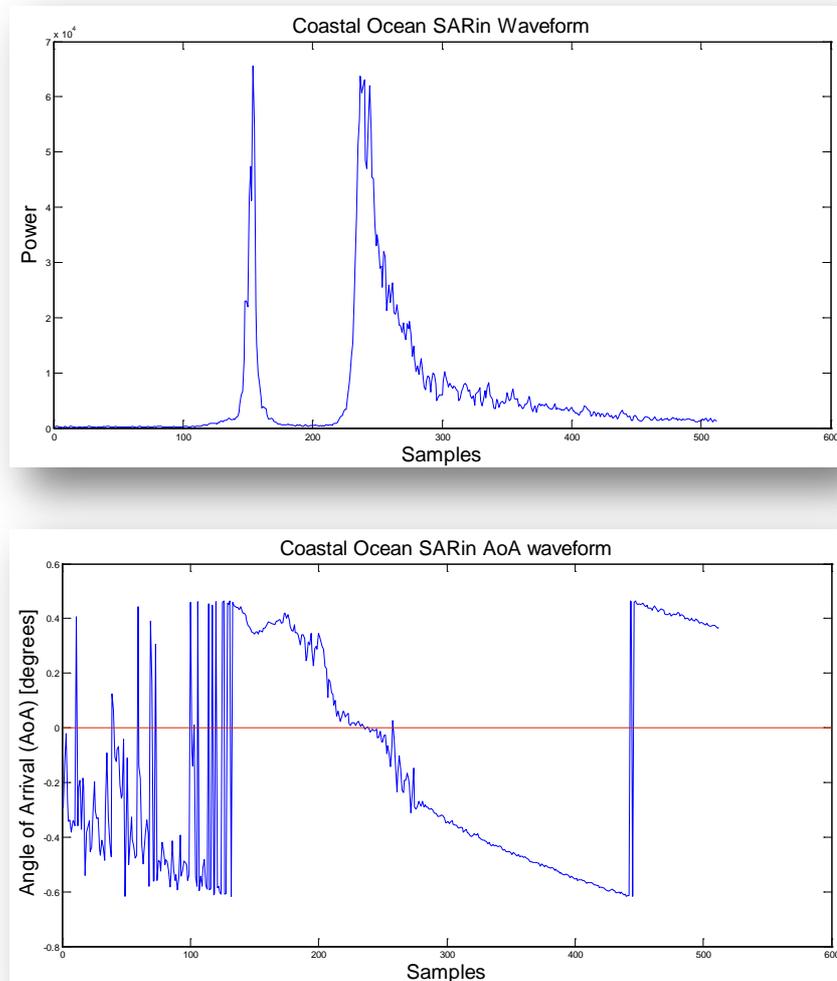


Figure 2. Example of a Coastal Ocean contaminated waveform. TOP: Power waveform. BOTTOM: Angle of Arrival waveform.

With this data in our hands, we can discriminate the different surface facets across-track, within the 300 meters surface stripe footprint. Therefore, in a track configuration like the one showed in Figure 3 following the same Figure 1 footprint schemes, the different facets from ocean or land (here at the right side of the sub-satellite track) will be present in the waveforms range bins with their correspondent AoA, permitting the discrimination of the specific nadir ocean signal within the waveform.

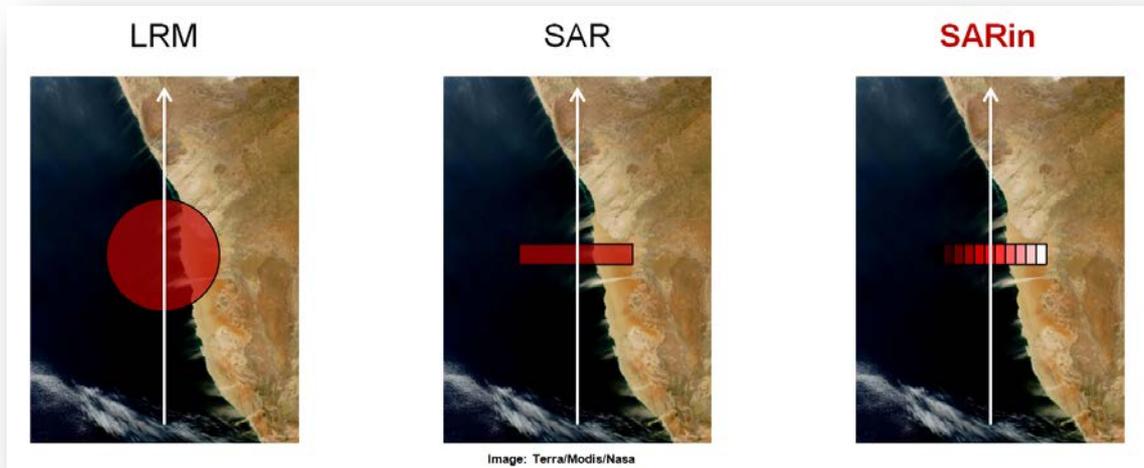


Figure 3. Superimposed ideal CS2 track and modes footprints on a Terra-MODIS image.

Knowing that our interest here is focused in the Sea Surface Height (SSH) estimates, and that a nadir AoA (AoA = zero degrees) is what we expect to detect from the sub-satellite track ocean surface while approaching the coast, any AoA response far from zero will be considered in this study as a target to be avoided.

This is the general approach of this small scientific investigation within the CP4O project that has the aim of improving the sea surface topography estimation near the coast with CS2 SARin data (AD-1).

3 Work Done

In this chapter we will explain the development of the work done in the investigation. The dataset used for the study and the geographical coverage are also hereafter described.

3.1 Input Data Set

A preliminary exercise done in this study was dedicated to analyse and select the set of products to be used for the investigation.

The different products were expected to show different scenarios in coastal zones. Two initial geometrical configurations were required: some of the passes with the sub-satellite track parallel to the coast, and some others perpendicular. This idea comes from an initial theoretical assumption of two extreme cases: a clean track with no visible land contamination in the perpendicular case; contrary to the previous, a parallel track to the coast that should show several signs of land contamination.

Once we checked the affected geophysical parameters at L2 products outputs, other variables were added to the selection exercise, such as the terrain topography characteristics. Therefore, we considered also different coastal land scenarios (low lands or cliffs), and diverse coastal zones with cays, reefs and others that can provide us with specular waters close to the track.

Several products were analysed from an initial set of more than 50 products in order to find a proper quality basis for the study. This second selection was mainly based on the analysis of the SSH results at level 2. First, we discarded those tracks showing stability in SSH estimation along the track. Second, for those not satisfying the first condition, thus having a SSH variable performance degradation, we selected several products from different topographic scenarios. From the selected track we analysed not only the sea level time series, but also we reconstructed the geolocation from the L1b computed through the analysis of the AoA. This re-geolocation permits us to easily detect in which cases the retracked surface facet do not correspond to the sub-satellite track (nadir).

The selection of the proper standard ESA product format for the study was addressed. The L2 products include the geophysical parameter SSH, but we also needed extra information as for instance the retracked sample of the waveform (Epoch) in order to perform a comparison with our CP4O product results. Is for this reason that we requested access to the L2I (Level 2 Intermediate / Interim) products where all this information is available.

In addition to the previous, L1b products are essential for performing the work on the different processing stages of the algorithms explained here further. In these products, among other information, we find the three available waveforms that characterize the received echo in the SARin mode, and permit a sample by sample analysis not possible with L2 product outputs; they are the power waveforms, the phase difference (thus AoA) waveforms, and the coherence waveforms.

3.2 Area of interest

In the CryoSat-2 mission there are several SARin mode masks created all around the globe, with a main focus in steeply sloping ice-sheet margins, over some geostrophic ocean currents and over small ice caps and areas of mountain glaciers. Therefore, our coastal ocean field of interest is not covered by SARin patches but by chance.

For the development of this study, a particular SARin mode area was requested: the Cuban Archipelago, between latitudes 19 and 24 North and longitudes 73 and 86 West (see Figure 4). It was approved by ESA and implemented on the 1st of October 2012, when the CS2 mode mask V3.4 was activated.



Figure 4. SARin patch covering the Cuban Archipelago.

This area is an excellent area for studying the coastal ocean. In it, we can find lowlands in the most part of the Cuban coasts, and also cliffs in the southeast coast near the Sierra Maestra mountain system. Moreover, there is a very complex combination of cays, reefs and islands all around the main island.

The Cuban Archipelago is the most exploited zone for the study, but not the only SARin patch considered. In Chile, although almost the whole SARin patch covering the Andes has no coasts in it, there is a little part of it in the North near Peru around latitude 20°S that has proven to be interesting for the study, and also some tracks in the South near Chiloé around latitude 44°S (see Figure 5).

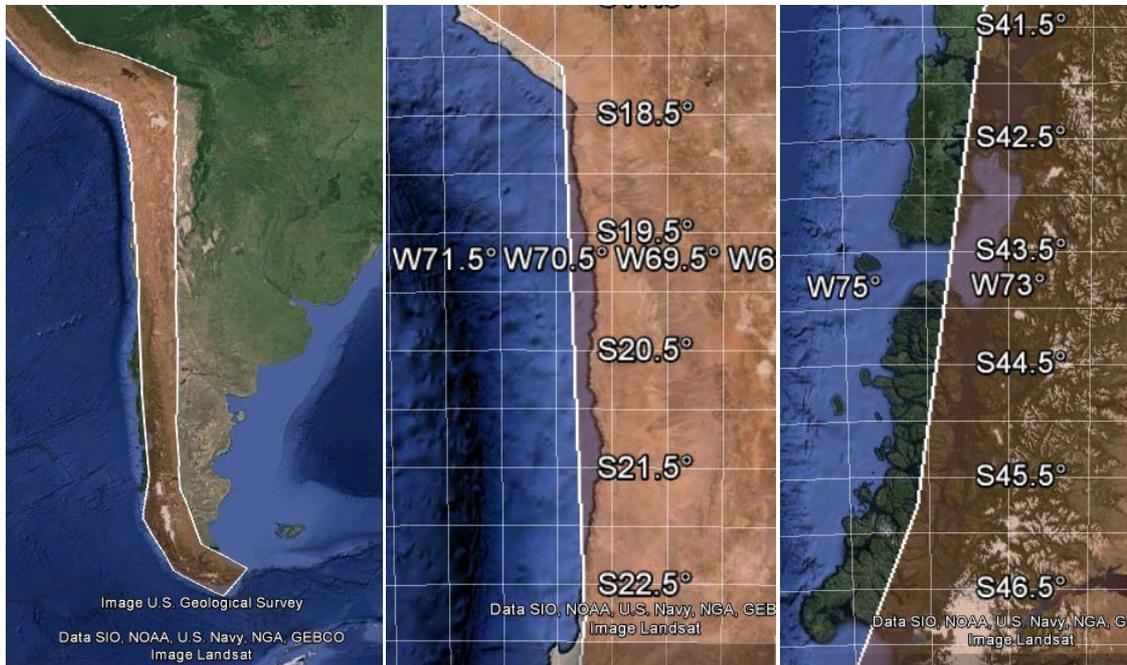


Figure 5. SARin patch covering the Chilean section of Los Andes (left), with coastal areas on the north (center) and in the south (right).

Also SARin patches over some coastal zones of Norway, Iceland and Canada (see Figure 6) have been considered, but with the constraint of having the risk of sea ice situations, that could invalidate the study.

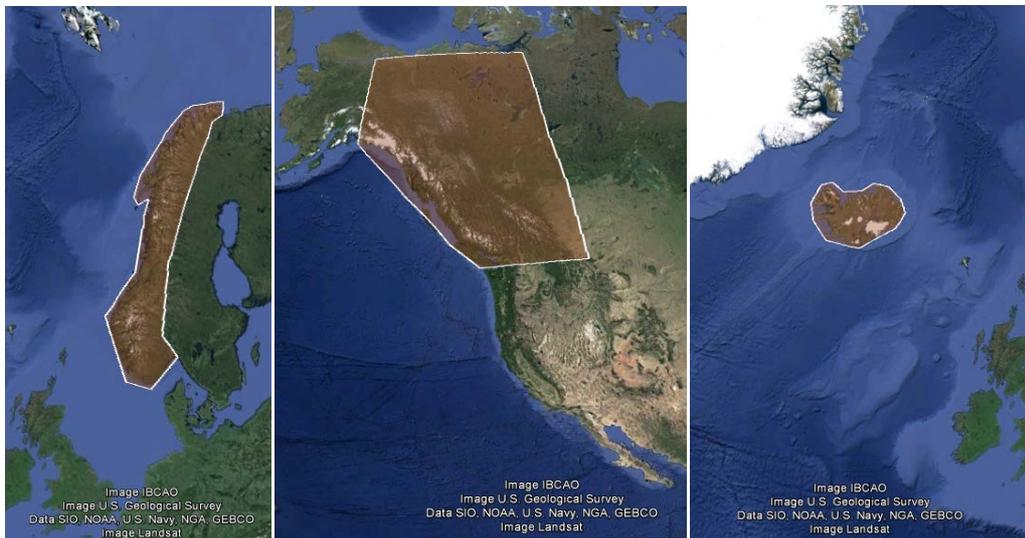


Figure 6. SARin additional zones over Norway (left), Iceland (centre) and Canada (right).

3.3 Algorithms

The whole investigation programming activities were developed in Matlab, from the ESA products reading routines up to the final results. The “L2 processing” tasks have been developed beyond the CP40 project, by isardSAT funding, with the aim of making possible a validation of the study results.

The data processing can be simplified into a two steps algorithm (see Figure 7). The L1b and L2l products are read and processed in a first stage, producing an input for the second stage that consists in a “seed” for the retracking. Then, the second step will retrack the waveforms coming from the L1b product, helped by the given seed. The L2l products are only used for the cross-comparison and validation of the final outputs.

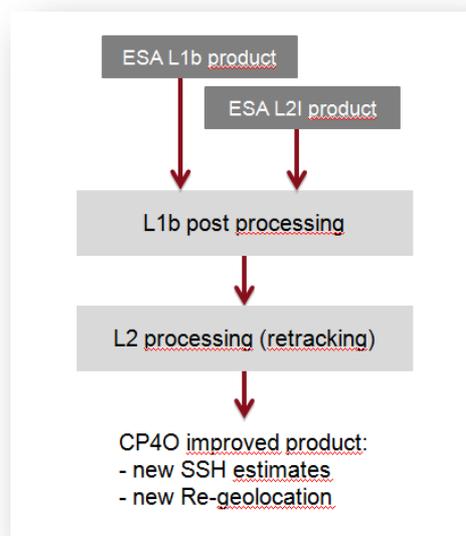


Figure 7. Simplified block diagram of the SARin investigation processing.

In Figure 8 it is showed a more detailed data flow of the algorithms that have been developed within the first stage of the processing, the “L1b post processing” block of Figure 7.

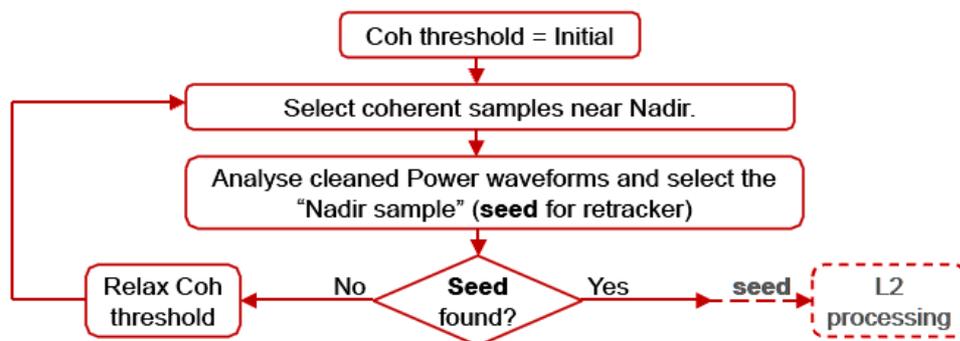


Figure 8. Block diagram showing the general algorithms logic of the L1b processing stage of the study.

The idea developed for the post-L1b processing stage is an iterative process that produces the final “seed” for each of the waveforms coming from the L1b product. The steps are:

- First, a “cleaning” of the waveform is done by selecting the coherent zones. This cleaning is based in a SARin mode specific information, that is the Coherence Waveform, showing how similar are the signals coming from the two antennas. If the Coherence threshold is too restrictive for the waveform cleaning, then the threshold is relaxed and another iteration begins. It has been necessary to dedicate a special effort to tune the coherence threshold for solving the trade off between covering a wide set of different scenarios (need of a relaxed coherence threshold) and selecting the precise waveform zone to work with (need of a restrictive threshold).
- Second, a new threshold is set, but now the variable in play is the Angle of Arrival (AoA) waveform. Here again, the tuning of the threshold needs a special dedication. It must let work in the next step with the zone of the waveform near nadir, but the meaning of “near” shall be well assessed and validated with the widest range of coastal scenarios. The AoA variable is the most important variable to be investigated for the purpose of this work, the variable that represents the meaning of the SARin mode. Without it we could not know which of the samples is near the sub-satellite track (nadir).
- Third, the Power Waveform is analysed, specifically the samples satisfying the above thresholds restrictions of Coherence and AoA. A unique sample has to be selected to serve the retracker (next step) as a guide to start the retracking process, the sample that is called “the seed”. This seed is the final objective of the whole L1b post processing stage, and has the coarse resolution that comes from a SARin waveform range bin. That is why we needed a second stage consisting in the L2 retracking operation to validate the results with the needed accuracy that enables a direct comparison with the ESA L2 results.

The representation of a specific study case in Figure 9 is worth for clarifying how the first two steps are working to discriminate the off-nadir targets. All the targets off-nadir on the left side figure are avoided and eliminated after the first two steps, as shown in the right side figure. The third step will work with the right side figure waveforms. In this example we see that the onboard tracking is following the most powerful peaks, but not necessarily the nadir sea surface signal, something that can be checked in Figure 10. The light orange nadir zone is what has to be selected, as done in Figure 9.

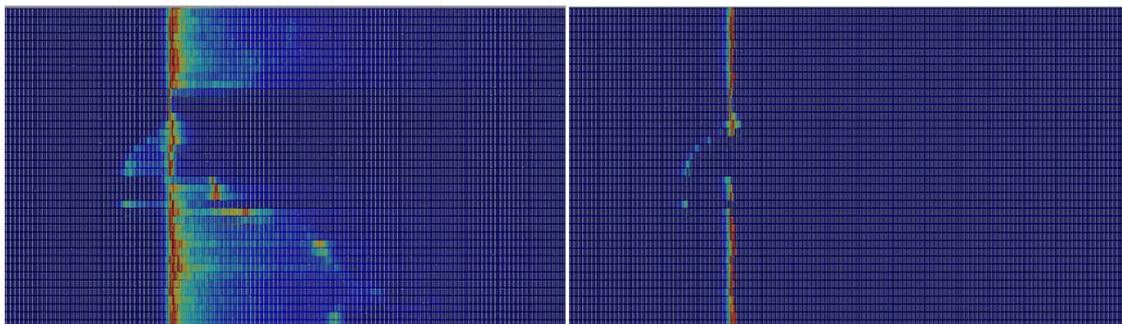


Figure 9. Power waveforms along the track in one of the study cases. Y axis is number of waveform; X axis is number of sample. The colour goes from cold (blue) meaning low power to hot (red) meaning high power. Figure on the left before and figure on the right after the waveforms cleaning.

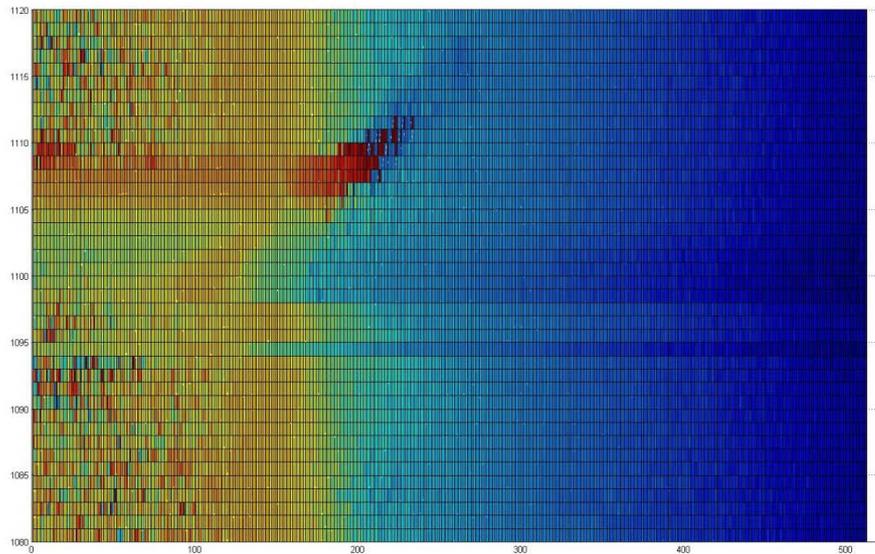


Figure 10. AoA waveform corresponding to the same case of Figure 9. Y axis is number of waveform; X axis is number of sample. The colour goes from cold (blue) meaning low AoA to hot (red) meaning high AoA. The light orange zone corresponds to near nadir echo.

The second stage of the overall algorithm is a retracking computation that is inherited from the SAMOSA model with a long series of modifications that makes it a really different new model, that we will call the isardSAT retracker in forthcoming sections.

The isardSAT retracker has been adapted from SAR mode to SARin mode, Baseline B, and other modifications. All its characteristics and development will be published soon in 2014 in a paper written by Cristina Martin-Puig from isardSAT. Cristina Martin-Puig is one of the main investigators that developed the well known SAMOSA SAR mode retracker. Her arrival to isardSAT gave her opportunity of gaining acknowledge about the whole L1b processing chain details, something that have been crucial for the improvements of the isardSAT retracker model.

Thanks to this L2 processing capability, the overall algorithm is able to estimate the SSH with an accuracy similar, if not better, than the one coming in the ESA products. Also the AoA determination will gain in accuracy from the retracker epoch computation, impacting the re-geolocation results.

The final output most important fields are:

- the SSH estimates, based in the satellite orbit altitude and the retracker epoch output.
- the latitude/longitude recomputed based in the retracked sample AoA and the original sub-satellite track.

The output results are showed in the next sub-chapter.

3.4 Results

This section aims at providing the performance quantification of the overall processing showing a series of final results from various case studies.

Results will be shown as means of a set of figures following a standard configuration that consists in superimposing information on the map of a particular coastal zone of interest. The superimposed information will be the ESA (read from the L2I products) and the CP4O (our processing output) retracked geolocations, including SSH estimates over the track's section of interest. This representation seems to be the best manner to provide evidence of erroneous SSH estimates due to the non proper focused retracking near the sub-satellite track (i.e. nadir).

A visual example of how the ESA L2 products retracked geolocations (computed from the AoA) can be far from nadir is showed in Figure 11. In this figure as the track is approaching the cays, very specular waters near the shore produces a powerful echo within the SRAL range window that is estimated in the retracking process as the backscatter from the surface, thus the output is not anymore giving the real SSH at nadir, but far from it. In principle this SSH estimation error should not occur and remedial actions should allow to solve for this error with some additional L1b processing. Note that SSH errors observed may be in the order of 2 meters when the retracked geolocation is not in nadir (AoA=0). In Figure 12 we can observe what is explained in terms of Power and AoA waveforms, as they have been read from the product of this specific example. The waveforms on the right are the ones impacted by the above named specular waters target, showing the clear peak off-nadir. The red (off-nadir) and green (nadir) lines are drawn for the sake of matching the possible retracked samples in the two waveforms (AoA and Power), making it evident that the AoA of the most powerful peak of the Power waveform is not near zero, meaning it is far from nadir.

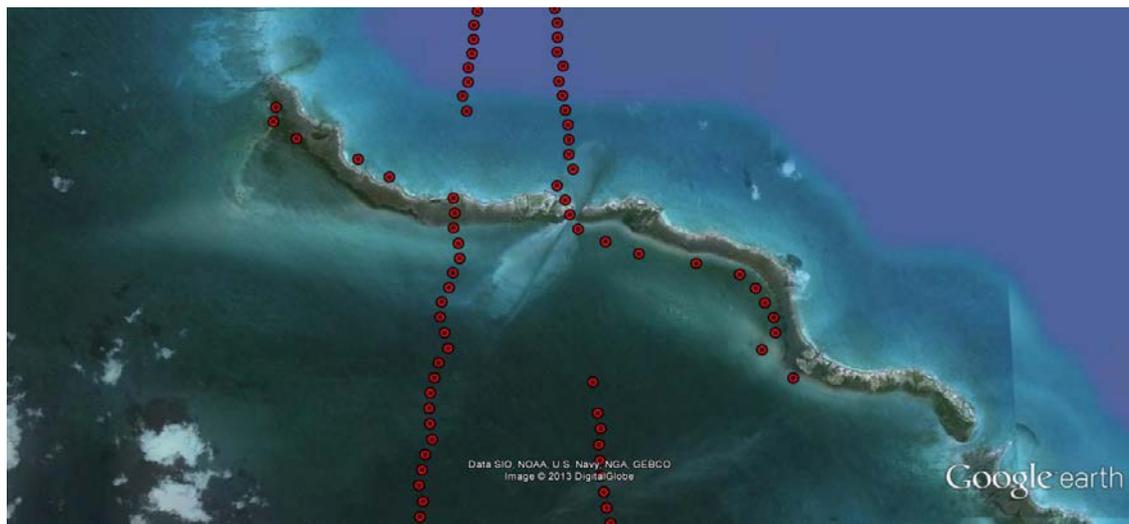


Figure 11. Superimposed ESA L2 product retracked geolocations on a track flying over Cuban cays.

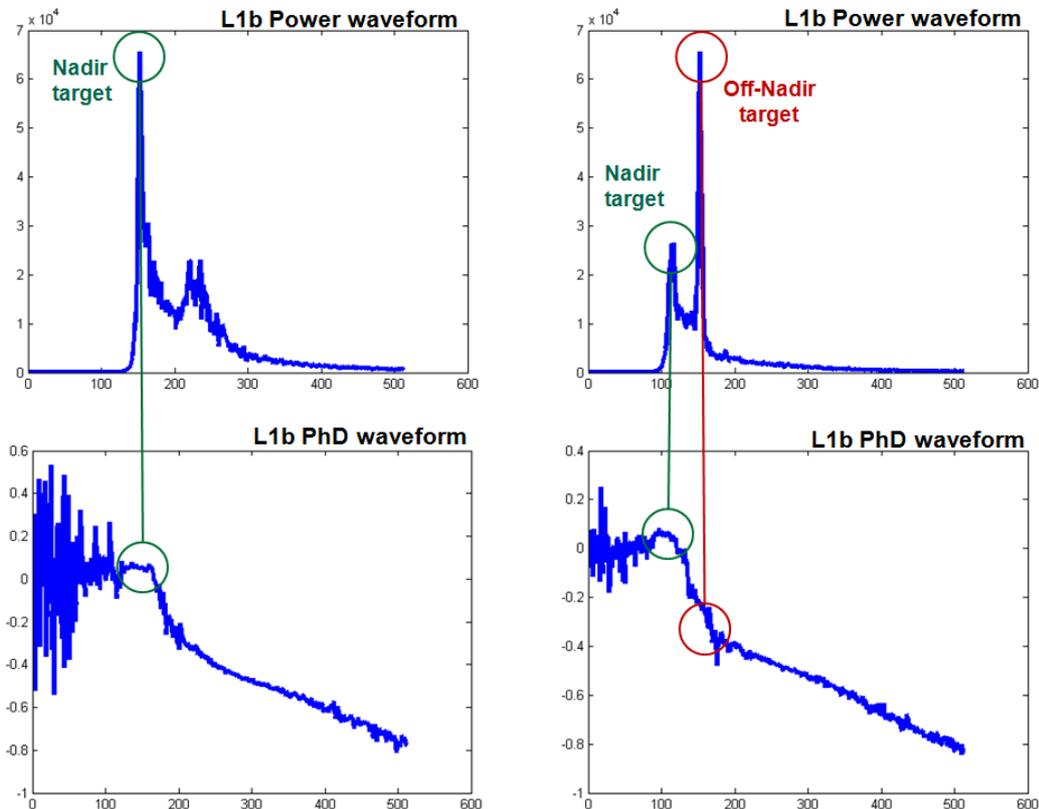


Figure 12. Power waveforms (top) and AoA waveforms (bottom) corresponding to an ocean scenario (left) and to a coastal scenario (right).

Here below, in Figure 13, Figure 14 and Figure 15 we show some examples of good results of the developed processing. In red, SSH as read from the ESA L2 products; note this show clear wrong estimates (in the form of jumps) when the track approaches the cays. This is well corrected by isardSAT processing strategy in CP4O; results are visible in yellow and show more stability than the previous all along the track.

Also the retracked geolocations are corrected from the red points (ESA L2I) to the yellow points (CP4O isardSAT processing). The ESA geolocations are detaching from the sub-satellite track from around 3 Km (case of Figure 15) to more than 5 Km (case of Figure 13).

The SSH estimates, being the focus of this study, is considerably corrected in the case of Figure 13, which shows SSH jumps of up to 6 meters. In the cases of Figure 14 and Figure 15, our processing SSH results show little errors compared to those from ESA L2, thus improvements can be appreciated. We have observed that SSH jumps in ESA results mostly correspond to waveforms with a Nadir Sea Surface target response close to a specular waters target, so near to each other that one is estimated erroneously instead of the other. Our solution allows for solving this problem, but when the targets are very close this is not possible as is visible in the last two figures where some jumps are not resolved with our approach.

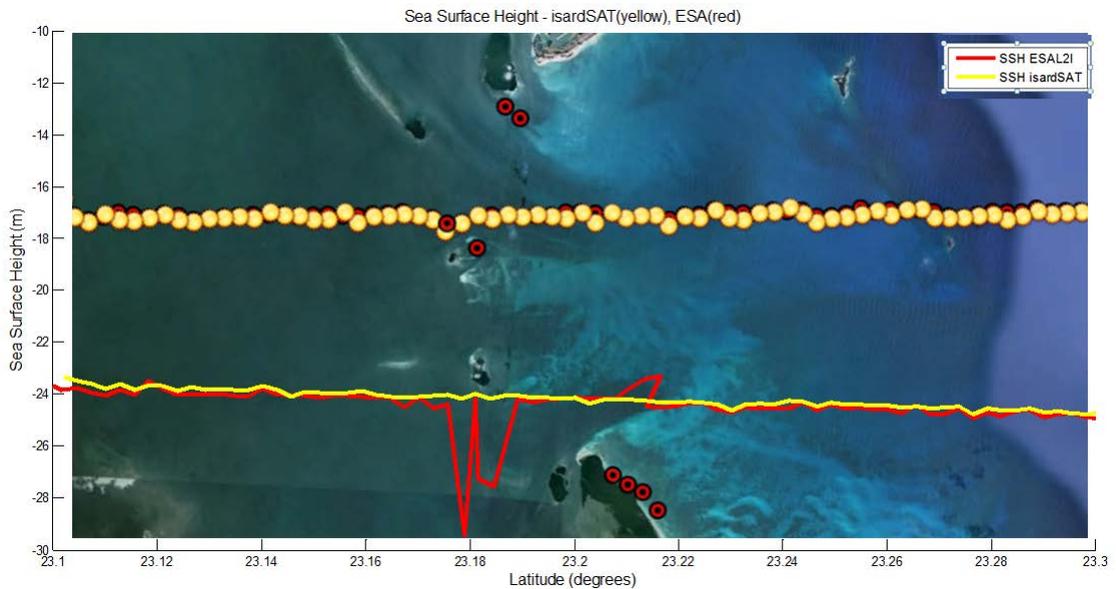


Figure 13. Example of processing results on Cuban North Cays near Varadero. Retracked geolocations are marked with points, and the SSH estimates (in meters) are marked with lines. Red colour is used for ESA L2 outputs, yellow colour is used for CP40 processing outputs.

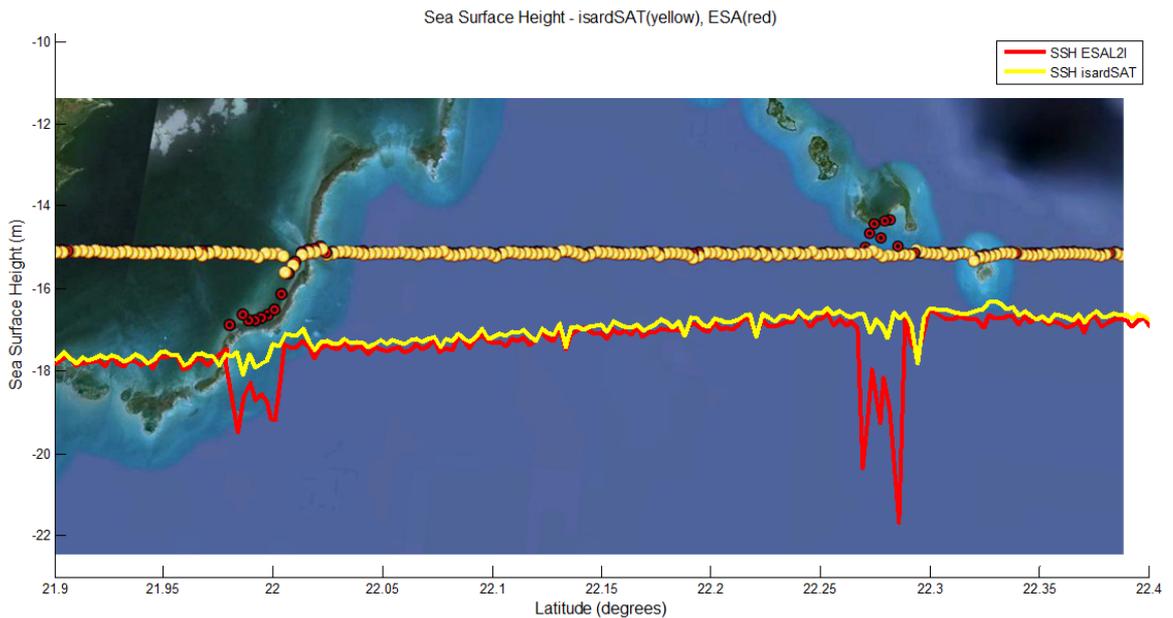


Figure 14. Example of processing results on Cuban South West Cays, ascending pass. Retracked geolocations are marked with points, and the SSH estimates (in meters) are marked with lines. Red colour is used for ESA L2 outputs, yellow colour is used for CP40 processing outputs.

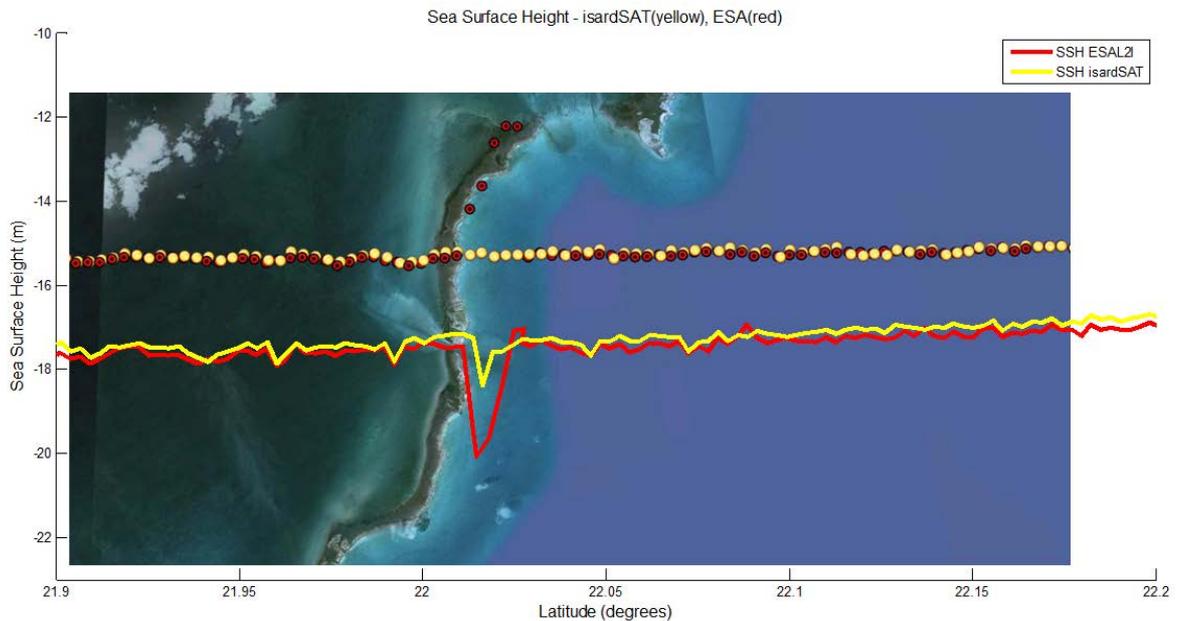


Figure 15. Example of processing results on Cuban South West Cays, descending pass. Retracked geolocations are marked with points, and the SSH estimates (in meters) are marked with lines. Red colour is used for ESA L2 outputs, yellow colour is used for CP4O processing outputs.

Although the solution proposed by isardSAT in CP4O has shown promising results, this is not always the case and hereafter we provide some evidence of false alarms (see Figure 16 to Figure 19).

Figure 16 shows a peculiar case in which although overall good performance along the track, there is a clear error; with a geolocation in line with other ESA results, near the shore of one of the little islands in latitude 22.62°N, possibly corresponding to a specular waters target. This record corresponds to the central waveform in Figure 17, where we can observe how the retracker is fitting (red line) the different waveforms. The retracked sample in the records before and after is the one corresponding to the nadir SSH (the peak on the left side of the waveform) but then, in the affected SSH record, the retracker fitted the leading edge of the peak on the right side of the waveform. The AoA of the peaks on the right side are far from zero, and the seed for the retracker was always indicating correctly the peak on the left side, but the retracker was unable in this case to follow the seed indications and avoid the wrong target. Note our retracker here is not applying any restriction to the lower and upper fitting parameters. If so, this error should be avoided, but it is rare to us that although providing the correct seed the fitting routine neglects it and iterates until the peak of higher intensity is retracked.

The second example, different from the previous but also interesting, shows the analysis of a parallel track to the coast-line on the Chilean North coast. This scenario includes cliffs and an airport relatively near the shore (see Figure 18). The combination of the ranges of the airport and the sea surface in the sub-satellite track, and the geolocations distance on ground between the two, make possible that both targets appear in the same range window. We observe in Figure 18 how the geolocations far from nadir are coincident with very specular airstrips surface locations.

Three records were corrected from five ESA L2 SSH errors. In Figure 19 we can see three consecutive waveforms fittings including the two CP4O algorithms bad SSH retrievals. In the three we can check the two clear peaks corresponding to the airport runway (lower range bin, specular surface echo shape) and to the sea surface (higher range bin, sea echo shape). The waveform on

the left is working as expected retracking the sea surface, while the other two waveforms on the middle and on the right are retracking the airstrips targets that are presenting a more powerful echo than the one coming from the sea in nadir.

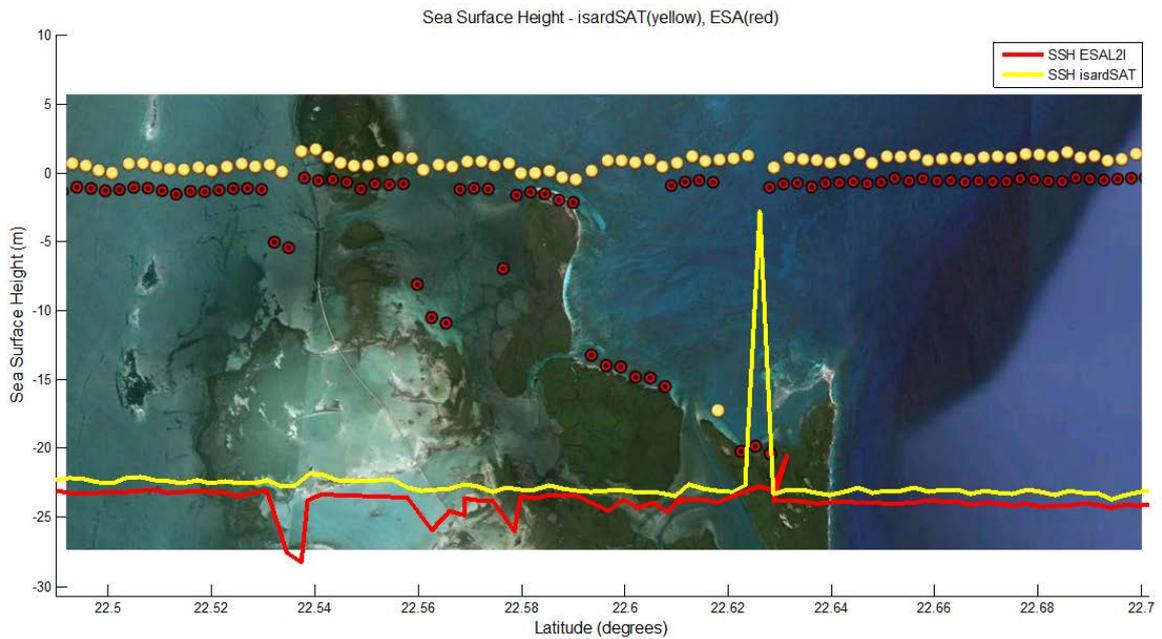


Figure 16. Example of processing results on Cuban North East Islands, descending pass. Retracked geolocations are marked with points, and the SSH estimates (in meters) are marked with lines. Red colour is used for ESA L2 outputs, yellow colour is used for CP40 processing outputs.

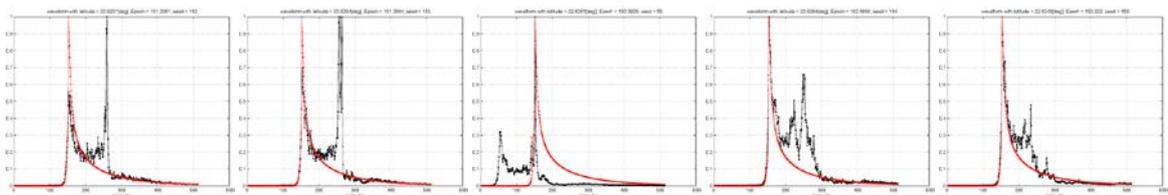


Figure 17. Five consecutive Power Waveforms corresponding to the track of Figure 16 around the CP40 processing SSH jump. The retracking fitted waveform is shown in red.

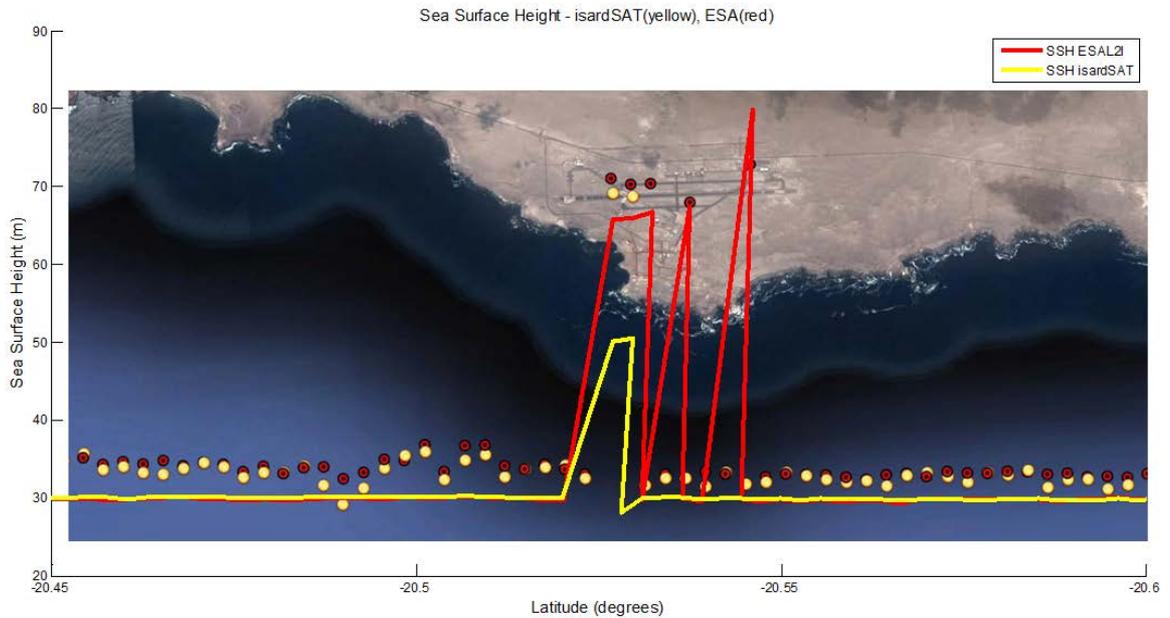


Figure 18. Example of processing results on the North Chilean coast. Retracked geolocations are marked with points, and the SSH estimates (in meters) are marked with lines. Red colour is used for ESA L2 outputs, yellow colour is used for CP40 processing outputs.

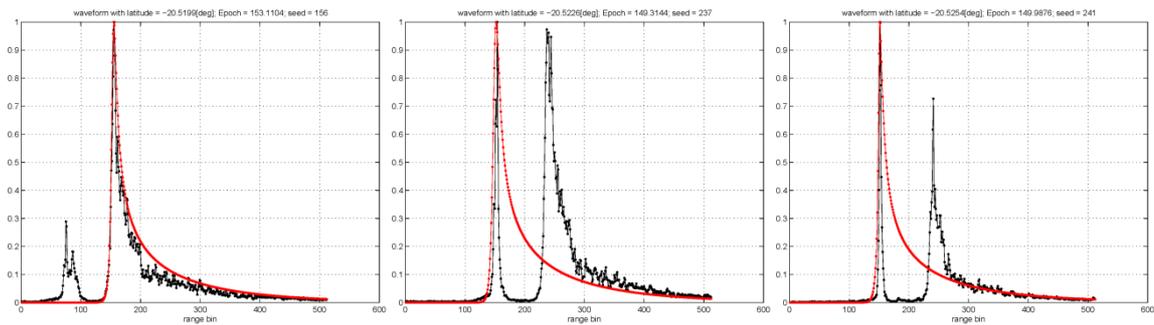


Figure 19. Three consecutive Power Waveforms corresponding to the track of Figure 18 around the first CP40 processing SSH jumps on the left side. The retracking fitted waveform is shown in red.

The cases of tracks perpendicular to the coast are not showed in this document because, as expected, they are not presenting so much signal degradation. In lowlands coasts the effect of possible land interferences is almost unappreciated, the SAR/SARin stripe footprint, in this case parallel to the coast already discriminates all the targets out of a 300m wide stripe. If the coast is abrupt (e.g. cliffs), in the cases of tracks crossing the coast line from the land to the sea the range window usually presents a certain delay until it follows the sea surface, so it is imposible to correct it because there is no sea signal in the waveform; and in the cases coming from the sea to the land, the waveforms are not usually so much affected as in the cases presented in this document.

Also the scenario of inland waters have been addressed, but it is not showed here because the ESA L2 results were good, so nothing had to be improved. It is the case of Titicaca Lake, in Andes.

Finally, there are cases with no solution with our approach. This is usually happening with waveforms with a too high level of contamination, or presenting a power peak very close to the ocean surface leading edge.

In Figure 20 we can see a power waveform corresponding to the track showed in Figure 21. It is an example of a contaminated coastal waveform that have been not solved by the isardSAT CP40 algorithms when trying to find the nadir SSH. In this particular echo we have several backscatter signals coming from different specular targets. Some of the possible targets are marked in Figure 21. This situation have no clear solution for the retracker to find the ocean surface height.

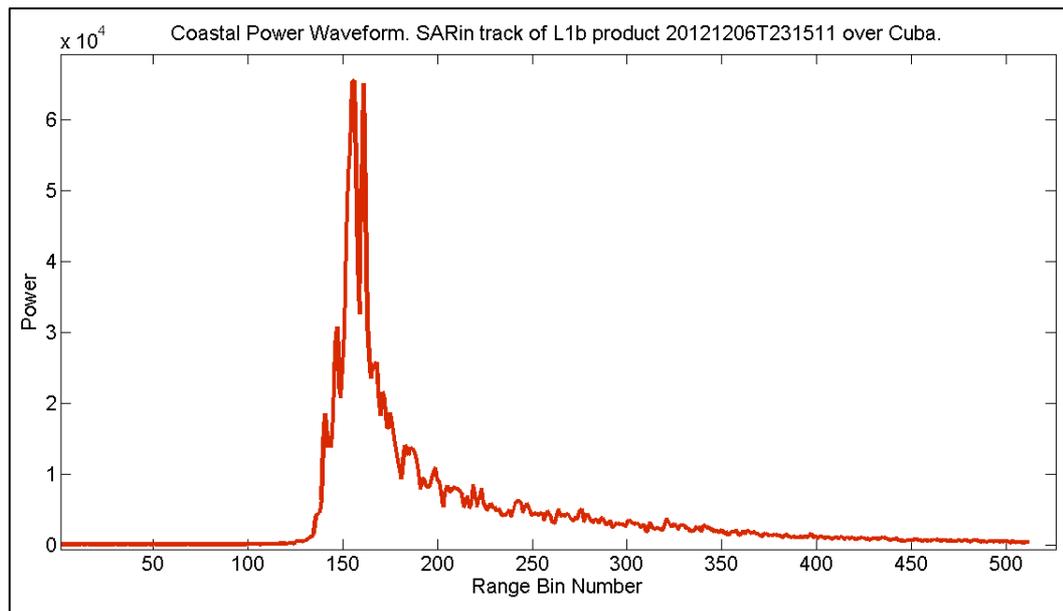


Figure 20. Example of distorted coastal echo. Track over the Cuban north coast near Puerto Escondido.

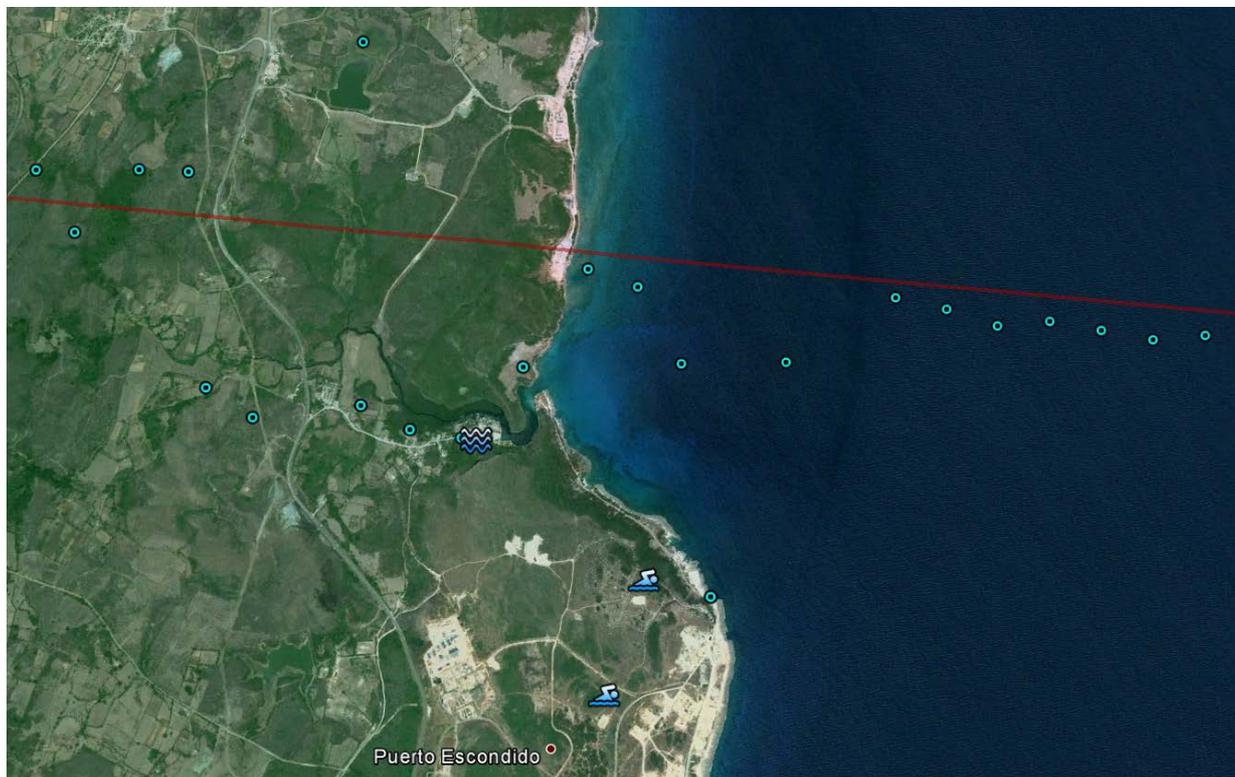


Figure 21. Track corresponding to the power waveform of Figure 20. Icons identifying different possible specular targets. Blue dots are retracked geolocations from ESA L21 products. Sub-satellite track in red solid line.

4 Conclusions

The main purpose of the study was to mitigate the impact on the sea level retrieval of the land contamination in coastal zones echoes using the CS-2 SARin mode products, specifically with the advantage of having the information of the Angle of Arrival of the different range bins.

It has been necessary to create a dedicated SARin patch for the study because of the lack of clear non sea-ice locations for the investigation. Many SARin zones were in high latitudes (e.g. Norway, Iceland, Canada, Finland). The possibility of studying the Cuban archipelago has been crucial for the developing of the study.

A large set of different coastal topography and track geometry cases has been considered in order to enable the use of this processing in a wide range of coastal areas.

Two different processing steps are developed. The first is a post-L1b iterative processing that produces a seed. The seed consists in a selected range bin that will be the input to the second processing: a retracking computation. This second stage is an activity that has been funded by isardSAT, and has proved to be essential for the direct comparison of the isardSAT CP40 work and the L2 ESA results. The final retracked range is expected to correspond to the nadir ocean surface.

The results presented in 3.4 show a clear improvement of the SSH estimates in several examples. We can say that the aim of avoiding the contamination has been fulfilled, and the sea level series along the tracks has been very much stabilised.

Several cases, as the three showed from Figure 13 to Figure 15, presented a very good performance of the overall processing. The SSH retrievals were highly improved, gaining in stability all along the track. Also the retracked geolocations are very much corrected, being all in line with the sub-satellite track.

Some other cases, as the two showed from Figure 16 to Figure 19, presented also improved SSH series, but with some records with errors, both in the retracked geolocations and the SSH estimates. The reason why it is happening is related to the second part of the processing: the retracking algorithm. The seed inputted from the post-L1b processing is correct, indicating the nadir range bin, but the retracker fitting routine neglects it and iterates until the highest power peak is retracked.

An extensive work has been done in tuning the post-L1b algorithm. The tolerances of the different parameters taken into account always presented a trade-off to be solved, compromising specific scenarios and favouring others. The best solution for covering the most of the cases was adopted.

But it is quite more complex to solve other cases. It is the case of tracks with very distorted waveforms by the contribution of many surface targets or also waveforms presenting a clear peak to be avoided but with its leading edge too close to the one corresponding to the nadir sea surface (see Figure 20 and Figure 21). The tuning of the algorithms of the study could not solve this kind of situations after a long series of trials, so it has been considered out of the reach of the current processing.

5 Work to be done

The examples of Figure 16 to Figure 19 were showing not acceptable results, although they presented several records corrected. These imperfections have to be addressed in a further phase of the investigation.

The situation of Figure 19 is identified as a clear issue that can be managed in further improvements of the algorithms, both in the post L1b processing and in the retracking algorithm. The visible detaching between the two targets signals in the range window indicates that there is space for a possible solution. The example of Figure 16 can be addressed in the same way, but here the ocean waveform trailing edge is not clean but distorted by another backscatter contribution.

The improvements to be done on the processing for solving the previous problems represent a different approach that have to be clearly developed and assessed in a later phase of the study. The limitations of the current solution can be possibly solved with a more detailed work on the available three waveforms in the third step of the post L1b processing (see 3.3), and also with the help of some auxiliary processing of along track range cross-correlations.

As stated in chapter 4, there will always be cases with no possible solution working with the adopted approach in this investigation, for instance too noisy echoes or power peaks too close to the one corresponding to the ocean surface.

Also the correction of the off-nadir elevations is a pendent issue, present both in the CP4O and the ESA L2 products. The results are not satisfactory in any of the two solutions, and also they are different. It is evident in the example of Figure 18, where the failed SSH estimates are substituted by results of the airstrips height, not corresponding to the real airport height; or also in the example of Figure 16 with jumps of more than 5 meters accounting for targets all coming from the sea surface.

The importance of the coastal zones for the human development makes it a valuable reason for boosting this kind of investigations and point out the need of a further evolution of the work here presented.