DEVELOPMENT OF RADAR ALTIMETRY DATA PROCESSING IN THE OCEANIC COASTAL ZONE

ESA/ESRIN Contract No. 21201/08/I-LG – CCN 3 (Phase2)

EWP3 – Deliverable D3.3

Development and Implementation of the Hyperbolic Pretracker

Graham Quartly

VERSION 2.0, 21 December 2011

National Oceanography Centre, European Way, Southampton, SO14 3ZH, United Kingdom

Tel: +44 (0)23 80596404 Fax: +44 (0)23 80596400 www.noc.ac.uk

Code COASTALT2-EWP3-D33 **Edition** 2.0 **Date** 21 December 2011

Client European Space Agency **Final User** -

	Name	**Signature**	**Date**
Written by	Graham Quartly		10/10/2011
Checked by	Paolo Cipollini		10/10/2011
Revised by	Graham Quartly		21/11/2011
Checked by	Paolo Cipollini		21/12/2011

DISTRIBUTION	**Affiliation**
Jérôme Benveniste, Salvatore Dinardo, Bruno Lucas	ESA
Paolo Cipollini, Helen Snaith, Phil Woodworth, Graham Quartly	NOC
Stefano Vignudelli	CNR
Jesus Gomez-Enri	U Cadiz
Cristina Martin	STARLAB
Joana Fernandes, Alexandra Nunes,	U Porto
Susana Barbosa	U Lisbon

Revision History

Issue	Date	Change
1.0	10 Oct 2011	Initial Release
2.0	21 Dec 2011	Revised documentation

TABLE OF CONTENTS

Acronyms used

AGC	Automatic Gain Control
COG	Centre of Gravity
ERS-1	European Remote-sensing Satellite - 1
ESA	European Space Agency
FTP	File Transfer Protocol
MDSR	Measurement Data Set Record
p.d.f.	probability distribution function
PRF	Pulse Repetition Frequency
RA-2	Radar Altimeter -2

Abstract

This report details the progress made on the design and implementation of a hyperbolic pretracker specifically for the Envisat RA-2 altimeter. First considerable care was taken to construct a waveform simulator whose characteristics matched those of the instrument — specifically, this meant some tuning to get leading edge slope and position of half-power point in exact agreement with the RA-2 altimeter. This simulator was used to generate waveforms as the virtual instrument overflies a small rectangular patch of enhanced variability (a "bright target") possibly representing glassy seas. From a number of such (noise-free) simulations spanning the narrow altimeter swath it was possible to use mathematical inversion techniques to produce a set of weights for estimating intensities of such hyperbolae, and then remove these features from the 2-D waveform space. This is demonstrated in a number of worked examples (datasets provided for validation). A key aspect of the technique is that it can equally well be used to compensate for weak targets (rain cells or absence of ocean due to land); this is particularly effective in the transits of Pianosa. This technique is much better at coping with multiple discrete targets than an approach that treats each waveform in isolation thus neglecting the contextual information from neighbouring waveforms. The implementation with real RA-2 altimeter data is still problematic, because there are movements in the tracker window, which can still be important at the sub-binwidth scale, yet are hard to correct for. A further challenge is estimating the intensity of hyperbolic features when not all of the feature is present in the waveform anomalies e.g. because of data over land generated using a different chirp bandwidth.

1. Introduction

1.1. Brown-like waveforms

The operation of an altimeter over a homogeneous surface is well-modelled, with the mean shape of the resultant waveform (see Fig. 1) being accurately described by the Brown model (Brown, 1977; Hayne, 1980). Indeed, when the mean shape from a large ensemble deviates from the idealised shape it informs us about minor imperfections in the space hardware of processing. A large number of algorithms have been developed to estimate certain geophysical parameters (chiefly range, significant wave height and backscatter strength) predicated on the assumption of a uniform reflecting surface.

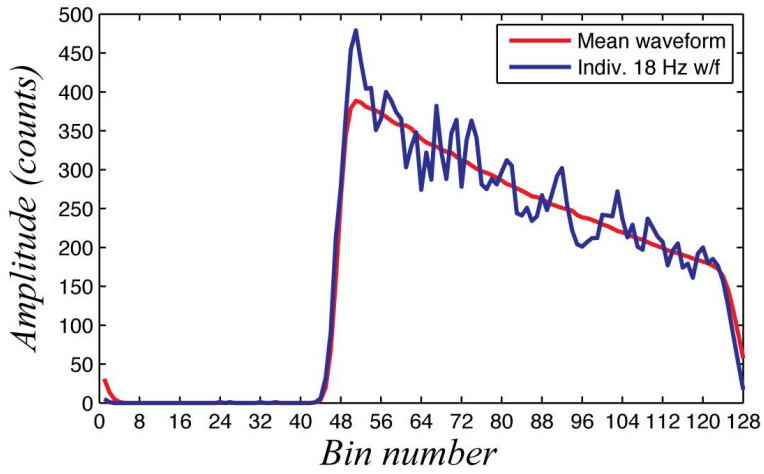

Fig. 1 : Typical ocean-like ('Brown') waveform. Example is for a wave height of 2.6 m, with red line showing mean expected shape and blue line the realization for a particular 18 Hz average waveform.

Individual average waveforms will differ from the mean due to the incoherent addition of signals with random phase, and thus the value for a specified waveform bin, *i*, will have a mean value w_i given by the Brown model, but a distribution governed by Rayleigh (fading) noise, with the distribution getting narrower as the number of pulses averaged together, N, increases. For an altimeter with a fixed pulse repetition frequency (PRF) and a given ground velocity, the choice of N is a compromise between greater accuracy (large N) and finer along-track sampling (shorter time intervals i.e. lower N). For the RA-2 on Envisat, the designed operating characteristics at K_u-band are to average 100 pulses together, and in those 0.057 seconds the sub-satellite point travels 374m on the Earth's surface.

1.2. Hyperbolic signature of weak/bright targets

The preceding section assumed all facets of the Earth's surface reflected with equal intensity. However, under certain circumstances, small patches (relative to the pulse-limited footprint of ~10 km diameter) may have greater or weaker reflections than their environs. Enhanced backscatter is typically associated with glassy seas either due to the sudden abatement of wind or sheltering of a region by land, or the presence of an oil slick (biogenic or from a tanker). A patch of apparent reduced backscatter may be due to a small island or attenuation by rain in the overlying atmospheric column. For any such case, the change in a series of waveforms can be modelled as the addition/reduction of a localised scatterer on top of the general homogenous scattering surface that gives rise to the Brown waveform.

As the altimeter track passes close to or over the affected region, the waveform bins to which it contributes follow a hyperbolic trajectory in waveform space (Fig. 2). The shape of the hyperbola is governed by the orbital and instrument characteristics (ground-track velocity, time interval per averaged waveform and bin width). Only the position of the apex of the hyperbola is governed by where the different patch lies along the altimeter track and how far it is off track.

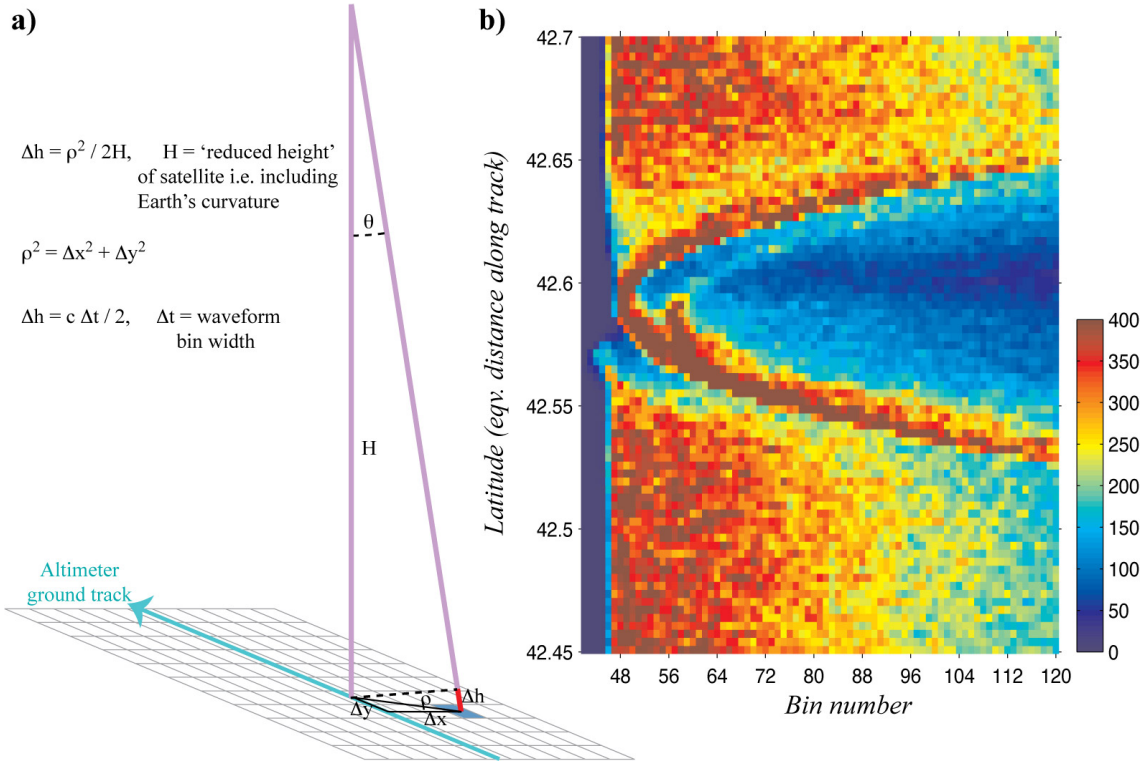

Fig. 2: a) Schematic of simulation, with sea surface represented by many square pixels, with properties of backscatter strength, height above sea level and wave height. The extra path length relative to nadir reflection, Δh, is a function of distance of the point from nadir, which is itself dependent on time along track. b) Simple example of signal from a bright target (Envisat ground track 128 near Pianosa, 7th October 2003).

2. Numerical modelling of waveforms

2.1. Achieving the correct mean shape

The numerical model simply adds together the delayed contributions of many reflecting facets from a grid of pixels covering a wide enough area (see Fig. 2a). The delay is calculated relative to reflection from nadir, with the delay to spherical spreading allowing for both the height of the altimeter above the surface of the Earth and the radius of curvature of the Earth's surface (see Quartly, 1998, Quartly et al., 1999). The effect of wave height is produced by giving each reflecting facet a Gaussian p.d.f. of expected delays, rather than the whole facet contributing at one single delay. (The width of the p.d.f. simulates both the contribution from wave height and the emitted pulse width.) The model can also incorporate changes in mean height to cover reflections from islands, but that is not used here, as islands (being generally weak reflectors) only contribute noticeably when their reflected contribution is ahead of the waveform leading edge. In such a case, the signals are easy to identify and have little effect themselves on waveform retracking as the added contribution within the main body of the waveform is very weak.

There is a common misconception that the RA-2 on-board tracker aims to position the half-power point at bin 46.5. In fact the on-board system makes window position adjustments simply to keep the majority of the useful signal within reception range. Thus a close examination of normal ocean returns from the RA-2 show the half-power point to be at about bin 46.3 for very low wave heights (<1m) stretching out to beyond bin 50 for wave heights exceeding 7m. In order to be able to produce fully realistic simulated waveforms for a wide range of surface conditions, there was some tuning of parameters to get the leading edge slope and position exactly as required for given wave height conditions. The result is shown in Fig. 3.

A further finessing in the implementation has been not to merely generate mean waveforms every 374m along track (corresponding to the 18 Hz sampling of the Envisat RA-2), but to produce them eight times as often (i.e. every 47m) and average in groups of eight to mimic the slight smearing that the instrument achieves in summing 100 pulses spread over 374m. This minor adaptation leads to a slight smearing of the sharp edges that would otherwise occur with the intended box-like features.

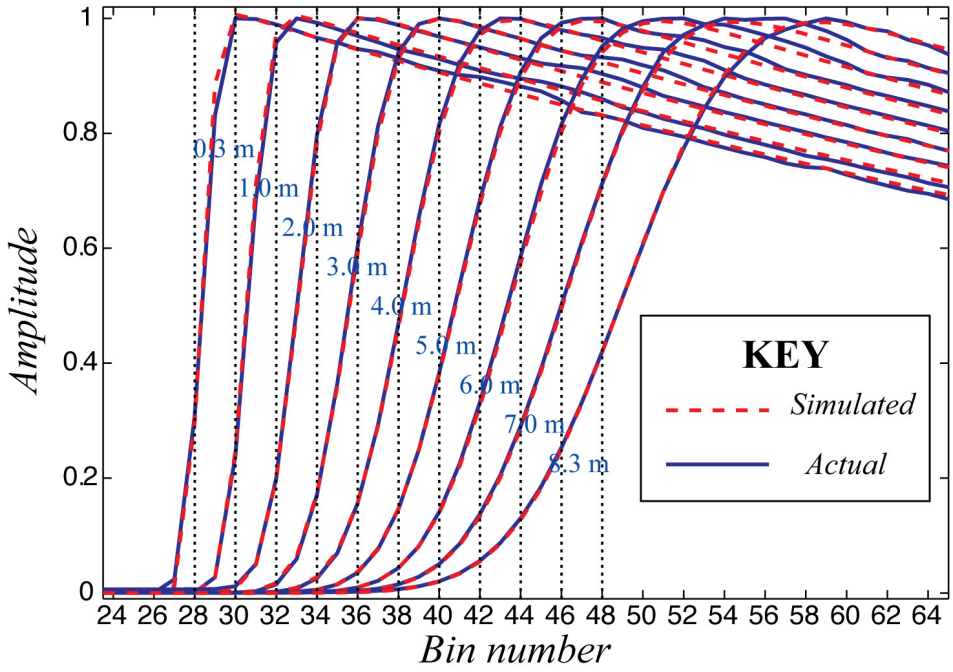

Fig. 3: Modelled and mean observed RA-2 waveforms superposed, with a close-up of leading edge shape, with each successive pair offset by 2 bins. (Vertical dotted lines 2 bins apart show that mid-point of leading edge recedes as wave height increases.) Note that each actual mean shows a slightly increased level (relative to its simulation) at the start of the trailing edge — this is an enhancement of 2% between actual bins 61 and 63, and is assumed to be related to instrument hardware or processing.

2.2. Modelling the effect of enhanced patches

Idealised simulations are then run with the altimeter flying over a surface of uniform backscatter apart from in one square or rectangular box where the reflectance is enhanced such as in Fig. 4. The model assumed is one with an isotropic scattering distribution within the narrow (1˚-2˚ wide) cone that can contribute timely returns. Square or rectangular cells are used because they can be tessellated to provide a representation of larger variable strength features. Given a set of such boxes spanning the modelled surface, with their associated characteristic hyperbolae, it is possible to develop a set of weights to estimate the anomaly in surface reflectance at each location, and use this estimated signal strength to remove the pertinent signal. The set of weights are determined by setting up a number of simultaneous equations and solving them, with the added constraint of minimising the sum of squares of the weights. This constraint is not simply added to lead to a unique solution for an otherwise under-determined set of equations; it corresponds to the minimum sensitivity to the effects of fading noise (Quartly, 1998).

In mathematical terms, using $a_{rs}(i,j)$ to signify the anomaly pattern in 2-D waveform space (denoted by bin no. *i* and waveform no. *j*) due to a region of

enhanced backscatter in 2-D physical space (denoted by pixel *r* away from nadir and index *s* along the sub-satellite track) and $w_{pq}(i,j)$ being the set of weights to evaluate the power excess at a pixel *p* away from nadir and index *q* along track, then

$$\sum_{i,j} w_{pq}(i,j)\, a_{rs}(i,j) \;=\; \delta(p-r)\,\delta(q-s) \qquad (1)$$

where $\delta(x)$ is the integer delta function having value 1 for $x=0$ and zero otherwise. The formula in (1) is effectively a set of simultaneous equations in $w_{pq}(i,j)$ to be solved for all relevant values of *r,s*, which is to be solved subject to:

$$\min \left\{ \sum_{i,j} w_{pq}^{2}(i,j) \right\} \qquad (2)$$

The above equations are solved to yield w, and then the 2-D summation on the left-hand side of (1) is used to calculate the strength of any potential bright target. A different set of weights is required for each position away from nadir (*p*), but for different positions along track (*q*), it is the same set of weights shifted along.

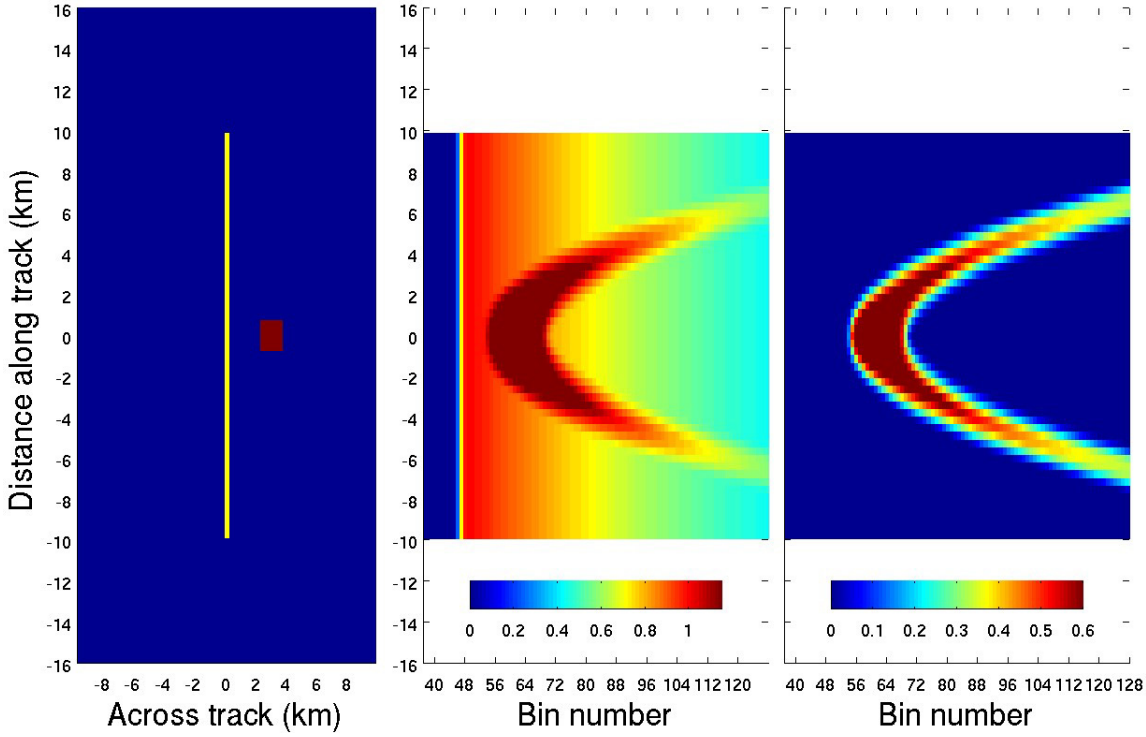

Fig. 4: a) Representation of surface reflectance model, showing square region of increased reflectance, with yellow line showing altimeter ground track over surface. b) Simulated waveforms, c) Calculated anomalies (see section 2.3). Power is normalised so that feature-free noise-free signals have a peak of one.

2.3. Defining waveform anomalies

The basic Brown waveform signal needs to be removed from individual waveforms, not only to enhance the visibility of the hyperbolic features, but also to

enable them to be estimated in a robust and unbiased way. For this a reference waveform shape is required. Although the numerical model can generate perfect waveforms for a completely homogenous surface, it is preferable to use one calculated from actual Envisat data, as that will contain the instrumental artefacts that will also be present in that data to be cleaned. A database could be constructed with appropriate large ensemble means for each possible wave height. However, here we implement the approach of calculating a representative mean from the unaffected waveform data. Thus a median is determined over 100 to 200 waveforms, using the assumption that intermittent features will make a negligible impact unless present in most of the waveform set being analysed.

As the contribution of fading noise is proportional to the expected value at each bin, it is advantageous to divide each observed waveform by the reference shape, rather than subtract it. This division renders the characteristics of the noise uniform throughout the useful waveform space. A miniscule amount, ε, amounting to 10^{-4} of the reference waveform amplitude, is added to both the individual waveforms and the median (reference) to prevent division by zero errors occurring in the thermal noise region. Thus

$$\text{Fractional anomaly} = (\text{waveform} + \varepsilon) / (\text{reference} + \varepsilon) - 1 \qquad (3)$$

The technique for removing bright target signals can then be summarised as:

1) convert individual waveforms to fractional anomalies (via eq. 3; see Figs. 4 & 5)
2) estimate and remove the hyperbolic features associated with localised bright/weak targets
3) convert the cleaned up fractional anomalies back into waveforms (using the reverse of eq. 3, with the same ε and reference waveform), ready for the application of standard waveform-fitting techniques.

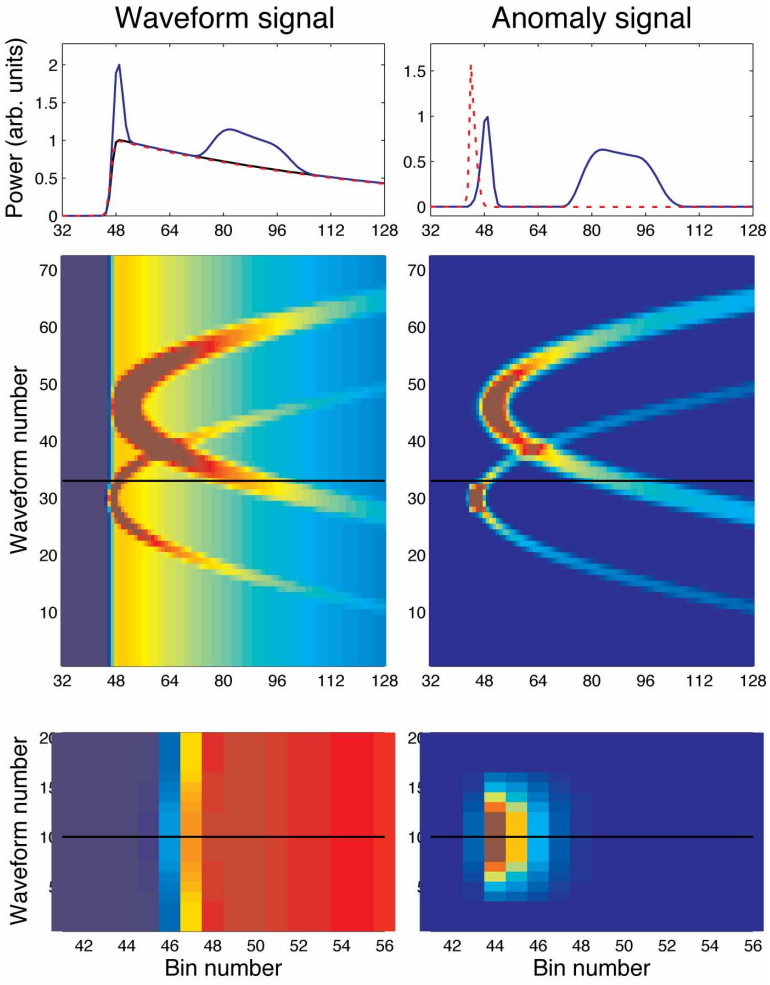

Fig. 5: Further illustration of calculation of waveform anomalies (fractional changes). Middle panels show simulation of a couple of simple "bright targets" (of different size, intensity and distance off track). Lower panels show a zoom for a separate simulation where the waveform reception window moves such that waveform is 0.25 bins earlier than in reference. Top panels show individual waveforms — black is waveform no. 1 (effectively the reference), blue is waveform no. 33 of "bright targets" example and red dashed is waveform no. 10 of example of waveform advance (which almost overlies the reference waveform).

3. Evaluation of Method

3.1. Idealized cases

In the simplest example, waveforms are simulated for an altimeter flying over an area of homogeneous backscatter except for a square region of enhanced reflectance. The hyperbolic-fitting algorithm is set to estimate backscatter strength for an array of squares of the same size, and the fitted features removed. Not surprisingly, with such idealized conditions (as were used in developing the weights for the inversion) the match is perfect. The case illustrated in fig. 6 is only complicated by the presence of fading noise (corresponding to the addition of 100 incoherent pulses); in such simple conditions the fit is still very good, although visually the 2-D plot of the residuals still has a hint of a hyperbola. The top panels of the figure illustrate a particular 18 Hz waveform before and after the hyperbola fitting and removal. Initially the waveform has a shape not easily matched by the Brown model (Fig. 6a), but after the pretracker algorithm has been applied (Fig. 6b) the resultant waveform data are much more amenable to a standard retracking algorithm, rather than needing a "Brown plus peak" model.

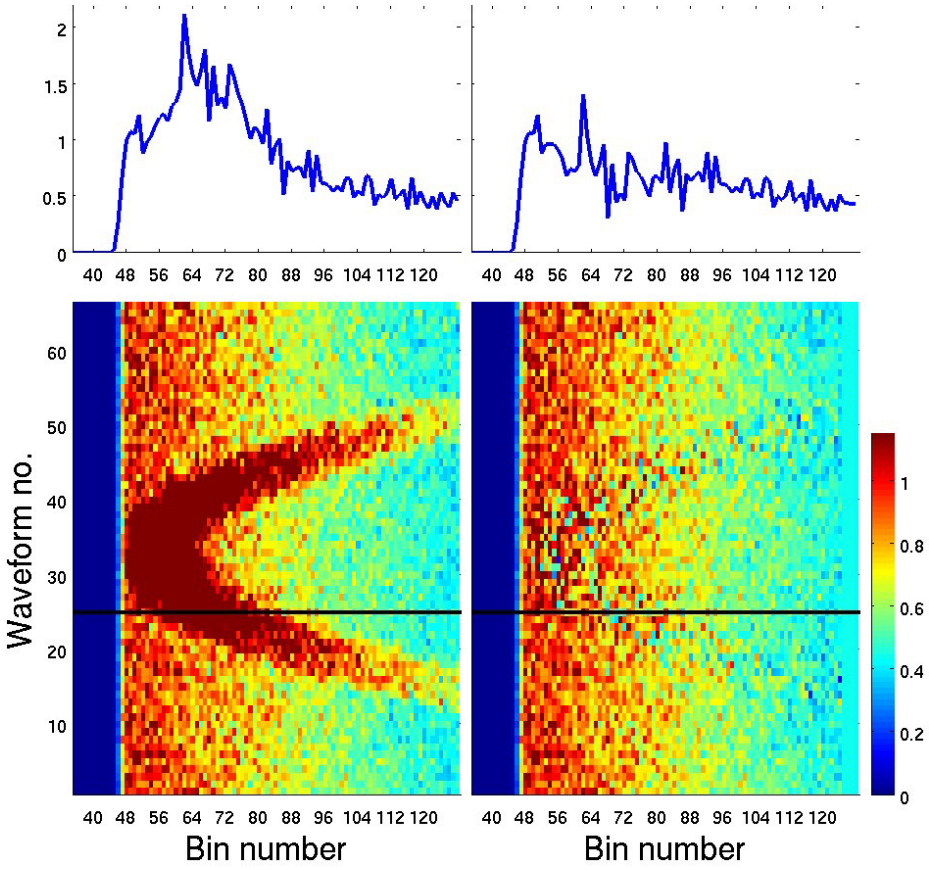

Fig. 6: Bottom left panel shows series of waveforms over a simple simulated "bright target" with fading noise added; bottom right shows the result after estimating and removing the hyperbola. Top panels show waveform no. 26 in detail. Power is normalised such that feature-free noise-free signals have peak = 1.

3.2. More complicated simulations

A more complicated (and realistic) situation is when the model for generating the waveforms is not congruent with that to be used for the fitting and removal of hyperbolae. If the fitted model corresponds to backscatter strength over large square tiles that do not exactly match that simulated, then although a considerable part of the signal will be matched and removed, narrow hyperbolic features will remain corresponding to the mismatch in grids. The solution is to have a fitting model that has many narrow backscatter cells so that a variety of different width features can be adequately matched. In this case the along-track size of a cell is chosen to match the distance travelled between successive average waveforms i.e. 374m. A wider across-track width can be employed and still give an adequate description of the spatial changes.

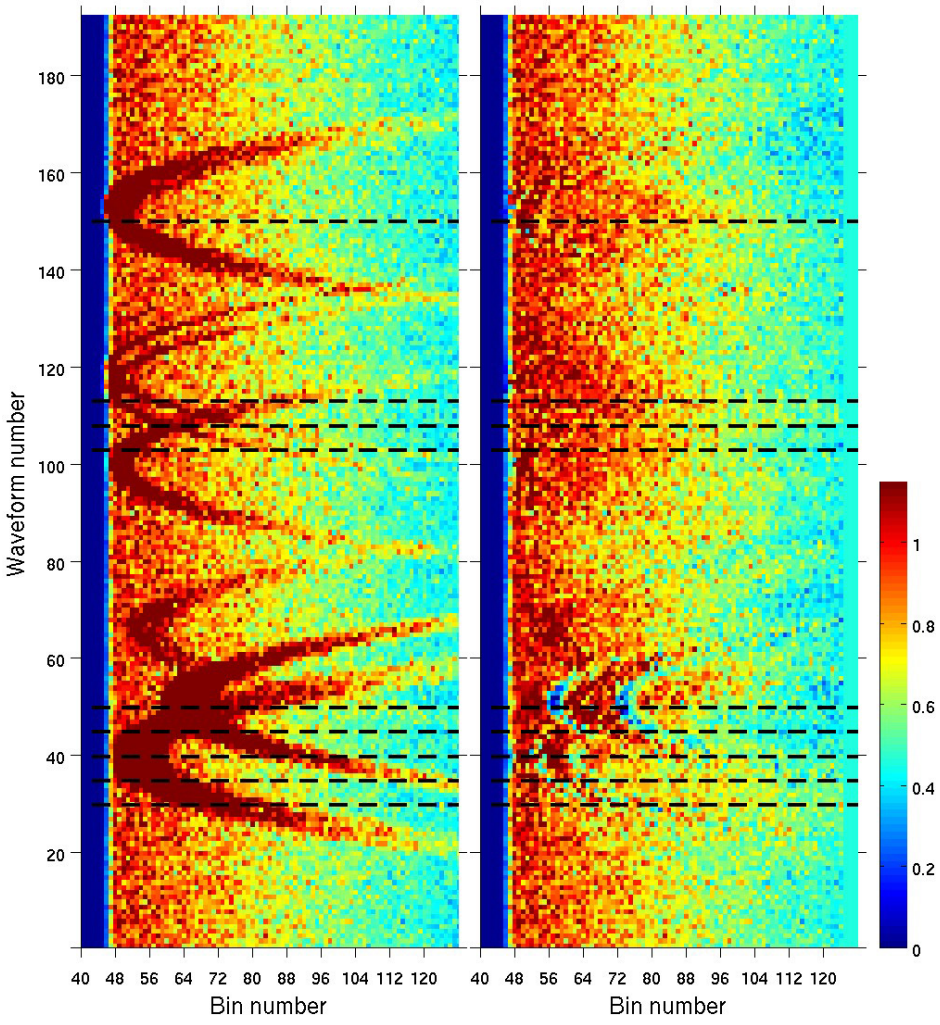

Fig. 7: a) Simulated waveforms for altimeter transit across a surface with a large number of patches of different strength of reflectance, whose shape and structure is not the same as the rectangular patches used in the inversion. Fading noise is equivalent to average of 100 independent pulses. b) Cleaned waveforms — note as the hyperbola centred on waveform 50 was not matched well, there is some vestigial signal. Dashed lines indicate specific waveforms shown in succeeding figure. Power is normalised such that feature-free noise-free signals have a peak of one.

In the second example (Fig. 7), the simulation is of an altimeter transiting across a number of features (narrow/broad, different shapes and both on the sub-satellite track and to one side). The hyperbola-fitting model represents the backscatter strength as an array of rectangular facets 374m along track and ~1km across track. As such a model is broadly able to encompass the range of conditions simulated, the pretracker is again able to clean up the waveforms allowing standard retracking algorithms to be applied to the resultant data.

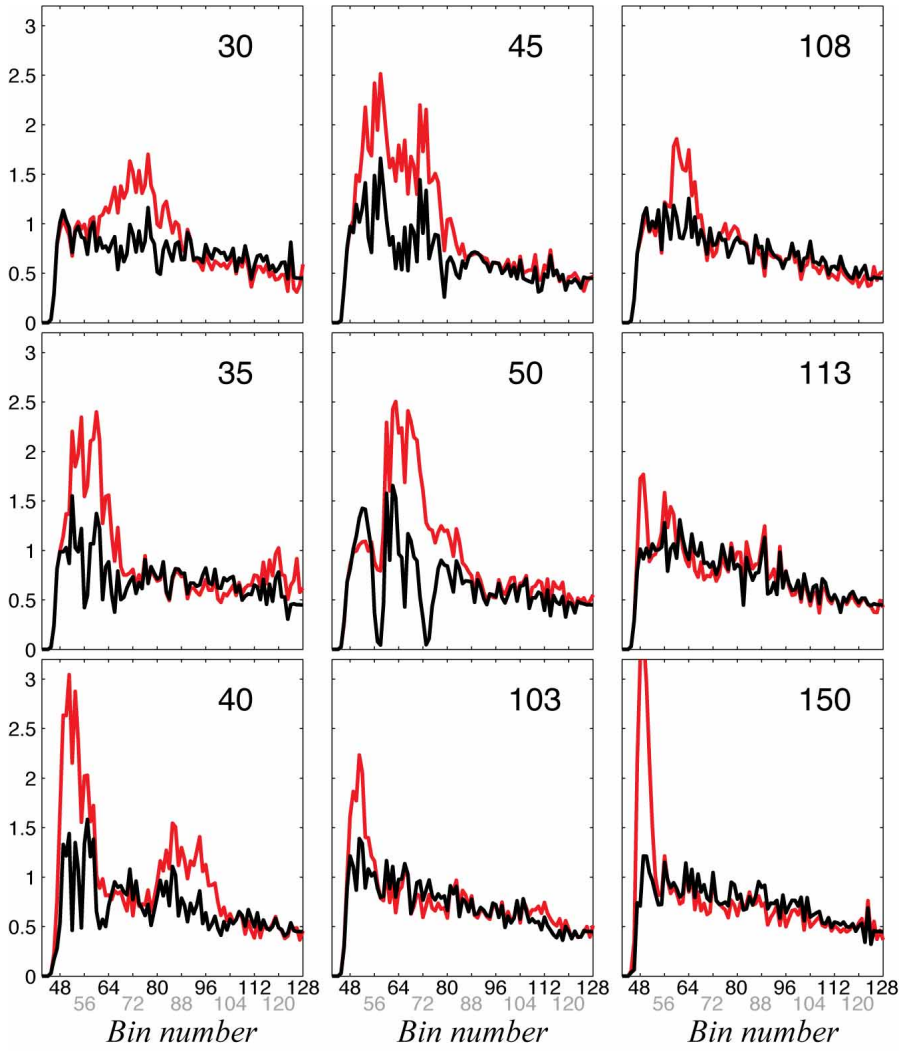

Fig. 8: Individual average waveforms from the complex simulation shown in Fig. 7, with red lines for original version and black lines for cleaned up dataset. In the majority of cases, the cleaned up version looks much more Brown-like and more amenable to a standard retracker, but it is probable that waveform no. 50 would still be problematic, as the trailing edge of the plateau region does not show a steady decline (as is expected for the Brown model). Power is normalised such that feature-free noise-free signals have a peak of one.

3.3. Rescaling and shifting of waveform data

In order to be able to apply the hyperbolic pretracker to real Envisat altimeter waveforms, these need to be forced to match the characteristics of the simulations. In effect this means that changes in waveform scaling due to the AGC setting need to be corrected for, and that the waveform positions should be free from any artefacts due to abrupt movement of the tracker window. The first element is easy to correct for using the high rate (18 Hz) AGC data in the waveform product. Each waveform is simply scaled by 10^(AGC/10). The second element has required much further investigation.

If the altimeter waveform reception window was suddenly, say, 3m to the right of normal (i.e. a later time delay), then the signal would appear 3 / δx earlier, where δx is the bin width (which has the value 0.46875 m). Consequently before applying the waveform analysis detailed herein it is necessary to shift that waveform later by 3 / δx bins. In practice this needs to be rounded to the nearest integer to fit the 2-D matrix formulation. A more challenging issue is determining the actual change in position of the tracker window. (Note, such cases cannot be merely dismissed as a rarity. There is much interest in the use of altimetry in the coastal zone; in such scenarios the window is likely to be migrating markedly in response to a land/ocean transition. Secondly the occurrence of bright targets or rain attenuation cells often leads to sharp movements in the tracker window (as shown for ERS-1, Guymer et al., 1995), and thus the very occasions when a hyperbolic pretracker would be useful are those likely to be exacerbated by movement of the tracker window.)

Initial correspondence with ESA had suggested that the 18Hz values of field 13 ("tracker range referenced to the COG") would suffice (once a minimum or mean value had been removed); however it quickly became apparent that this wasn't valid. The tracker range is often changing markedly in a few seconds as the instrument approaches or recedes from the ocean surface — such could be overcome by removing a linear trend from the 18 Hz range values.

As well as 18 Hz tracker range values, the waveform record also contains 18 Hz estimates of altimeter altitude (relative to ellipsoid) and retracked ocean range (field 19). In practice a reliable estimator of the positional adjustments of the tracker window will be based on the difference of two of these terms. That the correct solution is altitude (fields 9 & 10 of MDSR) minus tracker window (field13) rather than some other combination is demonstrated in the next couple of examples. Both come from an ascending Envisat pass over the North Pacific crossing a narrow isthmus before reaching the Alaskan mainland (Fig. 9). Thus the first case (Fig. 10) is a transition from open ocean (with typical Brown waveforms) to land, whilst the second shows specular echoes from the near glassy surface in the sheltered Cook Inlet (Fig. 11).

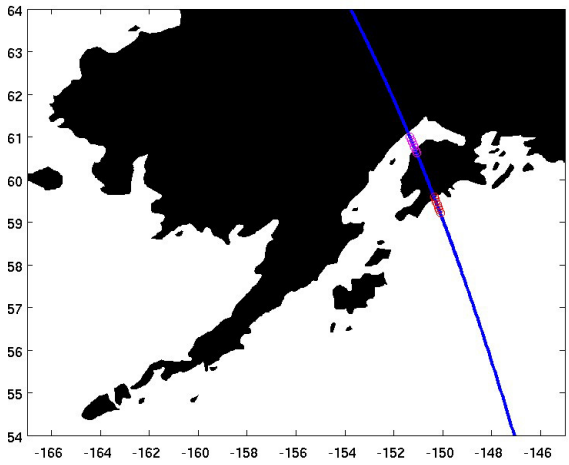

Fig. 9: An ascending Envisat pass over Cook Inlet (near Anchorage, Alaska), with data from the red and magenta parts of the track analysed in Figs. 10 & 11.

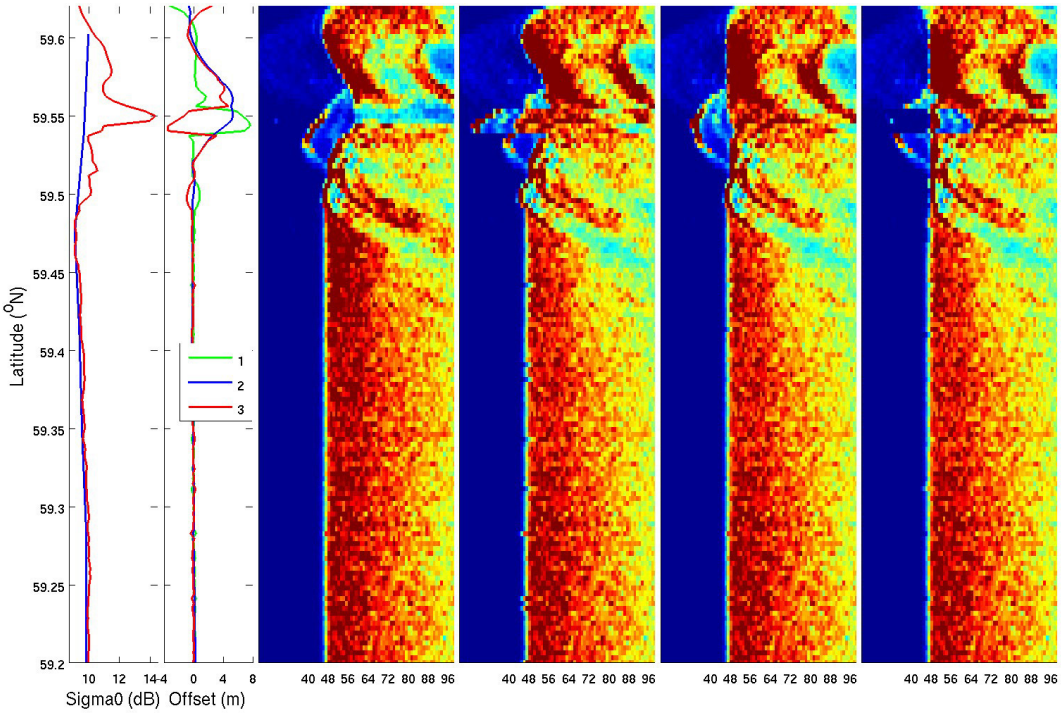

Fig. 10: Section along track in NE Pacific to northeast of Kodiak Island (red segment Fig. 9) as altimeter moves from open ocean to land. a) K_u-band backscatter values (blue is 1 Hz σ^0 values, red is 18 Hz AGC minus 16.6 dB). b) Derived 18 Hz offset values in metres (1 represent "altitude - range - geoid", 2 is "altitude - tracker window - geoid", 3 is "range - tracker window"). c) Original waveform data without AGC scaling or waveform shifts. d), e) and f) show waveform data after AGC rescaling with shifts calculated using offset 1, 2 or 3 respectively. Note the penultimate panel (using offset 2) gives the correct match for the extraneous feature at 59.55˚N

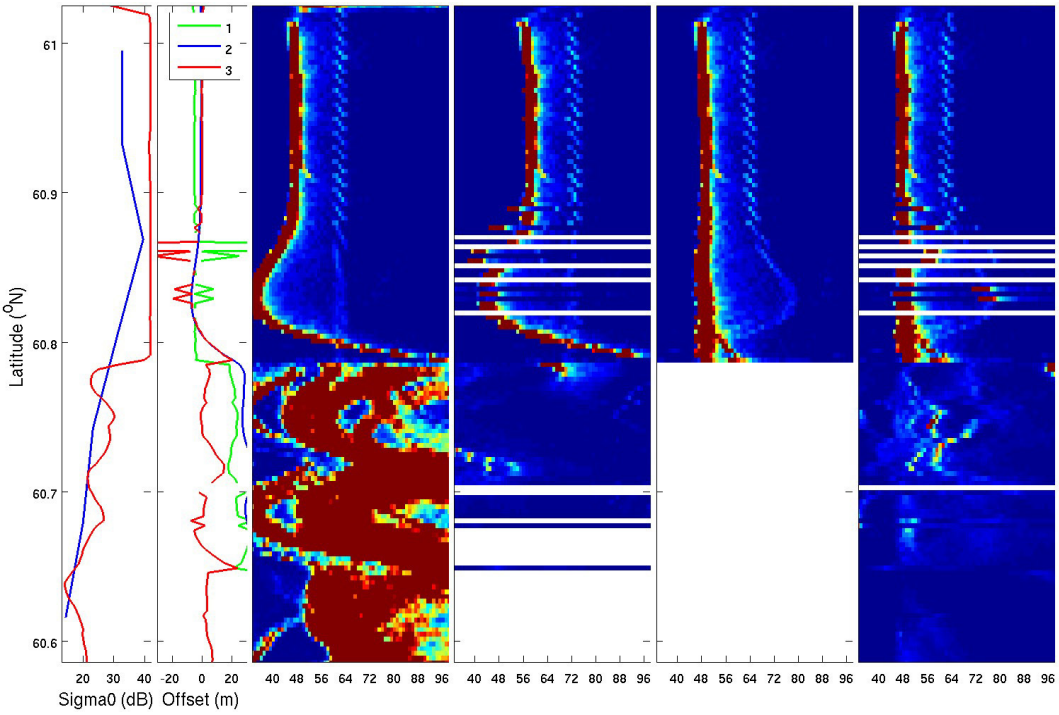

Fig. 11: Section along track in NE Pacific as altimeter moves from land to glassy sea (Cook Inlet, magenta segment in Fig. 9). Panels as for Fig. 10, but note rescaled waveform data are blanked out where implied shift is more than 25 bins.

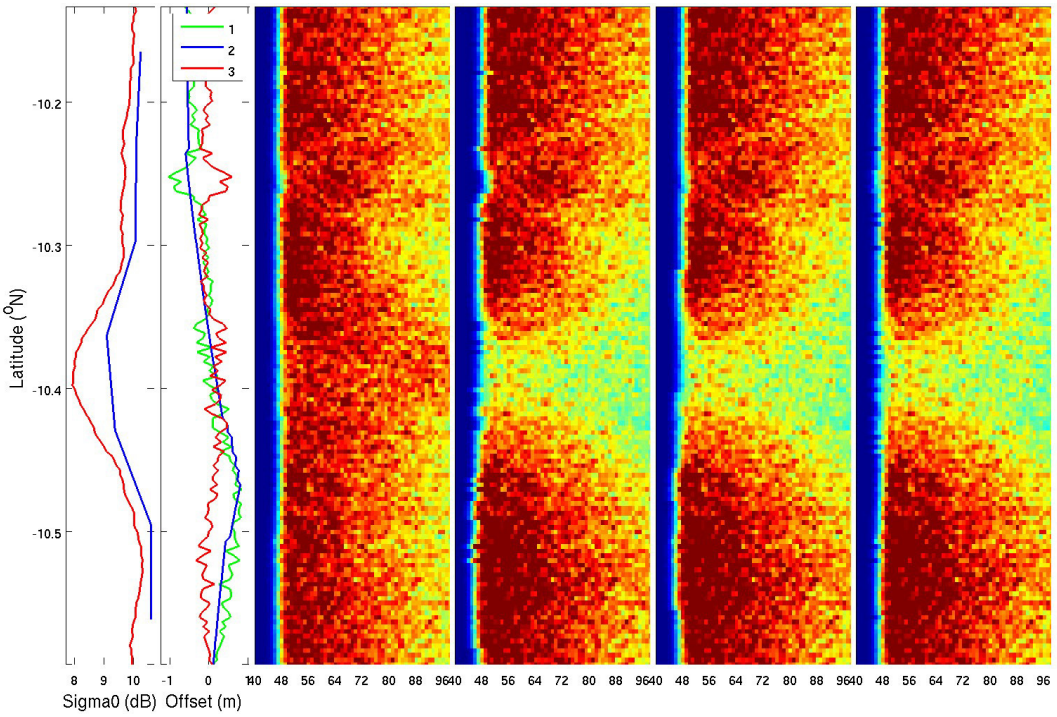

Fig. 12: Example from open ocean (with details of panels as for Fig. 10), where all three estimates of offset lead to significant shift (2 bins) of leading edge away from expected position.

The third example (Fig. 12) shows that "altitude - tracker window - geoid" is not the full solution. As the set of weights developed for the hyperbolic pretracker is predicated on the leading edges of the waveforms being close to 46.5, poor results will be achieved if the waveform repositioning has them two bins earlier. Also it emphasises the peculiar problems when the intended offset has a slow trend: the shifts may show multiple fluctuation once converted to integer values. The information to correctly calculate the tracker position relative to a mean sea surface should be present within the waveform data. In this case, a quick work around is to apply a high-pass filter to the derived offset values before calculating the necessary integer shifts. Here I have calculated a running 41-point median and subtracted this. The final example of waveform adjustment (Fig. 13) shows that the hyperbolic feature associated with a land reflection (earlier than the bins of the leading edge) has the same shape as used in the simulations.

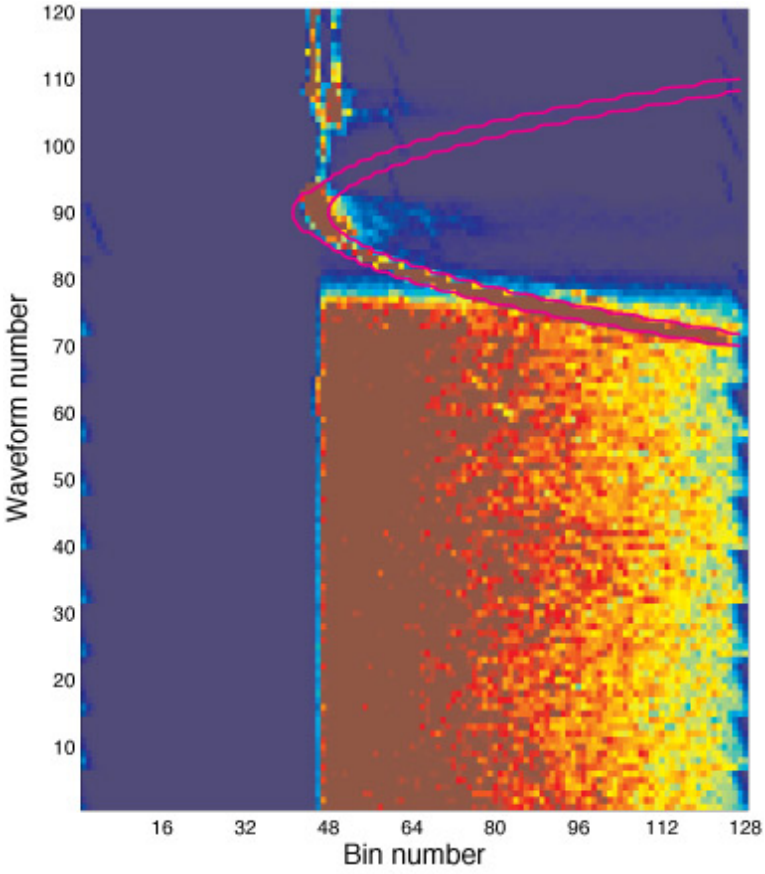

Fig. 13: Shifted waveforms for an ocean-land transition, with pink contours being the shape generated by a separate simulation for a narrow strongly reflecting cell.

3.4. Application to real data

This section shows some results where the hyperbolic pretracker does a good job or removing the unwanted features. Indeed, it can be used to provide an estimate of where the features are and of their intensity, although the emphasis here is on their efficacy at extracting unwanted contaminants from marine altimeter waveform data. It should be noted, that the present implementation does not work well in all cases, with a number of the reasons discussed in section 4.

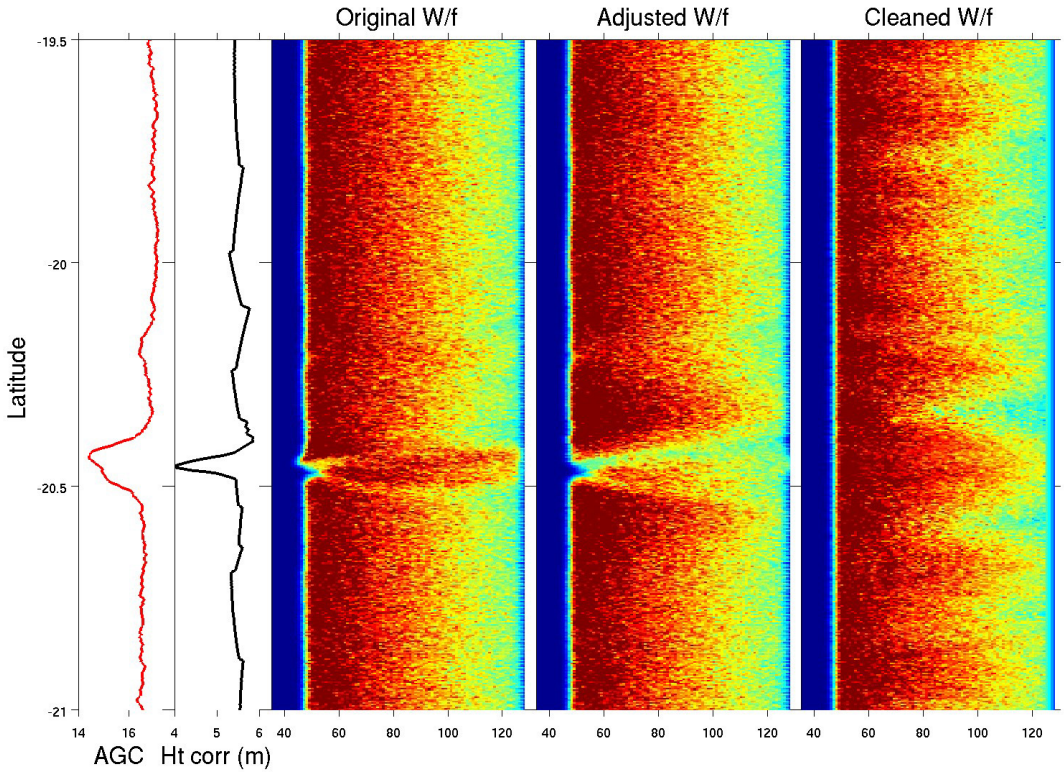

Fig. 14: A simple mid-ocean case with a reduction in AGC (first panel) and movement of the tracker (second panel). The other three panels show original raw waveform data, same after AGC scaling and shifts applied and finally after the hyperbolic pretracker has removed the "weak target" feature. (The value used for ht_corr is altitude - tracker window - geoid.)

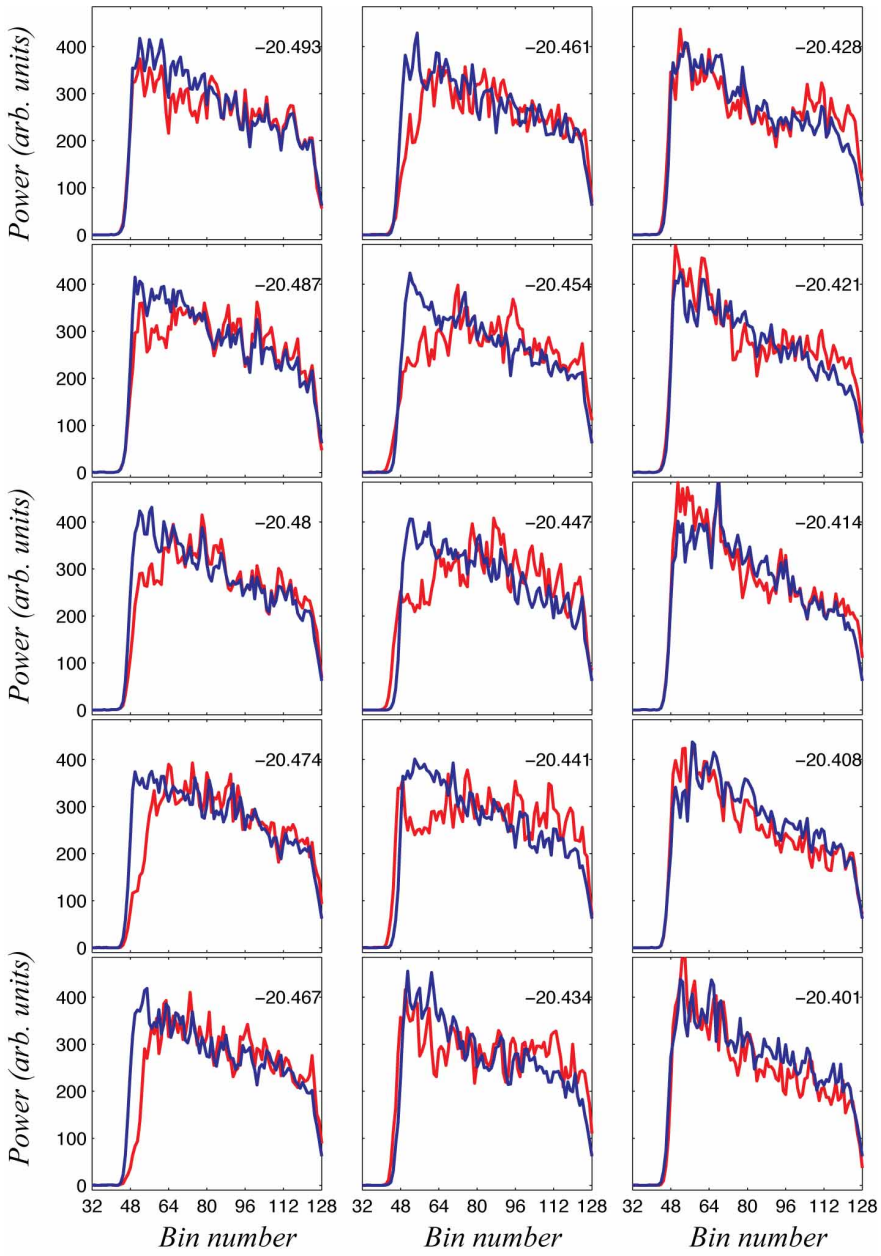

Fig. 15: Portrayal of every second waveform across the weak target shown in Fig. 14. Red line shows original signal and blue that after the pretracker applied to clean it up.

The succeeding examples (Fig. 16-19) are taken from Envisat overpasses of Pianosa Island in the Tuscan Archipelago, which has proven to be of great interest because of its frequent but intermittent occurrence of a bright target just to the north of the island coupled with a permanent weak target feature due to the island itself (Gómez-Enri et al., 2010, Quartly, 2010).

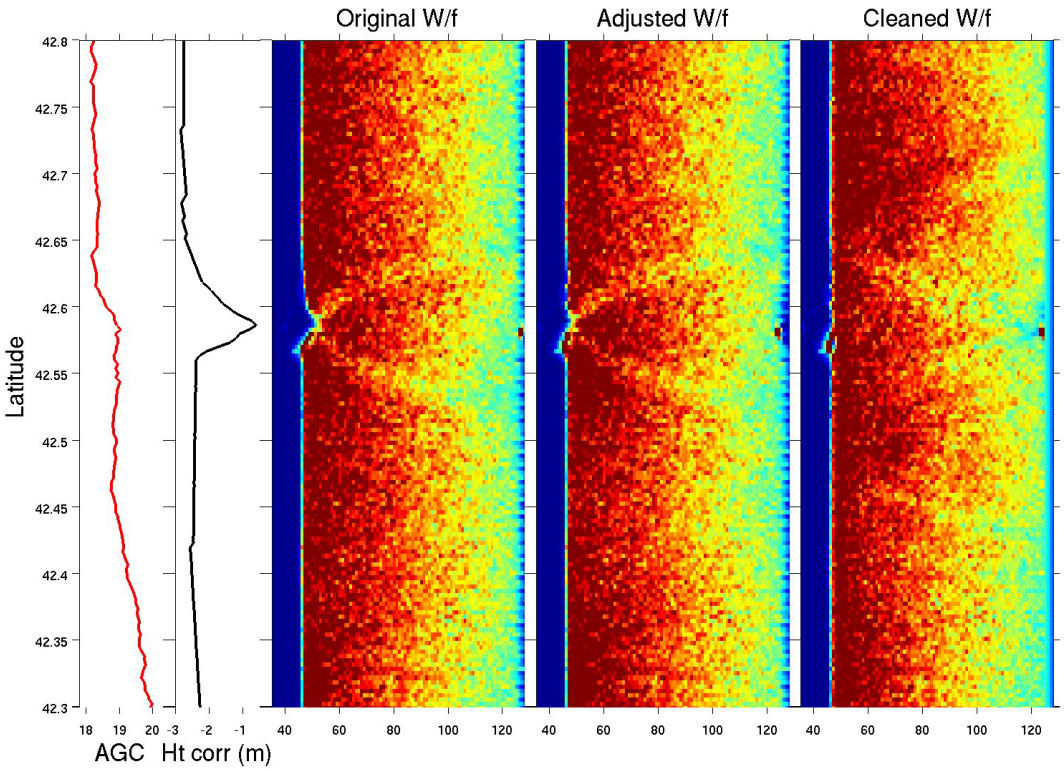

Fig. 16: Another example (case P04) — Envisat pass near Pianosa, cycle 040

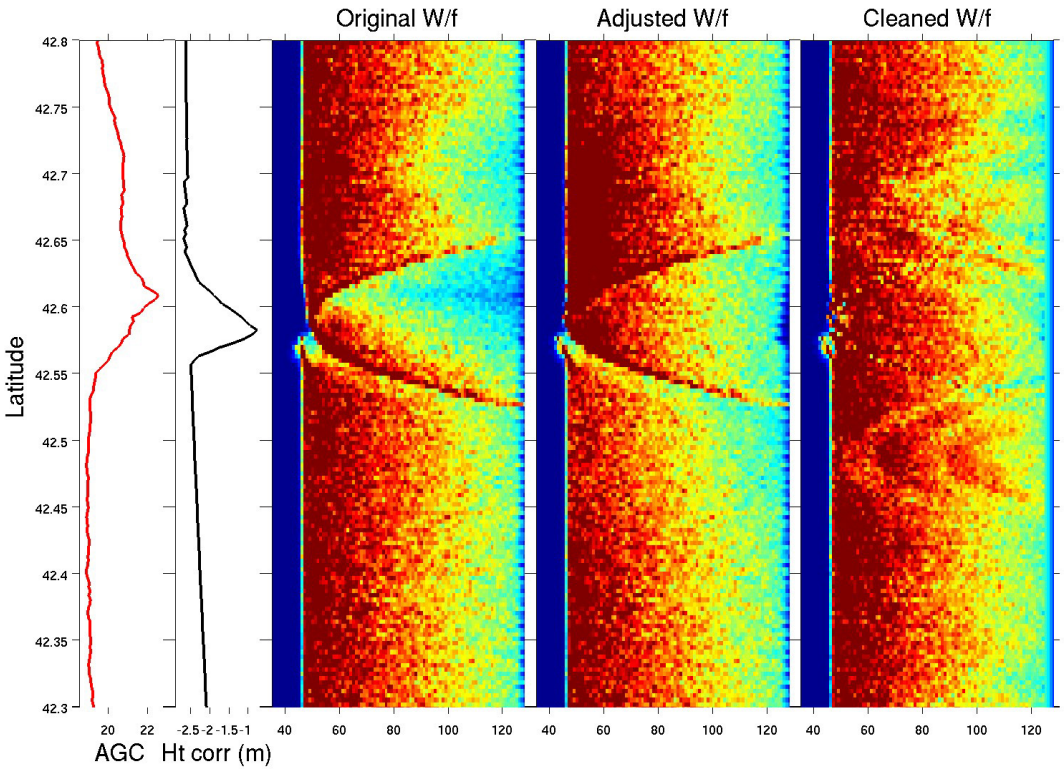

Fig. 17: Another example (case P08) — Envisat pass near Pianosa, cycle 068

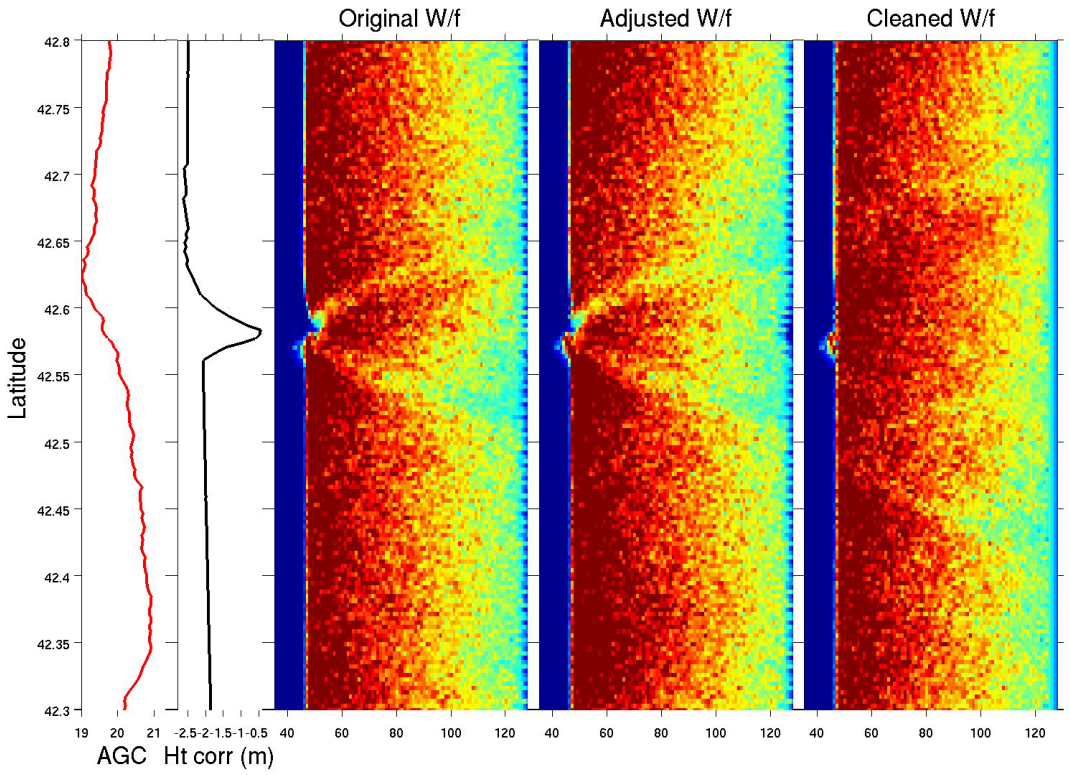

Fig. 18: Another example (case P09) — Envisat pass near Pianosa, cycle 069

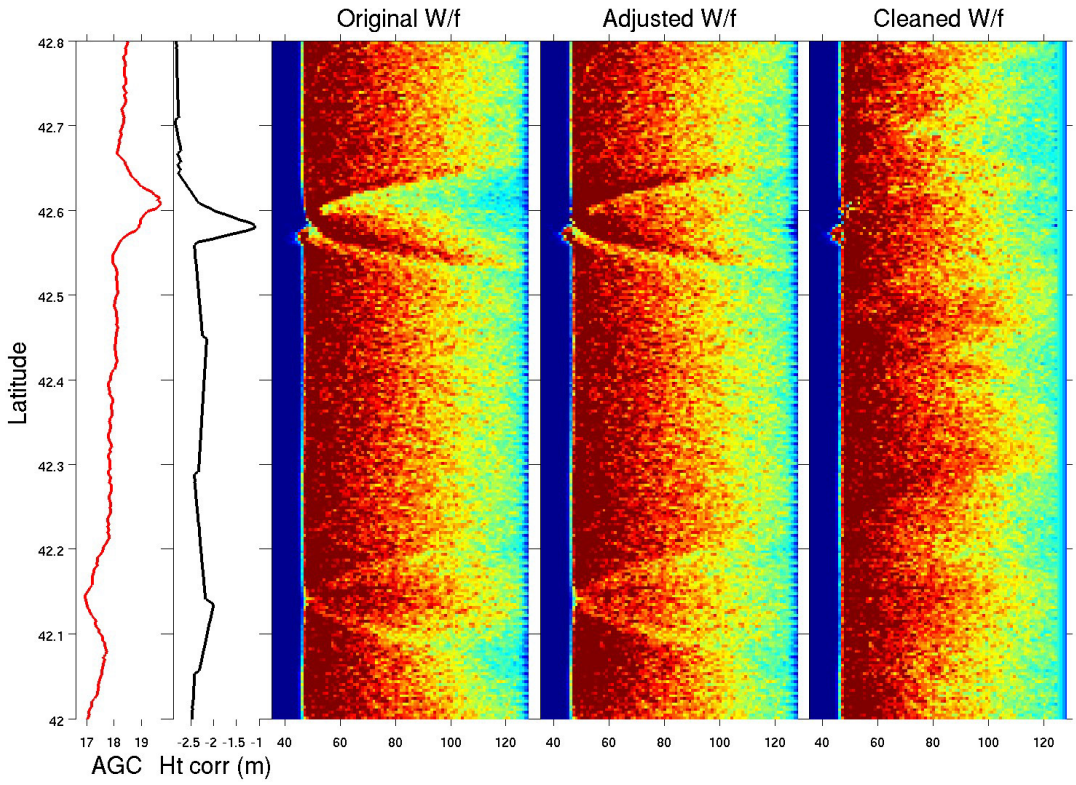

Fig. 19: Another example (case P13) — Envisat pass near Pianosa, cycle 073

4. Concluding remarks

4.1. Summary

The hyperbolic pretracker has been shown to be a robust feature for locating and removing hyperbolic features within the 2-D waveform data, performing the job very reliably with simulated data over complex scenarios with realistic levels of fading noise. The resultant cleaned waveforms (e.g. Fig. 15) are much more amenable to a conventional waveform retracker. The use of this technique as a pre-processor means that it can be implemented along with other waveform-fitting techniques (independent of whether they are fitting for skewness, mispointing etc.). However, the technique, in its present form, does not restore absolutely all input data into Brown-like waveforms, for example the complex input waveform no. 50 in Fig. 8 is improved by the cleaning operation (producing the black curve), but that output is still so far from that expected from scatter from a homogeneous surface, as to be unsuitable for retracking.

For use with real altimeter data, the waveforms need a degree of "massaging" — correcting for changing AGC and movement of tracker window; this has not been as simple as initially believed. And indeed a number of problems persist when there are small movements of tracker leading to alternately positive and negative anomalies at the leading edge. Despite this, the hyperbolic pretracker is able to work with many examples of real waveform data containing bright target features. A number of the concerns for a full implementation are raised in the caveats section (4.3).

4.2. MATLAB implementation & worked examples

The code to implement the pretracker is provided. The overall calls are to i) **HYP_wform_adj.m**, which applies AGC and tracker position corrections to the waveform data, and ii) **HYP_surgery.m**, which converts the waveform data into waveform anomalies (via **HYP_anom_calc.m**), then reads in a set of reference hyperbolae and weights from **HYP_data0.mat**, and applies these in turn to estimate intensity of hyperbolic features for each along- and across-track position and remove an appropriately-scaled version from the contaminated waveform anomaly data. Finally that program converts the cleaned up anomalies back into actual waveforms for use with a standard ocean-tracker algorithm. Specimen case studies are included **HYP_CaseP04.mat** etc., which contain waveform, AGC and tracker position data at 18 Hz for various short segments. Implementation:

```
rescaled_wform = HYP_wform_adj(wf,agc18,alt18-track_cog18-
geoid18);
cleaned_wform = HYP_surgery(rescaled_wf);
```

The MATLAB code and the examples shown are available on the COASTALT FTP site under `/utilities/hyperbolic_pretracker`. The data file contains a series of reference hyperbolae, given by their anomaly in 2-D waveform (i,j) space for pixel index r away from nadir (see text prior to Eq. 1 for full explanation of i, j & r). The weights to estimate each of these hyperbolae are of the same dimensions. (Note there is no need for reference hyperbolae or weights for different spatial pixels along track, as this is accomplished by shifting the series of waveform anomalies being processed.) The execution time for rescaling is 0.12 ms per waveform and the "cleaning" averages 2.47 ms per waveform on a Linux box with Intel Xeon X5660 2.8 GHz chips, running MATLAB 2009a (although of course the operation is only meaningful on a pass of many waveforms all together). These times are only indicative, and will be bettered when i) all weights and reference hyperbolae are kept in memory rather than read in during each call, and ii) code optimised to run in a compiled language such as C.

4.3. Caveats

Although it might seem inconsequential, it should be pointed out that the pre-tracker works very well with the simulation data. A number of cases are included of the successful application of the pretracker to real RA2 altimeter waveforms; however, in the present form it does not work well with all samples of real data. Much of this is due to the practicalities of arranging the received waveforms exactly as they should be for maintaining the tracker window such that the half-power point of perfect waveforms would be at exactly the expected location. In particular:

a) We have not managed to get a precise formula for where the tracker window is. The use of altitude - (tracker_cog + geoid) gives something slowly varying for benign conditions, but still contains an offset. Possibly additional terms of mean sea surface, expected tides, tropospheric and ionospheric corrections would give a near-zero correction for feature-free waveform data. Instead here, we have extracted the short-term variability by differencing with respect to a 41-point running median.

b) In order to calculate anomalies a reference waveform is required, for which the median of the ensemble is used, as this should match the typical shape for feature-free conditions. However, this assumption is invalid if the data segment is dominated by contaminated waveforms. Alternatively if very large segments of data are used, a reference calculated over the whole will be for a different wave height (and thus slope of leading edge) than is appropriate for the location of the hyperbolic features. Thus changes in wave height from the ensemble median will generate spurious waveform anomalies that may remove non-existent features. If the pretracker is

implemented within the full processor, ancillary information such as an initial wave height estimate may be used to define a local reference waveform that is pertinent to the situation.

c) Even on benign stretches, with no apparent features, the waveform anomalies involved in the calculation can show large positive and negative signals at the leading edge, persisting for 20 or 30 consecutive waveforms. This is presumably due to small variations in the position of the waveform window, causing large relative changes in the power for individual bins at the leading edge. (The large anomaly signal generated by an offset of 0.25 bins is illustrated in Fig. 5). To avoid these significant leading edge power anomalies, we need either effectively to shift the waveforms by fractions of a binwidth (involving interpolation) or develop a system of weights that does not use the leading edge bins in its estimation. However, the leading edge is the region of strongest signal in the waveform data and thus should enable the most robust estimations.

d) The systems of weights implemented was developed using simulations with a significant wave height of 1m. This is a good initial choice as most of the bright target signals will be associated with very calm conditions. However anomalous patches (enhancements or diminutions) can occur for passes across seas with greater wave heights. Ideally a slightly different set of weights and reference hyperbolic anomalies should be used for greater wave heights. Firstly the larger wave height spreads the effect of a modified patch over a greater number of bins, and secondly the position of the leading edge is different by up to one bin or more, implying that different bins should be showing the anomaly signal.

e) The pretracker approach is based on the idea that these hyperbolic features will occur on top of otherwise ordinary marine waveforms, both before and after the feature. However this 2-D approach, which looks at the waveforms in context does not easily cope with missing data e.g. if the reflections over land cause a change in tracker mode so that no meaningful 320 MHz waveform data bracket the feature — the exact position and strength of the feature is then hard to estimate for removal.

f) In practice many of these bright targets correspond to small patches with backscatter strengths 5-10 dB or more above those nearby. These give very pronounced effects on the waveforms, and are reliably extracted. An interesting side issue has been the faint hyperbolae ahead of the leading edge due to land, ships etc. (see Tournadre, 2007; Quartly, 2010). These are not dealt with here, because although interesting and intriguing, the reflected signal is nearly always so weak that it would not adversely affect the ability to fit a standard Brown model to the data.

4.4. Acknowledgements

Thanks to Salvatore Dinardo for advice on tracker movement and pointing out the *circshift* function, to Helen Snaith for further discussion, and to Yiannis Andrianakis for guidance on minimization routines in MATLAB.

5. References

Brown, G. S. (1977) The average impulse response of a rough surface and its applications. *IEEE Trans. Antennas Propag*., **AP-25**, 67–74.

Gómez-Enrí, J., S. Vignudelli, G.D. Quartly, C.P. Gommenginger, P. Cipollini, P.G. Challenor and J. Benveniste (2010) Modeling Envisat RA-2 waveforms in the coastal zone: Case-study of calm water contamination, *IEEE Geosci. Rem. Sensing Lett*., **7 (3)**, 474-478 doi: 10.1109/LGRS.2009.2039193

Guymer T.H., G.D. Quartly and M.A. Srokosz (1995) The effects of rain on ERS-1 radar altimeter data, *J. Atmos. Oceanic Tech*. **12**, 1229-1247

Hayne, G. S. (1980) Radar altimeter mean return waveforms from near-normal-incidence ocean scattering. *IEEE Trans. Antennas Propag*., **AP-28**, 687–692.

Quartly G.D. (1998) Determination of oceanic rain rate and rain cell structure from altimeter waveform data. Part I: Theory, *J. Atmos. Oceanic Tech*. **15**, 1361-1378.

Quartly, G.D. (2010) Hyperbolic retracker: Removing bright target artefacts from altimetric waveform data. in, *ESA Living Planet Symposium*, Bergen, 28th Jun-2nd July 2010. Noordwijkerhout, NL, European Space Agency **ESA SP-686**. (6pp).

Quartly G.D., M.A. Srokosz and T.H. Guymer (1999) Understanding the effects of rain on radar altimeter waveforms, *Adv. in Space Res*. **22**, 1567-1570.

Tournadre, J. (2007) Signature of lighthouses, ships, and small islands in altimeter waveforms, *J. Atmos. Oceanic Tech*. **24**, 1143-1149.