

Validation of the ERS-2 Scatterometer Ground Processor Upgrade

Pauline Pettiaux, Xavier Neyt and Marc Acheroy
Royal Military Academy, Brussels, Belgium

ABSTRACT

Due to gyroscopes malfunctions, the ERS-2 spacecraft cannot accurately be yaw steered. Moreover, the actual yaw angle of the spacecraft is unknown. Since the yaw angle is not known in advance and not periodic, the look-up tables-based original scatterometer processor is not able to compute accurate values for the backscattering coefficients from the measurements made. This implies the need for an upgraded wind scatterometer ground processor in order to obtain accurate backscattering coefficients. Moreover, the upgraded processor includes several other enhancements. This paper presents the results of the validation of the upgraded processor. The validation of the geometrical model is performed by comparing geometric parameters such as the incidence angle and the sub-satellite track heading. The radiometric performance of the upgraded processor is first evaluated with data acquired in nominal attitude, by comparing the backscattering coefficients. The radiometric performance of the upgraded processor is then further assessed with data acquired in degraded attitude. This is only possible over the calibration test site and over other selected land areas.

Keywords: Scatterometry, attitude estimation, validation

1. INTRODUCTION

Due to gyroscopes malfunctions, a new Attitude and Orbit Control System (AOCS) was loaded onto the ERS-2 spacecraft begin January 2000. This new AOCS introduces a new piloting mode, termed Zero-Gyro Mode (ZGM) where the attitude of the spacecraft is acquired using the on-board Digital Earth Sensor and Digital Sun Sensor. The pitch and roll angle are close to their nominal value but the yaw angle of the satellite is unknown. In practice, it can be kept within $\pm 2^\circ$ around its nominal zero Doppler attitude (also called Yaw Steering Mode — YSM) and the rate of change of the yaw angle is small.

The non-nominal yaw angle value introduces a Doppler frequency shift on the signal received by the antennas. The existing on-board and on-ground compensations are computed for a nominal yaw angle and are unable to correct the shift. This causes a significant part of the signal energy to remain outside the system pass-band. This had a non negligible impact on the output product, namely the backscattering coefficients and was the main driver for a review of the ground processing chain.

The results of the review are described in another article¹ to be found in these proceedings. While some relatively minor parts remained untouched, other areas were thoroughly modified. The changes were not only limited to blocks directly affected by the spacecraft's attitude change, but affected also other blocks where new signal processing methods were applied. Moreover, new features were introduced such as the enhanced σ^0 spatial resolution. Hence, the ground processing was totally rewritten which prompted for a validation of this new processing chain. The validation was conducted using the existing processor as reference in the cases where it still provided accurate results, and using on-ground reference targets in the other cases.

This paper presents the validation of the upgraded processor and is organized in two main sections. Section 2 reports about the validation of the geometric quantities, such as the estimation of the yaw angle, the position of the nodes on ground, etc. Section 3 deals with the validation of the radiometric quantities. The performance of the high-resolution mode are also reported in that section.

Further author information: (Send correspondence to Pauline Pettiaux) E-mail: Pauline.Pettiaux@rma.ac.be, Telephone: +32 2 737 6474, Address: Signal and Image Center, Electrical Engineering Dept, Royal Military Academy, 30 av de la Renaissance, B-1000 Brussels, Belgium

2. GEOMETRICAL VALIDATION

2.1. Satellite simulator

The existing Scatterometer ground processor relies heavily on pre-computed tables that were established for a spacecraft in perfect yaw steering mode. These tables contained for instance the along-track and across-track phasor used to perform the Doppler frequency shift correction. Due to the geometrical symmetry of the system, most of the tables were established for one orbit. Anyway, since the yaw is not periodic, the mechanism of look-up tables can not be used at all anymore.

In the new processor, all the geometric quantities such as target location, velocity, incidence angle, ... are computed on the fly. The computations are driven by the acquisition time of each data packet. An orbit propagator is used to compute the true position and velocity of the spacecraft at that time. Various geometric relations are then used to compute the other geometric quantities. This part of the processor was validated through the values of particular parameters output in the products. This validation is described below.

2.2. Yaw angle estimation

The yaw angle is obtained from the estimation of the residual Doppler frequency shift, as explained in.¹ This estimation is conducted on the signals received by the 3 antennae. Since the pulse length emitted by the mid-beam antenna is much smaller than that of the pulses emitted by the two other antennae, the spectrum of the corresponding received echo is much wider. This causes that spectrum to be shifted outside the bandwidth of the on-board anti-aliasing filter and the spectrum of the mid-beam data received on ground is irremediably distorted. This renders the residual Doppler frequency shift performed on the mid-beam signal less accurate, even at moderate yaw angles. For this reasons, only the estimated residual Doppler frequency shift performed on the signals received from the fore and aft antennae are considered. These two estimates are averaged together and along-track to yield the final yaw angle estimate.

The results of the estimation of the yaw angles are presented in figure 1 and 2. The yaw angle in yaw steering mode

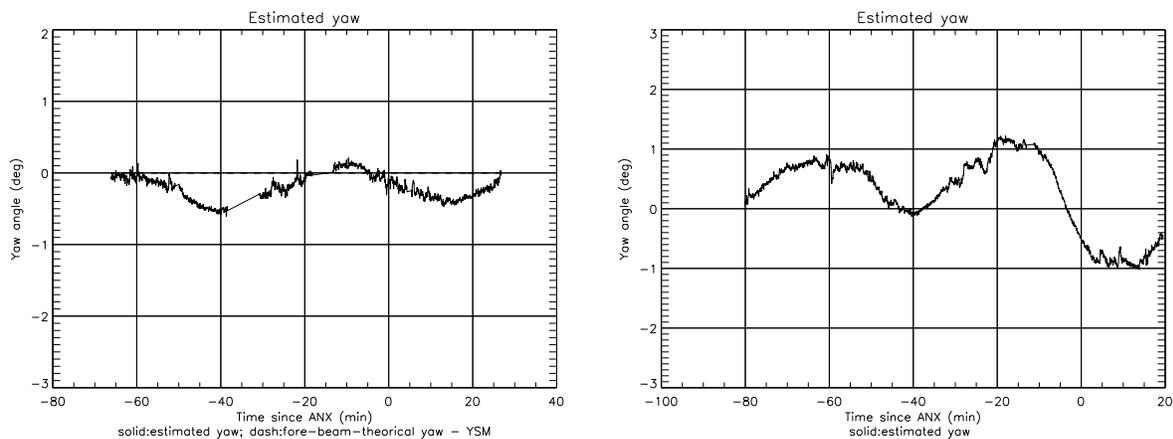


Figure 1. Along-track evolution of the yaw angle estimation for one orbit of data acquired when the spacecraft was in resp. yaw steering mode (left) and in zero-gyro mode (right). In yaw steering mode, the yaw angle should be zero while it can be arbitrary in ZGM.

should be zero. The estimation in fig. 1 shows that the estimated angle oscillates between 0° and -0.5° around the theoretical perfect YSM attitude. This behaviour is general for data acquired in YSM. Investigations are still on-going to determine whether the cause is related to an imprecision in the actual yaw steering mode of the spacecraft or due to a model error in the satellite simulator, and particularly in the expression relating the estimated Doppler frequency shift to the yaw angle. The right part of figure 1 shows the estimation of the yaw angle for data acquired when the spacecraft was in zero-gyro mode. As can be seen, the yaw angle exhibits rather an erratic evolution. In particular, there is no periodicity and the yaw angle cannot reliably be predicted for future orbits.

Figure 2 shows the yaw estimation and the estimation error for data acquired when the spacecraft was in fine pointing mode (FPM). In FPM attitude, the roll axis of the satellite is located in the plane containing the absolute velocity vector and the earth center. The corresponding theoretical yaw angle can thus be computed from the satellite simulator and

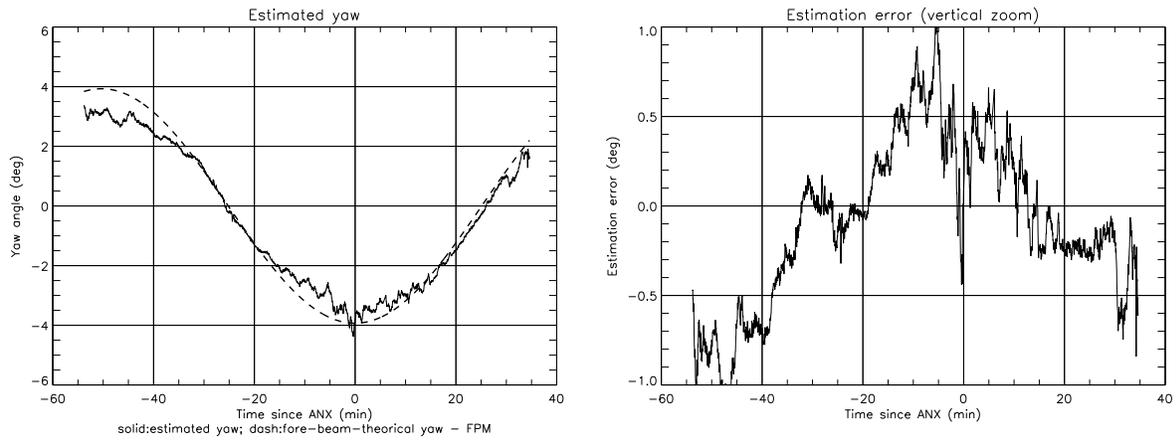


Figure 2. Along-track evolution of the yaw angle estimation for one orbit of data acquired when the spacecraft was in fine pointing mode (FPM). The dashed line in the left graphs shows the theoretical yaw angle, while the right graphs shows the estimation error (difference between the actual estimation and the theoretical prediction).

varies from -4° to $+4^\circ$, the maximas being reached at the equator. The difference graph shows that the estimation is quite accurate at low yaw angles but becomes less performant for large yaw angle, where the yaw angle is underestimated. Since the backscattered echo signal received by the mid-beam antenna has a much wider spectrum than that of the for/aft-beam antennae, the residual Doppler shift caused by moderated yaw angles causes the signal spectrum to be shifted outside the bandwidth of the on-board anti-aliasing filter. In practice, this limits the yaw angles for which accurate mid-beam measurements can be made to $\pm 2^\circ$. From a practical point of view, the inaccuracy yaw angle estimation at large yaw angle is thus less an issue.

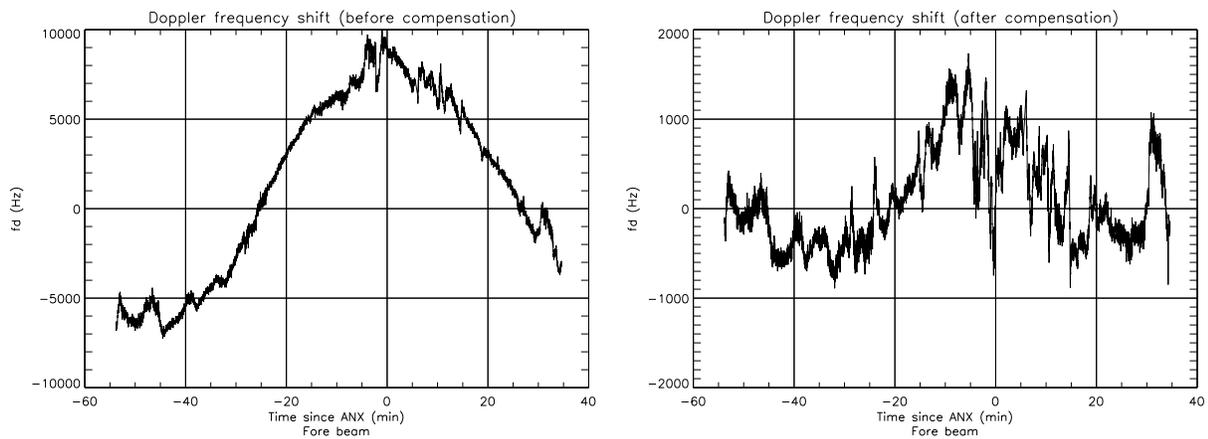


Figure 3. Illustration of the Doppler frequency shift correction on data acquired when the satellite was in fine pointing mode and received by the fore antenna. The along-track evolution of the mean signal spectrum center before and after residual Doppler frequency shift correction is shown respectively on the left and right graphs.

Once the yaw angle is obtained, it is used to compute the across-track evolution of the residual Doppler frequency shift. This result is then used to correct that frequency shift in order to center the received signal inside the pass-band of the subsequent low-pass filter. Figure 3 illustrates this process and shows that the correction is quite effective, even for very large shifts. It should be noticed that the signals are upsampled before undergoing the Doppler frequency shift correction in order to limit the effect of spectral aliasing.

The yaw angle is also used to determine the actual on-ground position of the raw measurement samples. Since the yaw angle can take arbitrary values, the appartenance of a sample to a node cannot be determined based on the across and

along-track time anymore. An individual computation of the relative position of the sample w.r.t. the considered node is required. This computation is performed in the plane tangent to the earth surface at the node location. The position of each sample is projected onto that plane. The sample is considered for the computation of the per node average if it fits in a rectangular window having one size parallel to the relative ground track of the spacecraft. The intrinsic consequence of the fact that the yaw angle can take arbitrary values is that the number of measurement samples contributing to nodes can also take arbitrary values, even if the considered nodes are located at the same range.

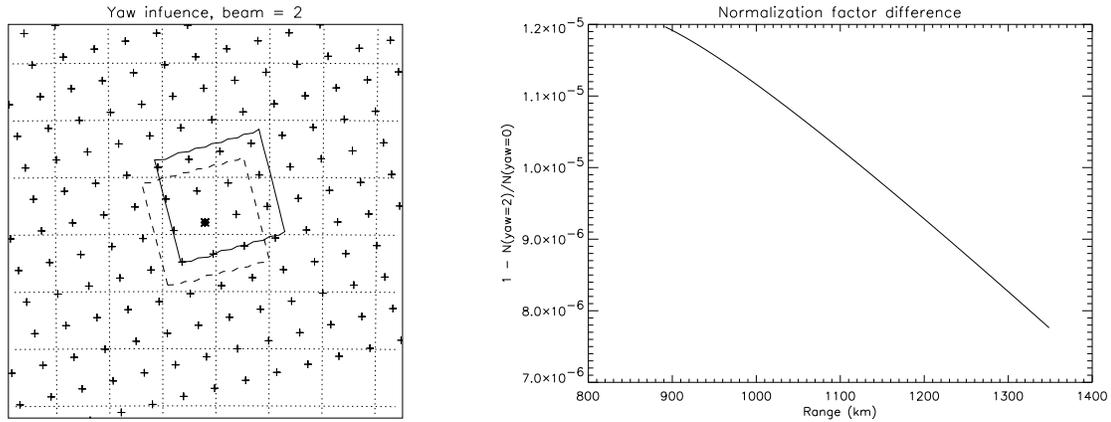


Figure 4. Effect of an error on the yaw angle of 2° . Left: geometrical effect. The nodes where σ^0 values are computed are indicated by crosses. The dashed-line polygon delimits the areas where measurement samples will contribute to the node of interest (diamond+cross in the center of the polygon) in the case no error is committed on the yaw angle. The solid-line polygon delimits the area where measurement samples will contribute to the node of interest in case the yaw is mis-estimated by 2° . Right: across-track relative difference in normalization factor if the yaw angle is mis-estimated by 2° .

Figure 4 shows the effect of an error in the assumed yaw angle. The left graph shows the location of the actual samples that would contribute to a node if the yaw angle is mis-estimated by 2° . The effect of such an error is that the processor is led to believe the samples are in the dashed-line polygon while their actual location is inside the solid-line polygon. The global net result is a displacement of the σ^0 map. The right graph shows the radiometric influence of the same error. The normalization factor is the factor used to correct for the range, elevation antenna gain and incidence angle influence on the value of the backscattered energy. The figure shows that the radiometric influence is negligible.

2.3. Node-related angles

Since the geometry was determined once and for all, the old processor used to precompute the node-related angles such as the incidence angle at a node and the node look-angle (angle between the local meridian and the pointing direction of the considered antenna). Actually, these angles depend explicitly on the yaw attitude of the spacecraft at the time the measurement samples contributing to the considered node were acquired. This means that the corresponding angle can only be computed once the measurement samples contributing to the node are known. Moreover, assigning a yaw angle to the center of a node would be quite complex given the fact that in practice, the center of a node rarely corresponds to an actual measurement.

The convention that was used to assign a value to the node-related angle was to use the yaw angle of the closest measurement sample. This is a pragmatic and very effective solution but causes some artefacts around data gaps as can be seen in fig. 5. These gaps are caused by either an instrument shutdown or a switch between functional modes. They are characterized by the absence of measurements. The nodes that are within a gap are totally “empty” and have an undefined σ^0 value. When the node is located at the boundary of a gap, it can happen that it is filled with less than half the nominal number of measurement samples. In that case, the measurement sample closest to the center of the node can be located quite far from the center of the node, and hence have quite different angular values. The consequences are however very limited since the σ^0 value is in these cases probably unusable anyway.

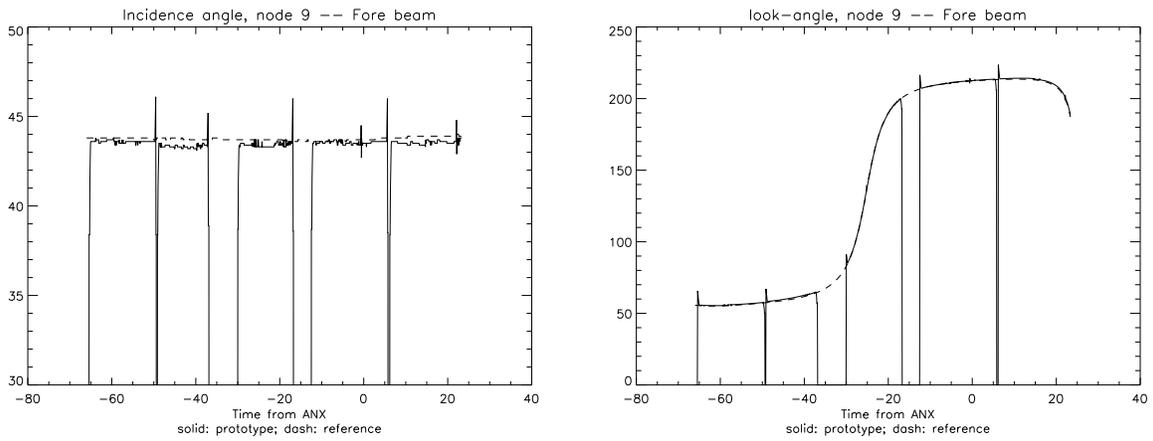


Figure 5. Incidence angle and look-angle at mid-swath. The solid line is the value obtained with the new processor while the dashed line is the value given by the old one.

2.4. Node position

The σ^0 values are given at nodes regularly spaced along the orbit and in the satellite's swath. The across-track position of the mid-swath node is given by a constant elevation angle. The position of the mid-swath node relative to the ground track thus varies along the orbit, with the height of the spacecraft. The other nodes of a row are placed at the intersection of a plane going through the spacecraft, the mid swath node and the earth center. The nodes are placed at a constant angular distance from the mid-swath node. The angle is measured between the vectors from the earth center to the mid-swath node and from the earth center to the considered node. The angle is computed such that the distance between two across-track nodes at mid-swath would be equal to 25km* if the earth was locally spheric.

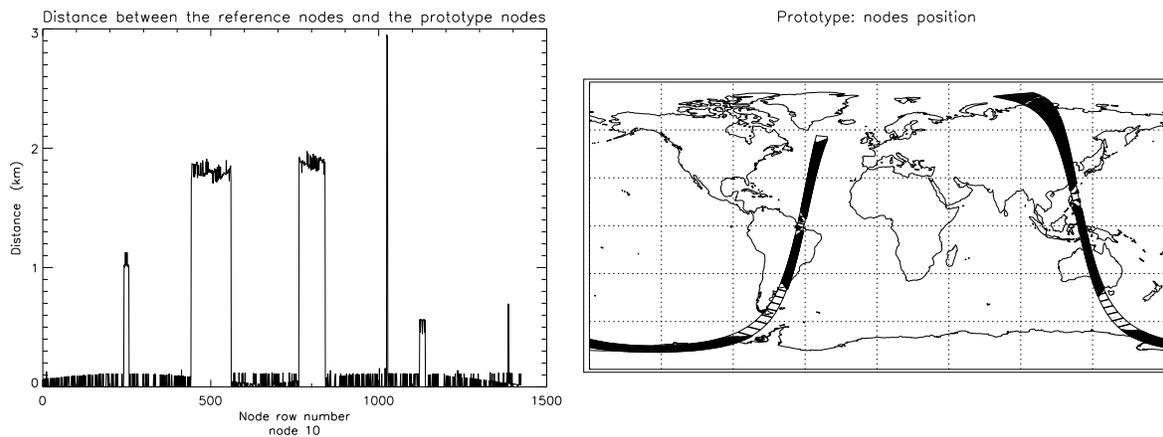


Figure 6. Left: great-circle distance between the mid-swath nodes generated by the old processor and the corresponding nodes generated by the new processor. The right graph shows the data gaps. The projected rectangles represent the products along the considered swath and the black dots inside the rectangles indicate valid nodes (nodes for which there were actually enough measurements).

The along-track spacing of the node rows is determined by the actual time of the measurements. More precisely, a node row is placed every fourth mid-beam measurement. The node rows are placed perpendicular to the relative subsatellite track. The along-track spacing between the node thus also varies along-track and is roughly equal to 25km. When there is no mid beam (inside a data gap) to use as reference, the node rows are placed at time intervals corresponding to 4 standard mid-beam sequences and are resynchronised on the available mid-beams when the instruments is switched back

*12.5km in the case of the enhanced resolution product

on. Figure 6 shows the distance between corresponding nodes generated respectively by the old and the new scatterometer processor. The relatively large differences occur inside data gaps, where there is no mid-beam to synchronise the position of the node row. The differences were found to be due to the fact that while the new processor computes the position of the node rows inside a data gap from the position of the last node row successfully synchronised, the old processor computes the position of the node rows inside a gap from the position of the first successfully synchronised node row after the end of the data gap. Since the length of a data gap is not necessarily equal to an entire multiple of the per-beam pulse repetition interval, both methods are not equivalent. This is illustrated in fig. 6, where the position of the large distance difference is clearly seen to correspond to “empty” products. It should be noticed that wind extraction at the concerned nodes is probably highly unprecise since the mid-beam measurement is at least partly missing. Besides the missing data, there is no other impact on the σ^0 values.

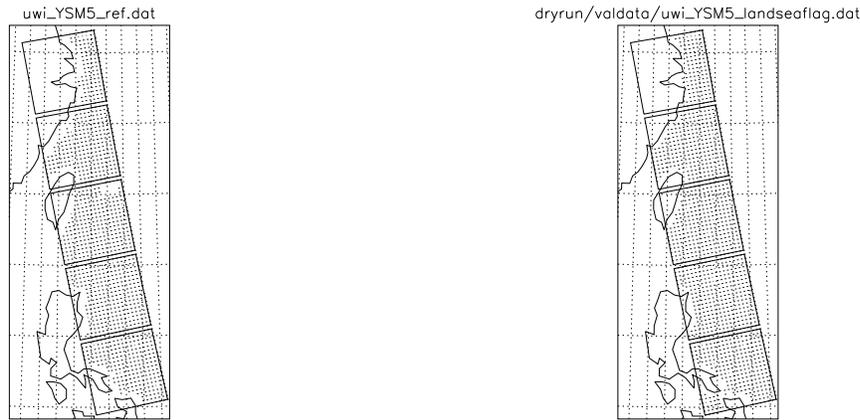


Figure 7. Comparison of the land/sea status of nodes. The rectangles indicate product boundaries and the black dots indicate that the node was considered over sea. Nodes considered over land are not plotted at all. The left map is obtained from products generated by the old processor while the right graph is obtained from products produced by the new processor.

The scatterometer was initially essentially designed to measure sea-surface winds. Obviously, the computation of wind vector from the σ^0 values can only be done if the measurements were actually done over sea. As can be seen from fig. 7, the old processor was exagerately pessimistic. The new processor is flexible in that respect: a node is flagged as being over sea if the relative number of contributing individual samples over land does not exceed a configurable threshold.

3. RADIOMETRIC VALIDATION

3.1. σ^0 values

The validation of the backscattering coefficients is performed by comparing the values computed by the new processor with the value obtained by the old one, through comparison of the output products. Figure 8 shows the difference between the σ^0 values computed using the old processor and the values computed using the new processor on data acquired when the spacecraft still had a nominal attitude. Although not necessary, the yaw estimation was left active. The differences are acceptable and probably caused by minor differences in the geometry or in the normalization of the raw measurements. The large differences around the gaps are due to a change in the criterion used to assess wether a node is considered valid or not. The new processor requires a minimum number of measurements samples contributing to a node to declare that node as valid.

To assess wether the Doppler frequency shift correction routines are functional, data acquired when the spacecraft did not have a nominal attitude was processed. As comparison point, σ^0 values obtained from data acquired when the spacecraft had a nominal attitude and processed by the old processor was used. Since the acquisition time of the data was different, the only meaningful comparison points are σ^0 values over land. While the ground is more stable than the sea surface, actual ground conditions will affect the backscattering coefficient. Other factors such as the look angle may also influence the result and justify for a larger difference observed in figures 9 and 10 in comparison with figure 8. The figure

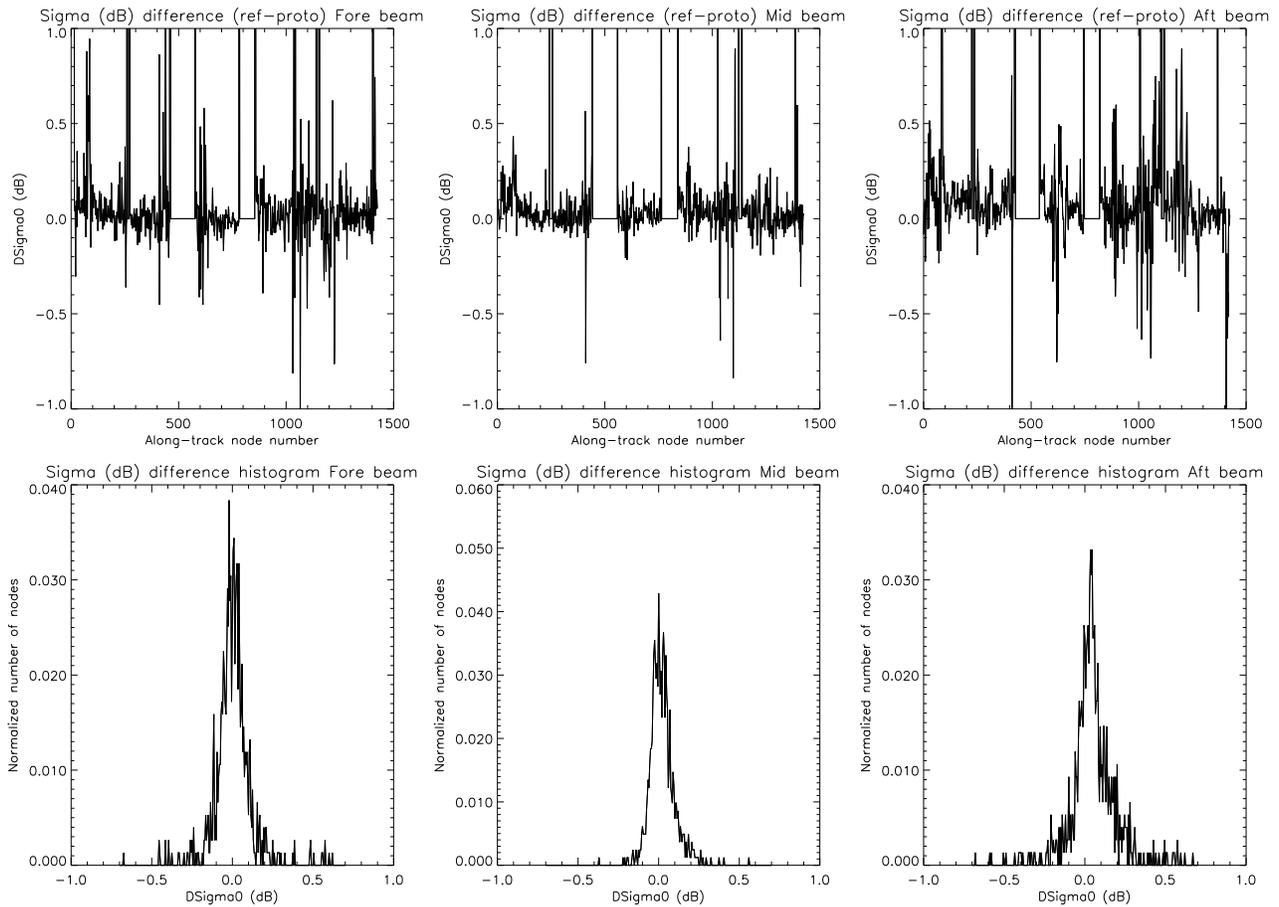


Figure 8. Upper graphs: along-track mid-swath nodes σ^0 difference between the old processor and the new one, for the three beams. Lower graphs: histogram of these differences.

9 shows the comparison results for data acquired when the spacecraft was in fine pointing mode (FPM) and hence had a known but non-nominal yaw angle. The comparison between the old processor and the new one over the data acquired in FPM shows that the new processor is effectively able to recover the backscattered energy that falls outside the system bandwidth of the old processor. Moreover, a geometrical effect is also visible in the form of a lateral shift of the σ^0 curve.

The figure 10 shows the comparison results for data acquired when the spacecraft was in zero-gyro mode (ZGM), the new “nominal” mode, and with a yaw angle a priori unknown.

3.2. Noise power

One measurement sequence consists in 32 pulses per antenna. Between the emission of a pulse and the corresponding return from the earth target, there is some dead-time. During 28 of these 32 pulses, that dead-time is used to measure the noise. During the remaining 4 pulses, the dead-time is spent measuring the emitted and delayed pulse for internal calibration purposes. The corresponding in phase and in quadrature noise measurements intensities are averaged together on-board and the corresponding averages are transmitted to earth. For monitoring purposes, an averaged version of the noise measurements is made available in the output wind product. Figure 11 shows the along-track evolution of the monitored noise power for one orbit for the two processors. A noise power level difference can be seen between the values reported by the two processors. This difference is due to the fact that the noise power reported by the new processor is compensated for the non linearity of the on-board ADC. A coupling between the I and the Q channels can also

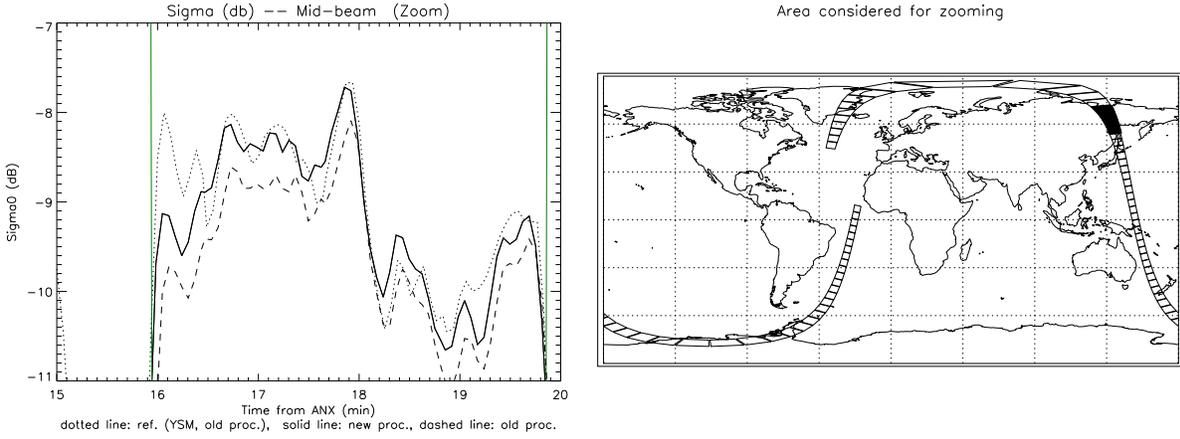


Figure 9. Left: comparison between the σ^0 value over land between the old and the new processor for a data set acquired in FPM mode. The nodes located between the two vertical lines are considered to be on earth. Right: orbit of which the data are extracted. The products filled black were considered for the comparison.

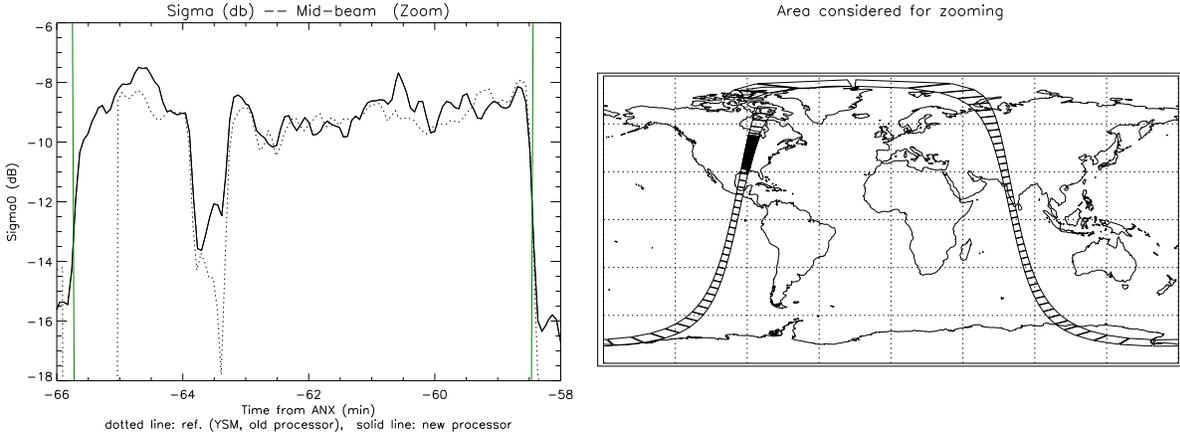


Figure 10. Left: comparison between the σ^0 value over land between the old and the new processor for a data set acquired in ZGM mode. Right: orbit of which the data are extracted. The products filled black were considered for the comparison.

be noticed. That coupling comes from the fact that the compensation coefficient used to correct the ADC non-linearity in both channels is obtained from the total noise power ($I^2 + Q^2$). The ADC non-linearity correction might seem important. This is due to the fact that the noise power barely exceeds one ADC count^{2,3} — as can be seen on the noise power level reported by the old processor — and the ADC non-linearity correction is computed to take into account a Rayleigh distribution for the noise amplitude. The holes in the noise power reported by the old processor are artefacts due to the presence of data gaps.

The processors also flags the products if the noise level for the beams inside the products exceeded a certain threshold. The standard UWI product is composed of 19×19 nodes and it was not possible to identify precisely the nodes to which the packets having a noise spike did contribute. A more detailed analysis shows the noise increase was probably due to a noise source on ground since it is geometrically coherent on the 3 antennas, as shown in figure 12.

3.3. K_p value

The normalized standard deviation of σ^0 measurements, K_p , is commonly⁴ used to evaluate the accuracy of the σ^0 measurements and is one of the input to the wind extraction process:

$$K_p = \frac{\sqrt{\text{var}\{\sigma_{\text{meas}}^0\}}}{\sigma^0}. \quad (1)$$

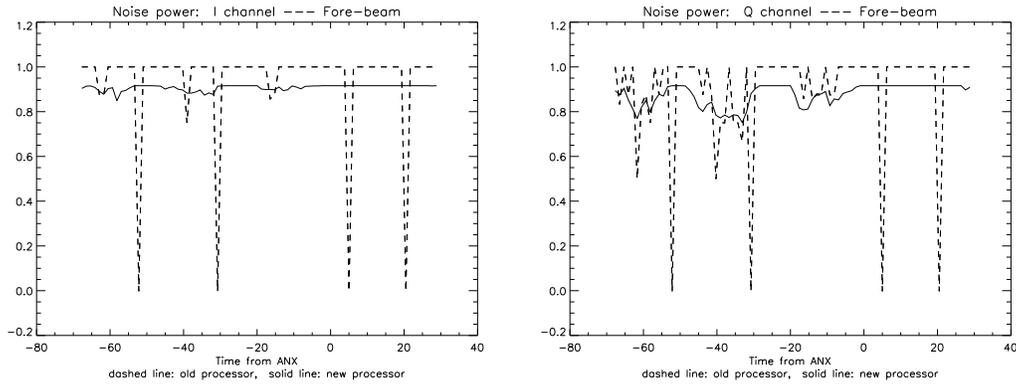


Figure 11. Typical along-track evolution of the per product noise power value for the old processor and for the new processor (left: in-phase (I) channel, right: in quadrature (Q) channel)

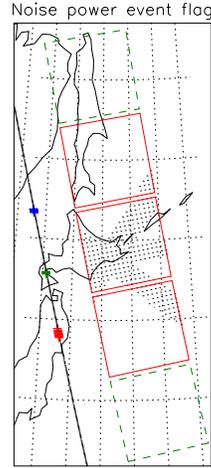


Figure 12. Noise power event. The solid rectangles show the flagged products while the dashed rectangles show the limits of unflagged products. The black dots shows the flagged nodes and the crosses show the flagged measurement sequences on the subsatellite track.

While the new processor uses the expression of eq. (1), the old processor assumes $\text{var}\{\sigma_{\text{meas}}^0\}$ is equal to the measured noise power corrected by a factor taking the correlation of the measurements into account. Doing so, the old processor only takes into account the noise level measured on-board while the new processor considers the variations of the σ^0 and thus is also able to take into account the noise due to speckle, inherent to coherent imaging systems.

3.4. Relative calibration

The relative calibration across the swath for a given antenna and between the different antennae can be performed using distributed targets,⁵ an area of known and constant backscatter. The tropical rain forest in South America is one of the test site for ERS-2 σ^0 measurements relative calibration. This test site is assumed to have a constant, uniform and isotropic backscattering coefficient. It is shown⁵ that the γ^0 defined as

$$\gamma^0 = \frac{\sigma^0}{\cos\theta} \quad (2)$$

where θ is the incidence angle corresponding to the σ^0 measurement, is independent of the incidence angle over the rain forest. This value can thus be used to perform the relative calibration of the σ^0 measurements.

Figure 13 shows the γ^0 histogram for data acquired over the test area. The histograms obtained from data produced by the processor with yaw estimation have consistently a smaller standard deviation than the one computed from data

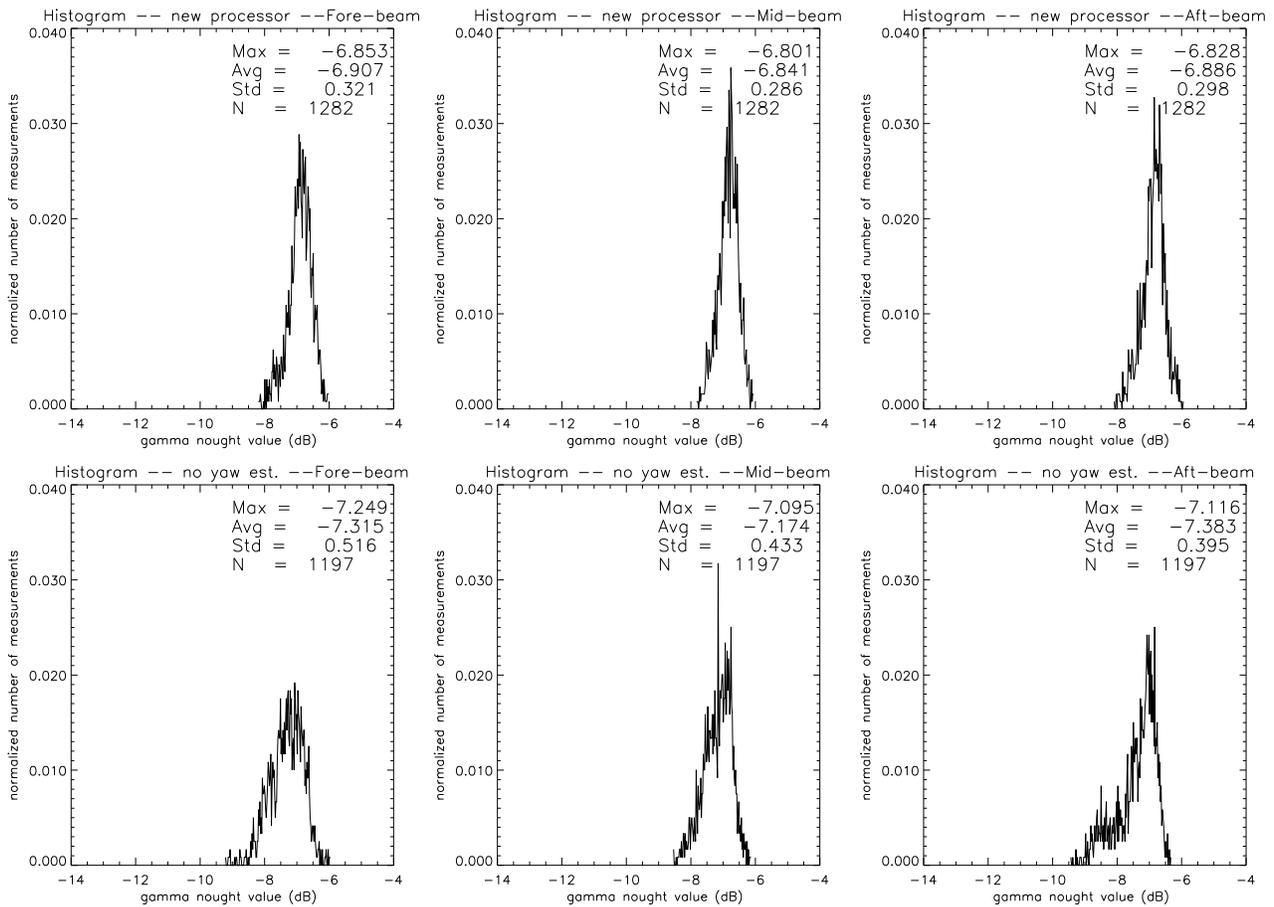


Figure 13. γ^0 histogram for data acquired during the ascending passes over the Brazilian rain forest for data acquired in zero-gyro mode. The upper graphs were obtained from products generated by the processor with yaw angle estimation while the lower graphs are from products produced by the processor without yaw angle estimation (yaw angle assumed equal to zero). The corresponding yaw angles made by the spacecraft are plotted in fig. 14.

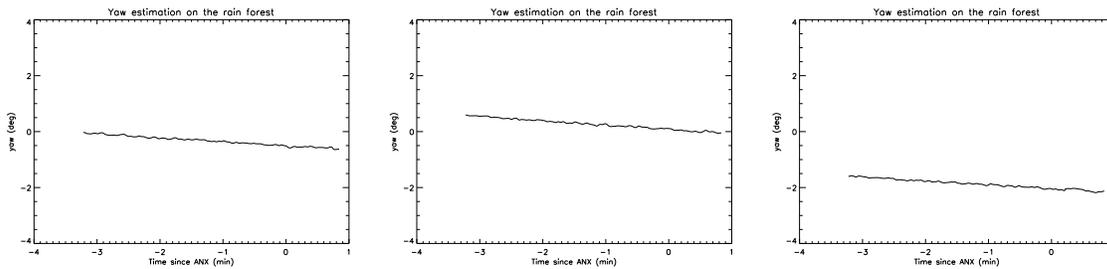


Figure 14. Yaw angle of the spacecraft over the Brazilian rain forest when the data used to produce the graphs of figure 13 were acquired.

produced by the processor without yaw estimation. Moreover, the variability of the maximum and average of the γ^0 histogram of data obtained with the new processor is within 0.02dB (fore/aft beams) and 0.1dB (fore/mid beams), which is consistent with the values reported in.⁵ While these results are very promising, they were however obtained on a limited data set. More extensive validation will have to be performed to confirm these figures.

3.5. High-resolution product

The nominal resolution of the output product of the scatterometer processor is of about 50km. This resolution is adequate for wind-retrieval applications over sea, the main application for scatterometer measurements. Scatterometer data is however also used for applications over land and ice where a higher resolution product could be useful.⁶

The output product is obtained by spatially averaging the individual measurement samples and sampling the filtered result on a regular grid of nodes. The characteristics of the spatial filter used determines the product's resolution, the distance between the nodes being chosen accordingly to avoid aliasing artefacts. While the nominal resolution product uses a separable Hamming window of 86km size, the enhanced resolution product will use the same type of window, but having only 43km size. The corresponding inter-node distance will nominally be of 12.5km.

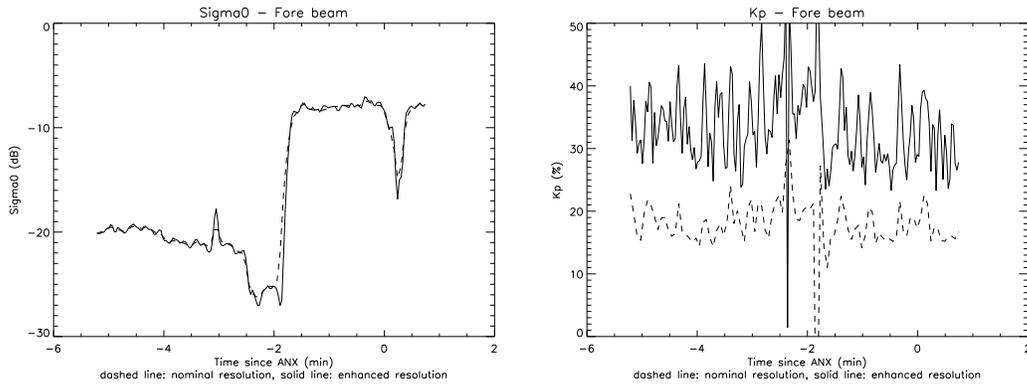


Figure 15. Comparison between the nominal resolution and the high-resolution products. The along-track evolution at mid-swath of respectively the σ^0 and the K_p are presented in the left and right graphs.

Figure 15 presents a comparison of the high-resolution product with the nominal resolution product. An along-track cut at mid-swath in the σ^0 surface is presented left. The increase of resolution is clearly noticeable, both inside the small “dip” as at the sea/land transition. The graph on the right presents the along-track evolution of the K_p at mid-swath for the two products. The K_p of the high-resolution product is roughly twice as high as the K_p of the low-resolution product. This was to be expected. The K_p is the variance of the σ^0 , and is computed from the individual measurement samples. Since the high-resolution σ^0 value is obtained by averaging roughly 4 times less measurement samples, the corresponding variance will double and hence the K_p . This is the price to pay for the high-resolution.

4. CONCLUSIONS

After the gyroscopes failure, the perfect yaw steering mode of the ERS-2 spacecraft could not be guaranteed anymore. Since actual values for the yaw angle vary around the nominal value, the existing scatterometer processing chain is not able to convert the measurement made to calibrated σ^0 . A processor thus had to be upgraded to support the processing of data acquired with the new attitude.

The upgraded processor was shown to be able to correctly estimate the actual yaw angle of the spacecraft with respect to its perfectly yaw-steered nominal attitude. The corresponding residual Doppler frequency shift is also correctly corrected, with a direct impact on the radiometric quality of the output product. The relative calibration of the instrument over the Brazilian tropical rain forest is also restored to values close to the nominal ones.

As part of the upgrade, several enhancements were introduced, such as the enhanced node land/sea determination or the enhanced noise power event detection. More fundamentally, a new high-resolution mode was introduced, essentially to support land and ice applications.

The results presented in this paper have been obtained on a limited but representative data set. They are very promising and will have to be confirmed by larger scale evaluation once the upgraded processor is put in operation.

ACKNOWLEDGMENTS

This work was performed under European Space Agency (ESA) contracts. We would like to thank the ESA for the use of data and the provision of industry-confidential information. This work would not have been possible without the help of Pascal Lecomte and Raffaele Crapolicchio from the European Space Agency/ESRIN.

REFERENCES

1. X. Neyt, P. Pettiaux, and M. Acheroy, "Scatterometer ground processing review for gyro-less operations," in *Proceedings of SPIE: Remote Sensing of the Ocean and Sea Ice 2002*, **4880**, 2002.
2. P. Lecomte, "The ers scatterometer instrument and the on-ground processing of its data," in *Proceedings of a Joint ESA-Eumetsat Workshop on Emerging Scatterometer Applications – From Research to Operations*, pp. 241–260, ESTEC, (The Netherlands), Nov. 1998.
3. R. Crapolicchio and P. Lecomte, "The ers wind scatterometer mission: routine monitoring activities and results," in *Proceedings of a Joint ESA-Eumetsat Workshop on Emerging Scatterometer Applications – From Research to Operations*, pp. 285–298, ESTEC, (The Netherlands), Nov. 1998.
4. D. G. Long and M. W. Spencer, "Radar backscatter measurement accuracy for a spaceborne pencil-beam wind scatterometer with transmit modulation," *IEEE Trans. on Geoscience and Remote Sensing* **35**, pp. 102–114, Jan. 1997.
5. P. Lecomte, "Ers wind scatterometer commissioning and in-flight calibration," in *Proceedings of a Joint ESA-Eumetsat Workshop on Emerging Scatterometer Applications – From Research to Operations*, pp. 261–270, ESTEC, (The Netherlands), Nov. 1998.
6. D. Early, D. Long, and T. Stein, "Scatterometer resolution enhancement," in *Geoscience and Remote Sensing Symposium Proceedings, IGARSS '98*, pp. 1970–1972, 1998.