To: File
From: Walter H. F. Smith, NOAA Geophysicist
Re: Near-real-time wind-wave and height products from CryoSat: progress to date
Date: 27 July 2011

Summary

We estimate significant wave height (SWH), wind speed, sea surface height, and other variables from CryoSat2 fast-delivery ("FDM") and/or Level-1b ("L1B") conventional ("low rate", or "LRM") mode data. We search for new files hourly, and when a new file is found, we: (1) retrack the 20 Hz waveforms; (2) average derived values to 1 Hz; (3) remove land values and reformat for the NOAA Advanced Weather Interactive Processing System (AWIPS); (4) export to NOAA forecasters (Figure 1, next page).

The Radar Altimeter Database System (RADS) also ingests the output of steps (1) and (2), sorts the data segments into a pass and cycle structure, updates orbits and geophysical corrections as they become available, and facilitates comparison of our results to those of Jason-1, Jason-2, and EnviSat. Inter-satellite comparison of the histograms and geographical distribution of our estimates provides validation of our approach. Examples are discussed in the following pages.

Preliminary findings:

- CryoSat provides wind speed, wave height, and sea level anomaly data that are essentially equivalent to those obtained from Jason-1, Jason-2, and EnviSat.
- Inter-satellite comparison between our CryoSat results and those from J-1, J-2 and EnviSat shows similar geographical patterns and histograms, validating our results.
- The range precision of the CryoSat altimeter is excellent: around 1.5 cm for a one-second average in conventional mode.
- Wind speed estimates derive from the radar backscatter coefficient known as “sigma nought”, σ_0. We compute orbit- and retracker-derived corrections that we apply to CryoSat’s reported automatic gain control (AGC) data to obtain an adjusted AGC. Adjusted AGC should be within an unknown constant of σ_0. We lack documentation of this constant, but inter-satellite comparison suggests that we have made a good guess.
- We find an unexpected pattern in the off-nadir pointing of CryoSat estimated from retracking. The magnitude and geographical pattern in these estimates disagrees with the platform attitude information available in the FDM/L1b and L2 data products. We have not yet been able to understand what we are seeing here. We cannot rule out the possibility that this is related in some way to CryoSat’s elliptical antenna pattern, which we have approximated as circular in our retracker.

We would be delighted to discuss our work with CryoSat experts, perhaps at the 5th Coastal Altimetry Workshop and Ocean Surface Topography Science Team meetings.

Our near-real-time production of wind and wave data products will stay within NOAA, and the RADS version of our results will stay on the non-public side of RADS, unless and until we have authorization from ESA to distribute our results more widely.

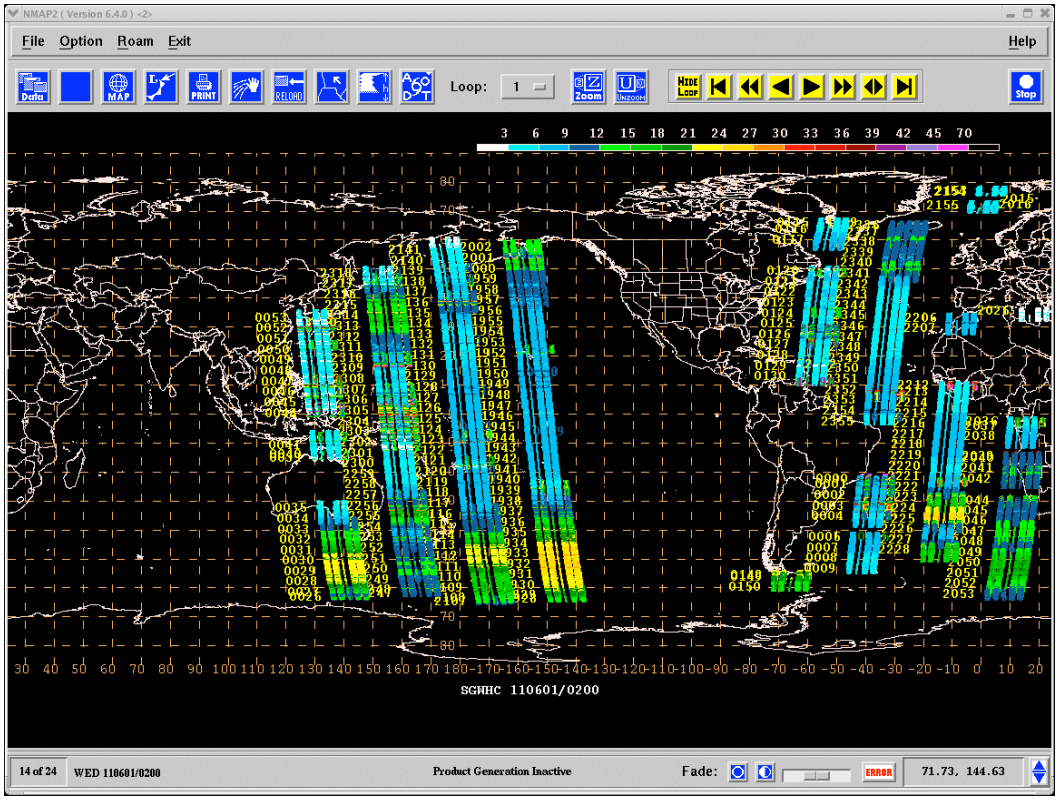

Figure 1. NOAA N-AWIPS display showing six hours of CryoSat wave height data. Image: G. McFadden.

Developers

The algorithm development and software coding of the retracking, CryoSat file reading, and ASCII output, was by Walter Smith. It is based on Walter's earlier work retracking other satellites, and building a library of CryoSat file read routines. Remko Scharroo modified Walter's algorithm to add NetCDF output and hooks for ingest into the Radar Altimeter Database System (RADS). John Lillibridge built a shell script using GMT tools and awk to reformat the data into the form ingested by the AWIPS system and create hourly files. Remko and John built cron jobs, makefiles, and other scripts to automate the process. Remko provided figures using RADS to compare CryoSat, Envisat, Jason-1 and Jason-2. Additional help with off-nadir analysis came from Eric Leuliette.

Leuliette, Lillibridge, Scharroo and Smith are PIs or Co-Is on the following CryoSat-2 investigations: #2689 (mean sea level), #2690 (comparing LRM and SAR), #2705 (marine gravity), #4510 (real-time wind and wave), #4518 (inter-satellite cross-validation).

Next Steps

We are now developing a retracker for SAR and SARIN ocean waveforms. If we are successful, we will investigate whether we can use these data to “patch the holes” in LRM ocean coverage. If SAR and SARIN data are available with fast-enough delivery, these modes may be added to our near-real-time product.

Retracking algorithm

We derive estimates of arrival time (radar range), SWH, returned power, noise power, and apparent off-nadir pointing, by fitting a parametric model to the radar echo waveform, a process known as "retracking". Our process closely follows that of the "MLE3" and "MLE4" retrackers used for Jason-1 and Jason-2. We retrack 104 gates from the middle of the 128-element CryoSat LRM waveform array, following the Jason-1&2 practice of removing the first and last 12 gates to avoid the portion biased by wrap-around from the deramp FFT. The fit minimizes the scaled sum of squares of errors, called "mean quadratic error" (MQE) in MLE3/4. We use the same search initialization, Gauss-Newton iteration steps, and convergence criterion as in MLE 3/4. The "Brown model" employs a Gaussian approximation for the point target response and an exponential approximation of the modified Bessel function I_0. For the results shown here, we used a first-order exponential approximation for I_0, but allowing off-nadir angle to be a free parameter. We followed MLE3/4 practice in holding the pre-arrival noise level fixed to the average of the first 10 gate samples. The "Brown model" assumes an axisymmetric antenna beam. We used a circular beam with a half power beam width (HPBW) equal to the harmonic mean of the major and minor HPBW values of CryoSat's elliptical antenna beam pattern.

Preliminary Validation of Wind, Wave, and Sea Level Anomaly

As indicated in the Summary above, wind speed and wave height can be obtained directly by retracking the FDM and L1B data. Thus these can be near-real-time products. To obtain the most accurate sea level anomaly, we use revised orbits and geophysical corrections available after real time. The results shown here validate all three quantities through inter-satellite comparisons. In real-time, we cannot make these comparisons; results shown here are analyses after the fact, but using results derived from CryoSat FDM (Fast Delivery Mode) data files.

The following figures, from Remko Scharroo, show results from CryoSat, Envisat, Jason-1, and Jason-2. For CryoSat, data are taken from April 19 to May 18, a time period called "subcycle 14" in RADS, and labeled "fdm1r"; note that this period is approximately 29 days. Comparison data from EnviSat (labeled "n1") are taken from cycle 102, spanning April 25 to May 25, a 30-day period centered about one week after CryoSat subcycle 14. Comparison data from Jason 1 cycle 344 and Jason 2 cycle 105 (labeled "j1j2") cover the May 4 to May 19 period, a span of about 15 days roughly coincident with the latter half of CryoSat subcycle 14. Since the space-time sampling of these altimeters is not coincident, we should not expect the figures below to show exact agreement between the various satellites.

Values from ascending and descending tracks are shown separately. Maps show all values, while histograms are drawn from only those values located over ocean. For CryoSat, "all values" means 1 Hz averages of all 20 Hz points for which the retracker converged to a valid ocean waveform solution. For Envisat, J-1 and J-2 we show the 1-Hz averages as they appear on the GDR products, with histograms again constructed from data over ocean only.

Significant Wave Height (SWH)

The most important parameter for some NOAA forecasters is Significant Wave Height. SWH exceeding 12 feet (approximately 3 m) is a threshold of interest in this context. We examined the overall SWH performance and also specifically the 3-m threshold. We find that the CryoSat values look essentially identical to EnviSat and the Jasons, although the Jasons report zero SWH values more frequently than EnviSat or CryoSat. There does not appear to be any systematic difference between ascending and descending SWH values in any of the altimeter systems.

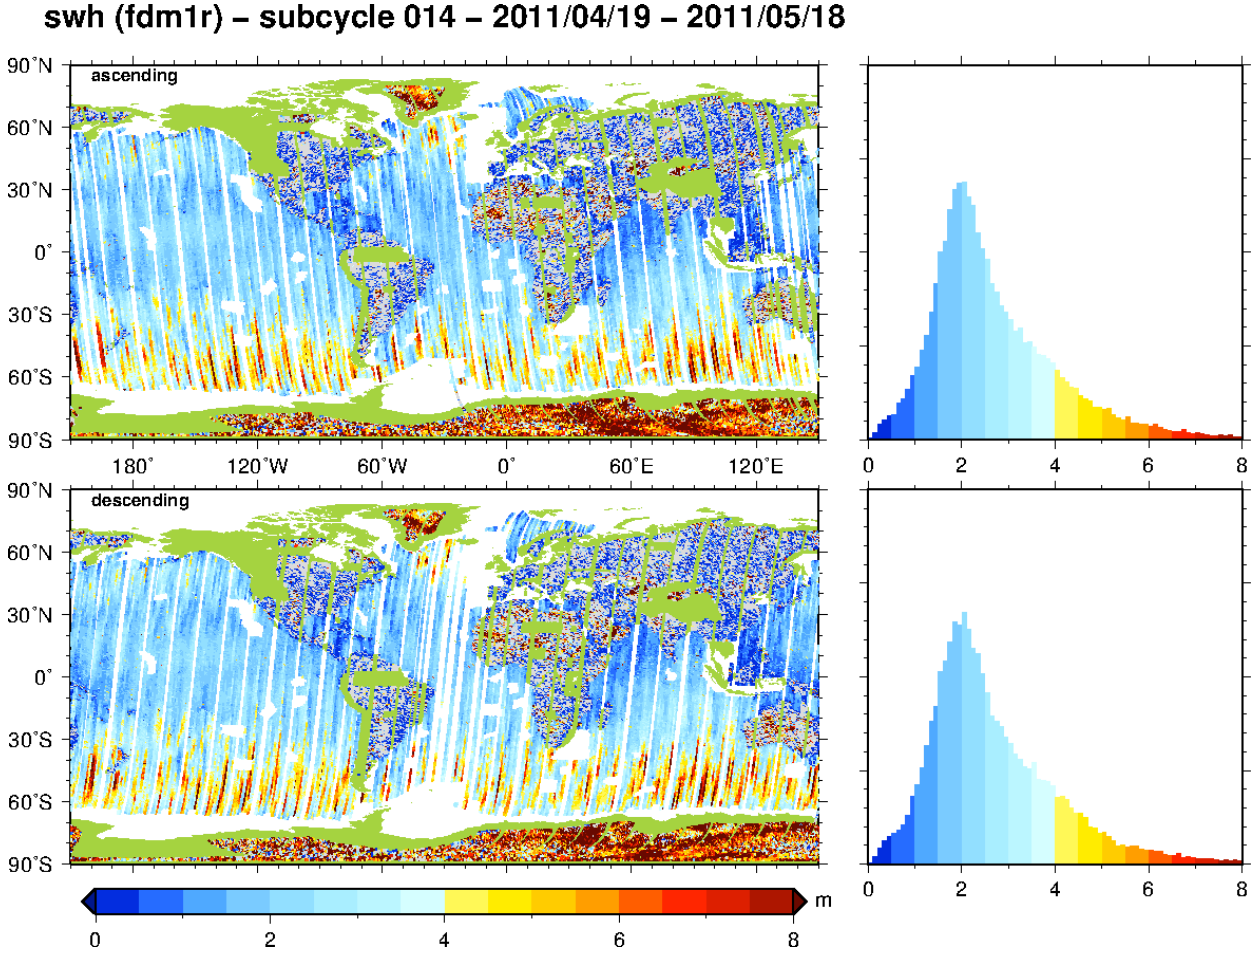

Figure 2. Significant Wave Height from CryoSat.

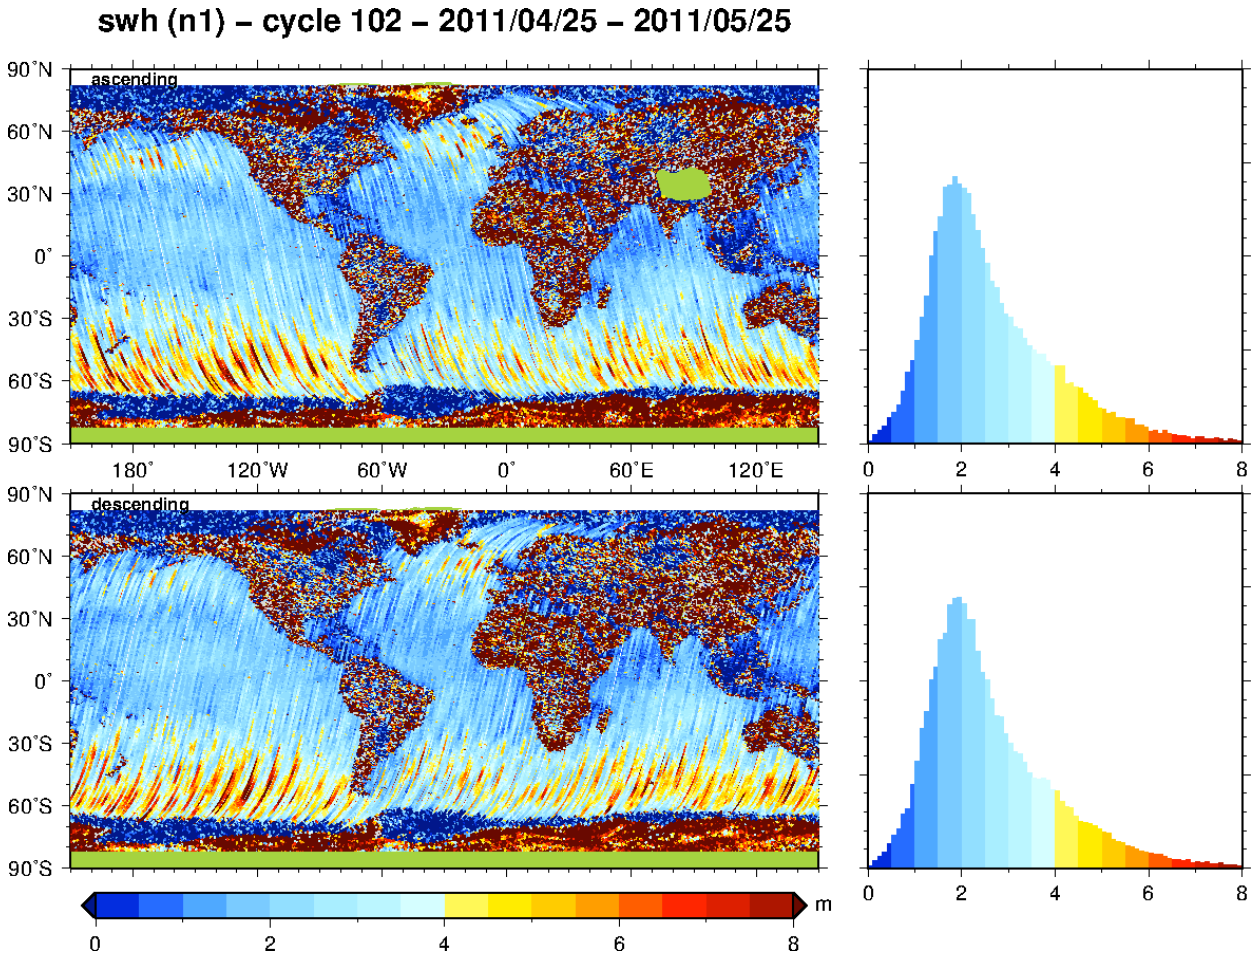

Figure 3. Significant Wave Height from EnviSat.

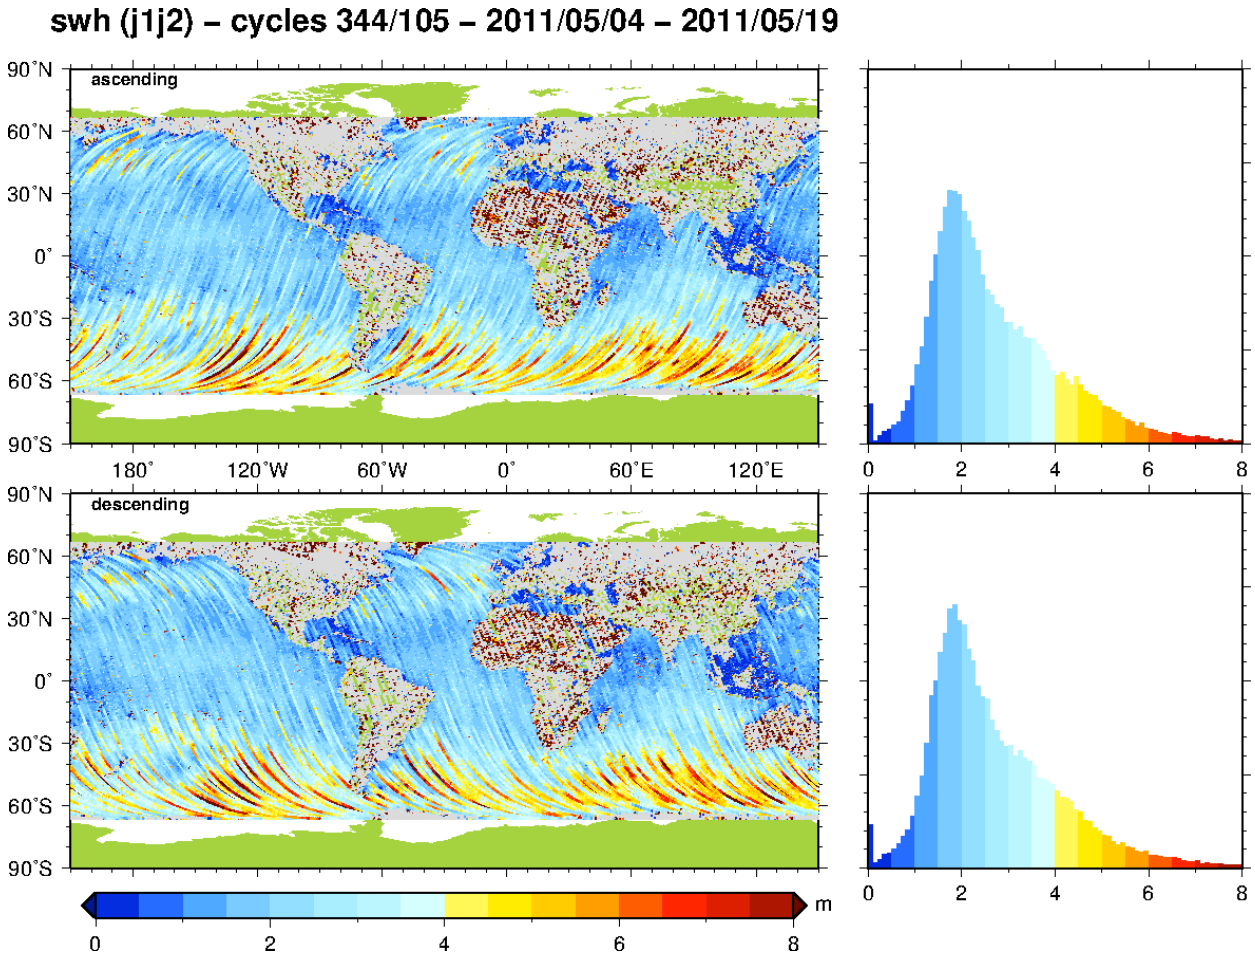

Figure 4. Significant Wave Height from Jason-1 and -2.

Sea Level Anomaly (SLA)

In order to achieve the highest accuracy in the 1-Hz averaged sea level anomaly, the radar range must be retracked at 20 Hz, after which the retracked range values are subtracted from the corresponding 20-Hz orbit heights. This yields raw sea surface height, equal to orbit height minus range, sampled at 20 Hz. These raw heights, taken over one second, are then fit with a least-squares straight line, and the value of the line at the time of the 1-Hz mean orbit is then taken as the 1-Hz averaged raw height. Corrections to raw height are then applied at 1 Hz.

The Sea Level Anomaly shown here is derived in RADS by correcting the CryoSat raw sea surface height using the DNSC08 mean sea surface model, GOT4.7 tides, MOG2D dynamic atmospheric correction, NCEP wet and dry propagation delays (these appear to be almost always invalid on the CryoSat product) and the precise CNES orbit. For Jason-1 & -2, and EnviSat, the corrections are the same, except that the ECMWF wet and dry delays are used. The near-real-time operational orbits for some of these satellites are not very accurate, so RADS searches for precise orbits after real time. Several of these correction fields are revised days to weeks after real time, and RADS makes updates as available. Therefore the best SLA is not available in near-real-time. What we show here is obtained after the best correction fields become available, but retaining the retracked range from the FDM CryoSat waveform.

We find SLA maps and histograms look similar for all the satellites, and there is no systematic difference between ascending and descending tracks.

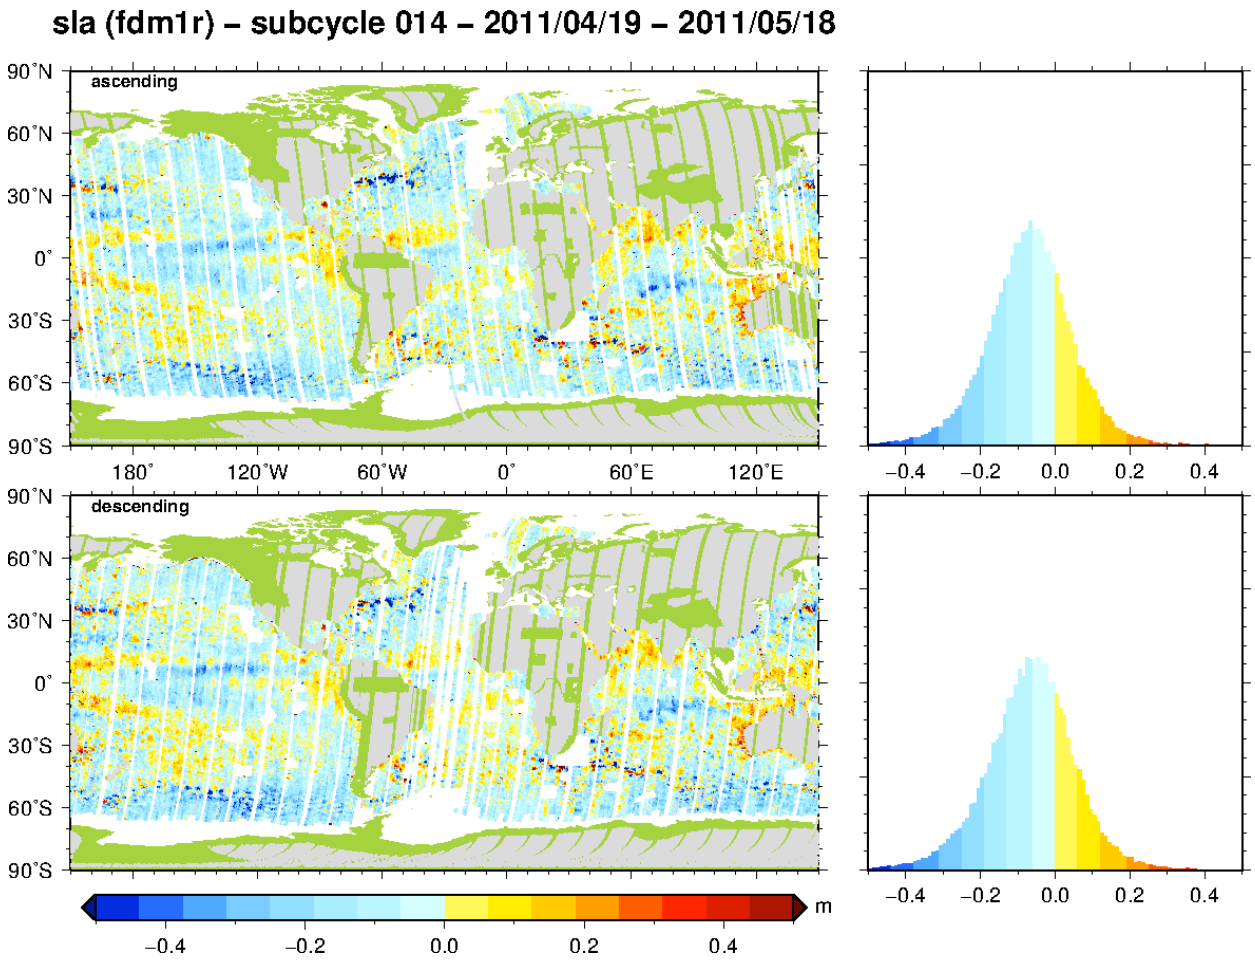

Figure 5. Sea Level Anomaly from CryoSat.

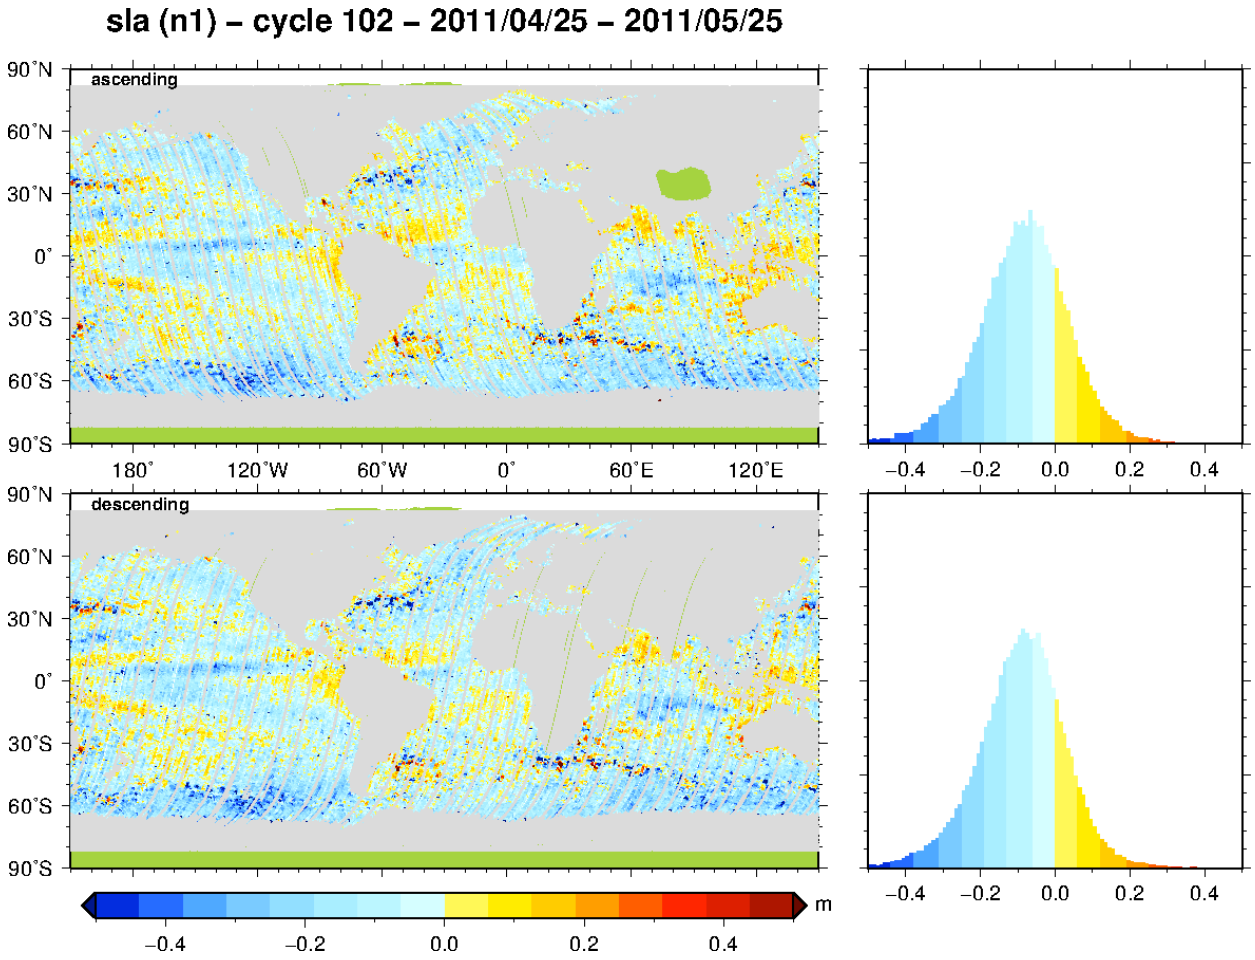

Figure 6. Sea Level Anomaly from Envisat.

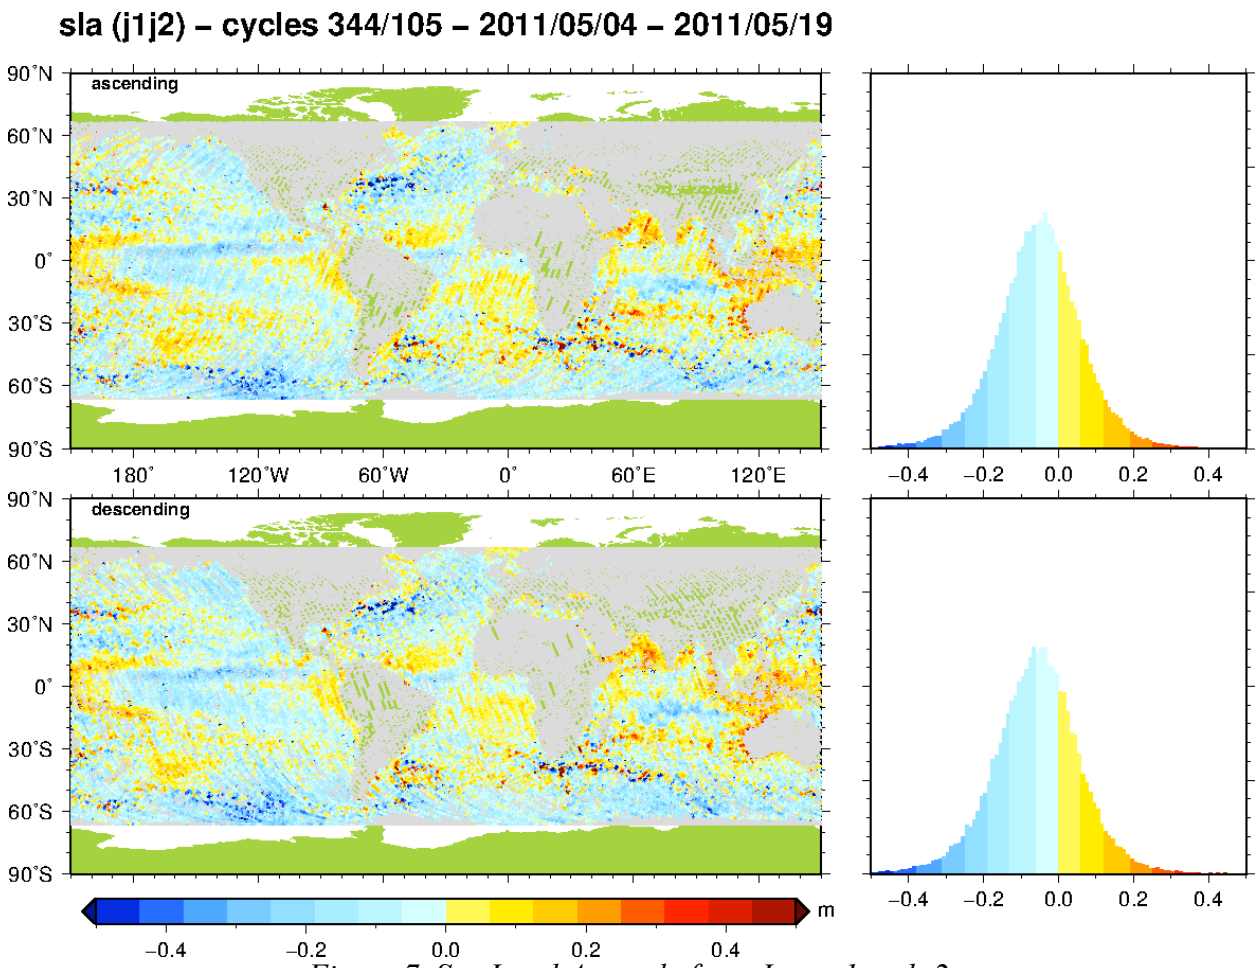

Figure 7. Sea Level Anomaly from Jason-1 and -2.

Wind Speed

Wind speed may be derived from the surface backscatter parameter σ_0 in various ways. Here we use the one-parameter ECMWF wind speed model of Abdalla (2007) that is used also for EnviSat. σ_0 derives from Automatic Gain Control (AGC) and various corrections to AGC. We do not know how to obtain σ_0 from CryoSat, so we have made a guess. We estimated a constant to apply to move retracker-corrected AGC to σ_0. The constant represents 10 $\log_{10}$ of the net gains and losses in the CryoSat system under typical operating conditions. Here we subtracted 24.0 from adjusted AGC to get σ_0. It appears that 24.3 may be better.

There is general agreement between the CryoSat estimates and the other satellites. Agreement is adequate for a near-real-time forecast application. However, detailed inspection suggests that there may be minor anomalies in our CryoSat estimates. The EnviSat and Jason maps show no differences between ascending and descending tracks, and the histograms look similar. CryoSat histograms are slightly different from the other satellites. In the maps below, CS ascending may show higher wind speed than CS descending in part of the North Pacific. The only CS parameter with an ascending/descending asymmetry we could find is the mispointing angle, but the dependence of σ_0 on mispointing seems too small to be the cause. Because wind speed is very sensitive to σ_0 and we had to guess at the σ_0 calibration, we should not read too much into this.

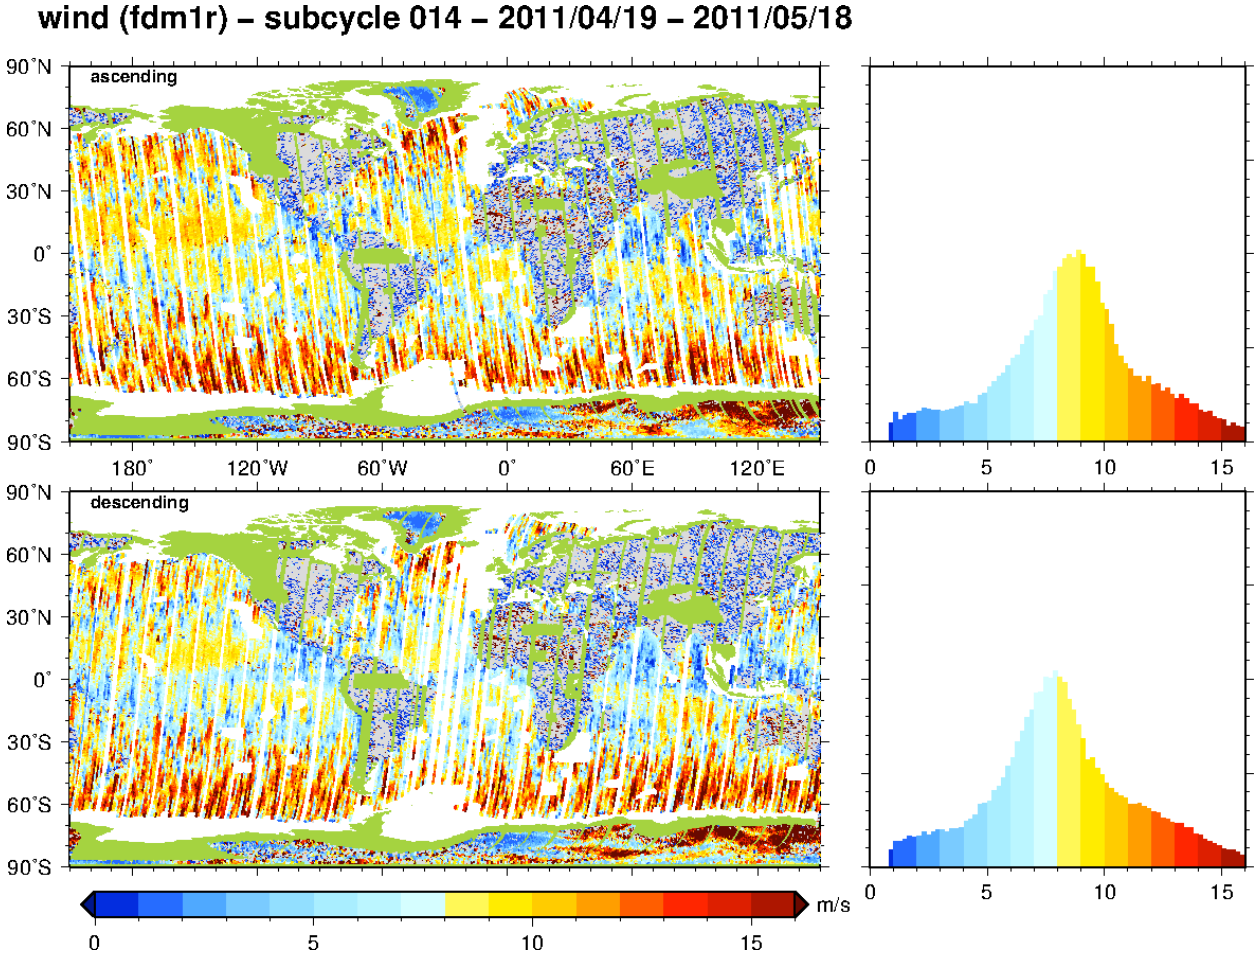

Figure 8. Wind speed estimated from CryoSat.

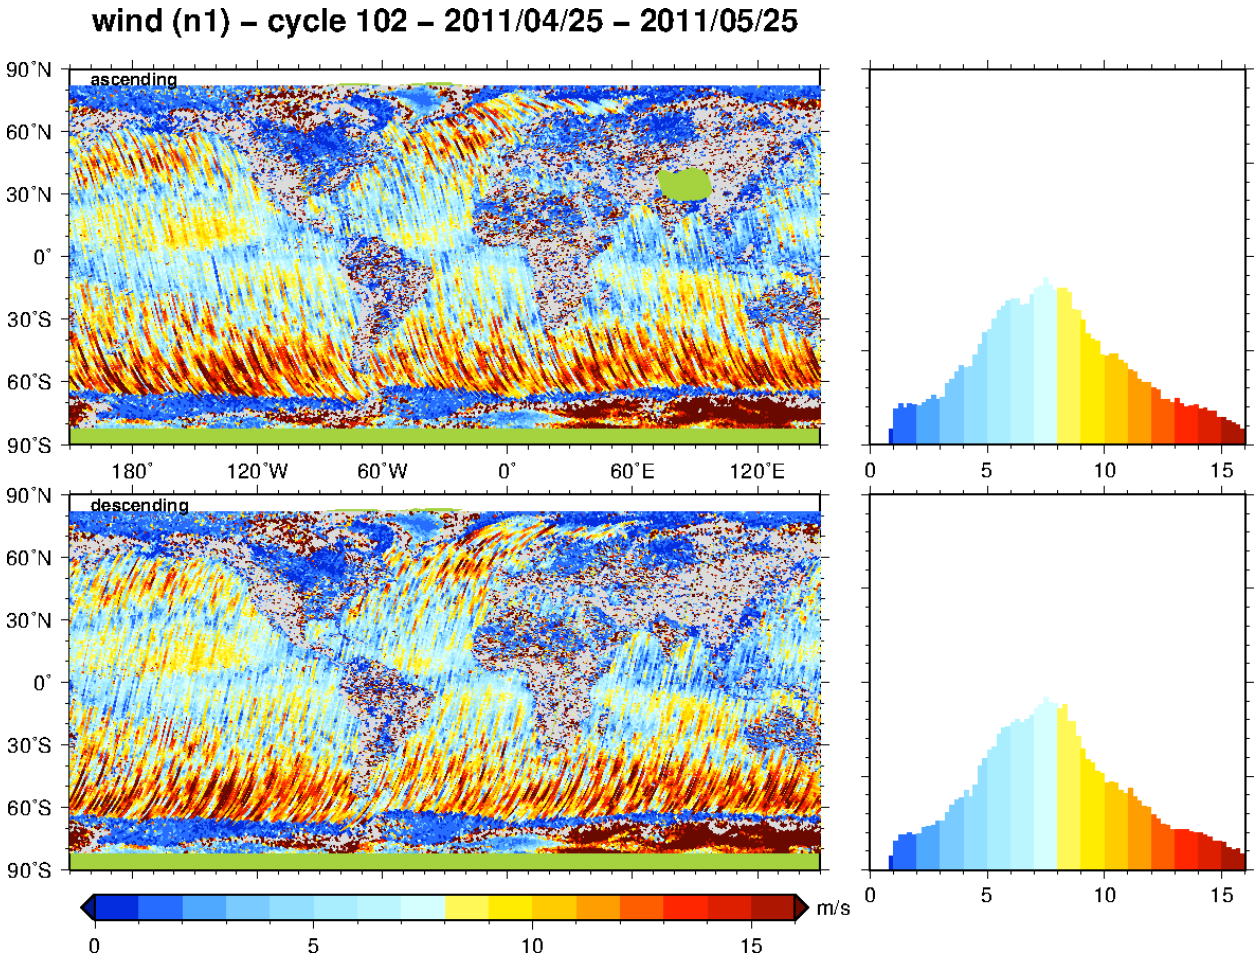

Figure 9. Wind speed from EnviSat.

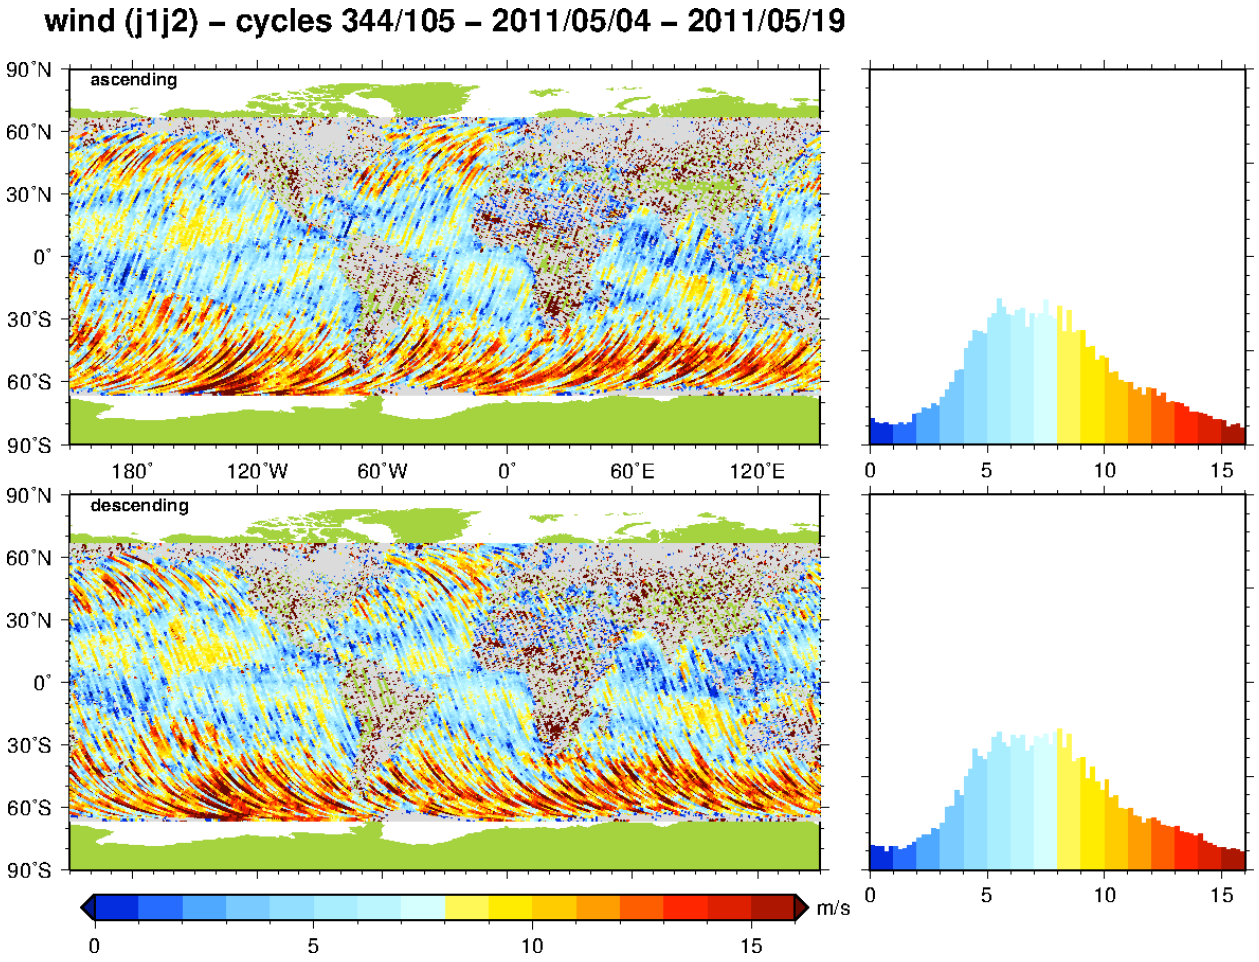

Figure 10. Wind speed from Jason-1 and -2.

Related parameters: off-nadir angle, and sigma-nought

Off-nadir angle

Our retracking procedure allows off-nadir angle to be a free parameter, as in the "MLE4" retracker used for Jason-1 and -2. CryoSat employs antennas with an elliptical beam pattern. Our procedure assumes CryoSat's beam pattern can be approximated with a circular pattern having a half power beam width (HPBW) equal to the harmonic mean of the short and long elliptical HPBW axes of CryoSat. For these values we used the "final" values for 13.575 GHz from the CryoSat-2 SAAB antenna calibration report.

Off-nadir angles we obtain are of the order of 0.15 degrees and display geographical and ascending/descending patterns. We also derived an off-nadir angle from the real beam and interferometer baseline vectors in the FDM/L1b data, and verified that it agreed with the mispointing attitude angle given in the L2 data. This platform attitude, derived from star trackers, shows a geographical pattern, but the angles are about a factor of 3 smaller than those we fit by retracking. It is possible that we are seeing a real effect of the elliptical antenna beam, or that our circular approximation of the beam is wrong.

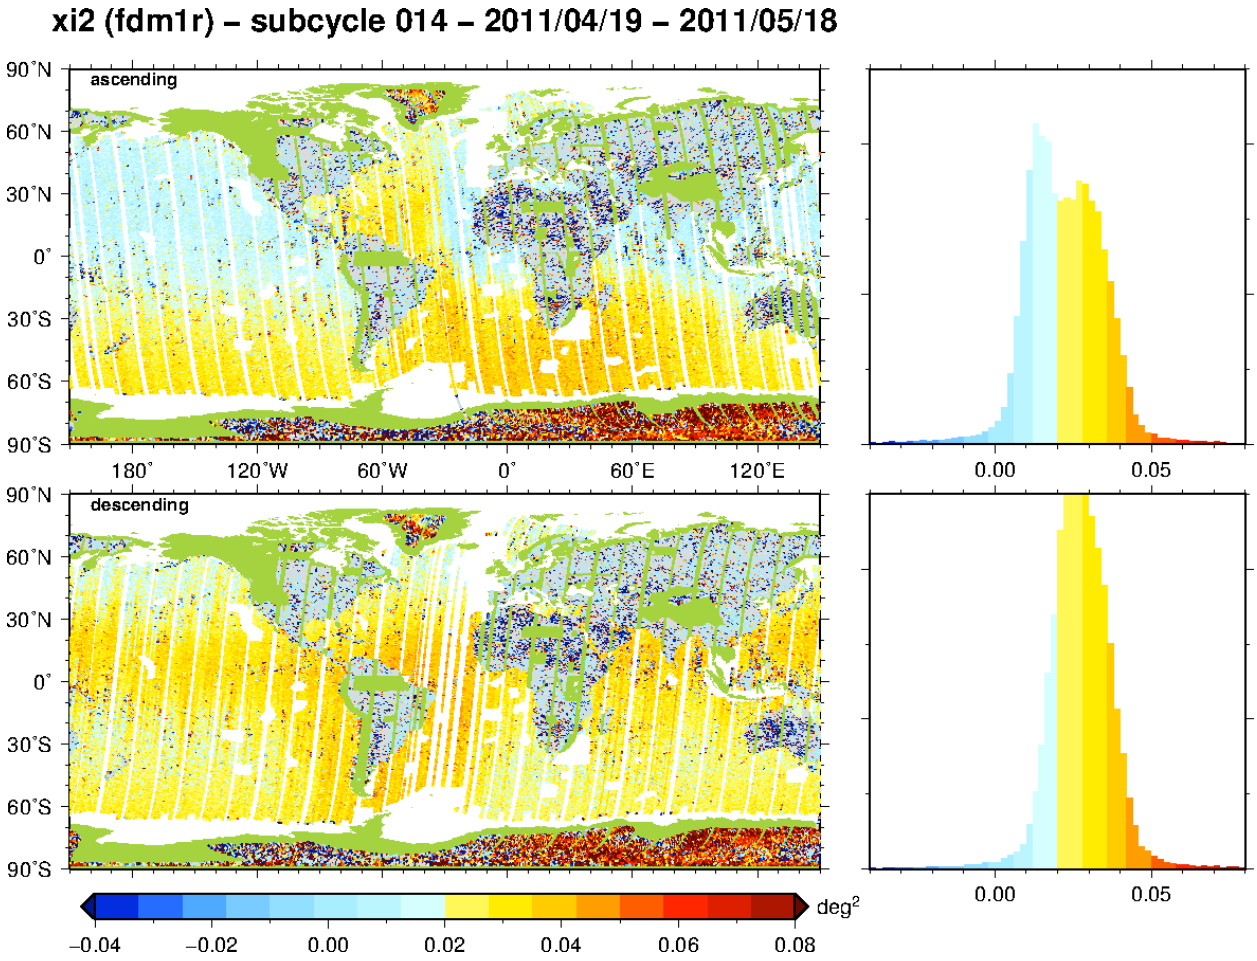

Figure 11. Off-nadir angle, squared, in degrees^2, obtained by retracking CryoSat waveforms assuming a circularly symmetric antenna approximation. Note geographical and ascending/descending patterns and asymmetry.

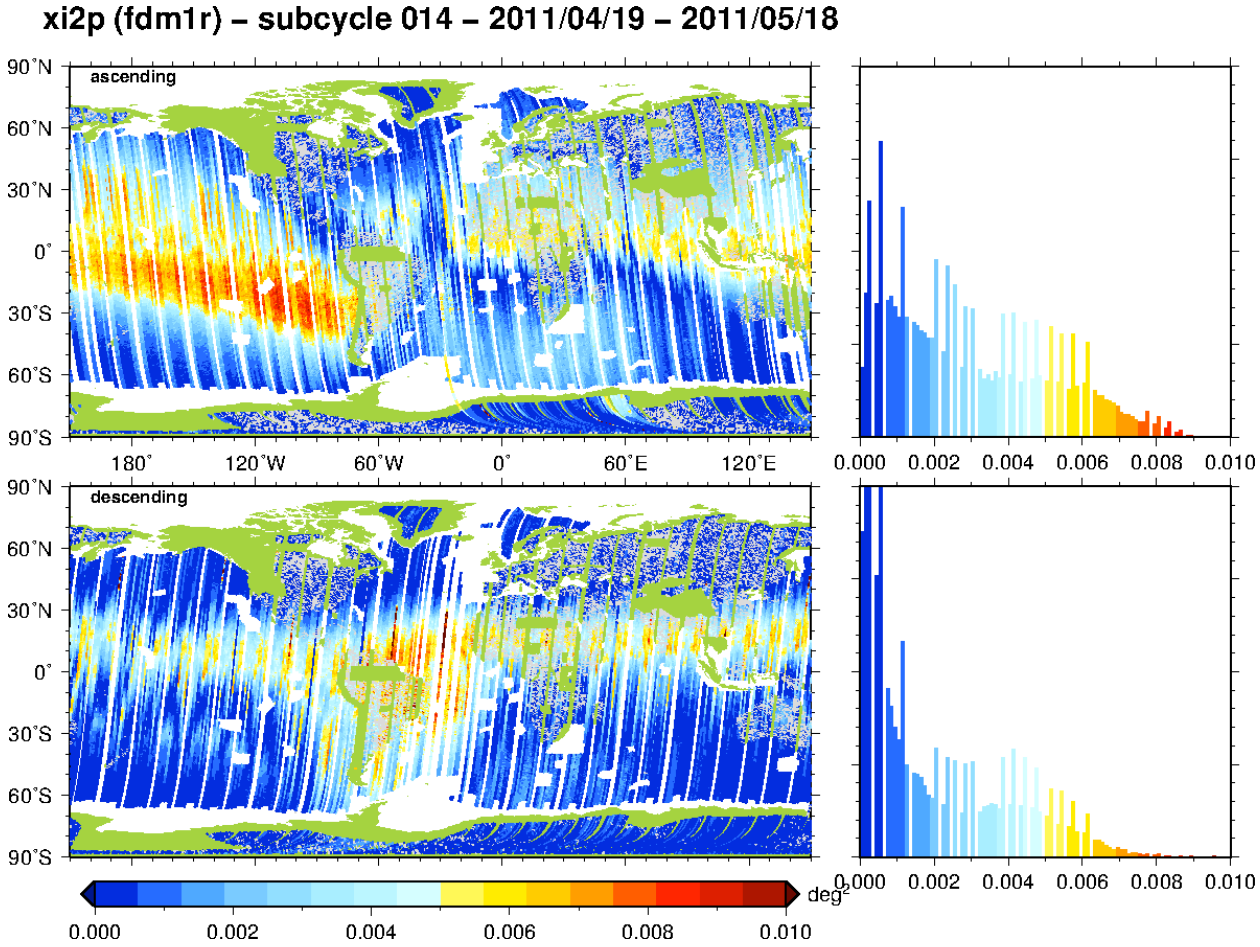

Figure 12. Off-nadir angle, squared, in degrees^2, of the CryoSat spacecraft platform, obtained from the beam direction and baseline direction vectors (ultimately derived from star trackers). There is a geographical pattern and an ascending/descending asymmetry, but the size of the mispointing angle is roughly a factor of 3 or more smaller than that obtained from retracking.

Sigma-nought

As noted above under Wind Speed, the wind speed estimate derives from σ_0. Retracking and orbit height furnish adjustments to Automatic Gain Control (AGC). Adjusted AGC must then be shifted by a constant to obtain σ_0. This constant is unknown (to us) for CryoSat; it depends on 10 $\log_{10}$ of gains and losses we do not know (these may be in cal files we do not have access to). We have estimated the constant by comparing with EnviSat and Jason-1 and -2. One difficulty here is that the other satellites do not fully agree on σ_0. The peak in the EnviSat histogram is a little below 11 dB, while for the Jason histogram it is a little above 11 dB. Jason-2 is known to have noisy σ_0 estimates relative to Jason-1; the waveform "compression" (averaging) employed in J-1 reduces the σ_0 noise. Note that the CryoSat histograms are narrower than the others.

All the σ_0 histograms are non-symmetric, with a bump on the lower side, reflecting the probability of high wind events. MLE3 output from Jason-1 resolves this bump but it is not resolved well in MLE4 output from Jason-2 due to the higher noise in J-2 σ_0.

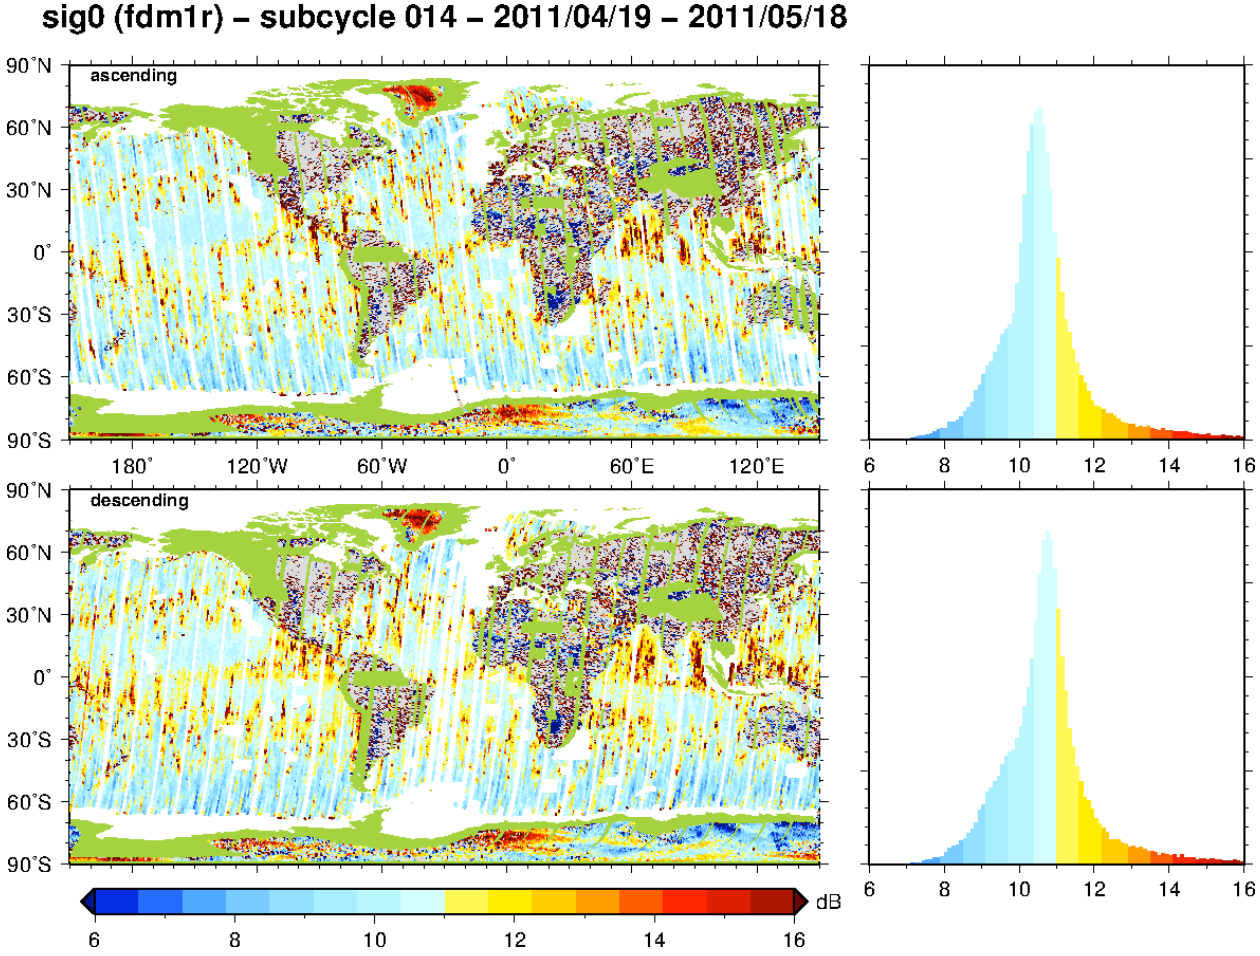

Figure 13. CryoSat σ_0 (dB) estimated from retracker-corrected AGC.

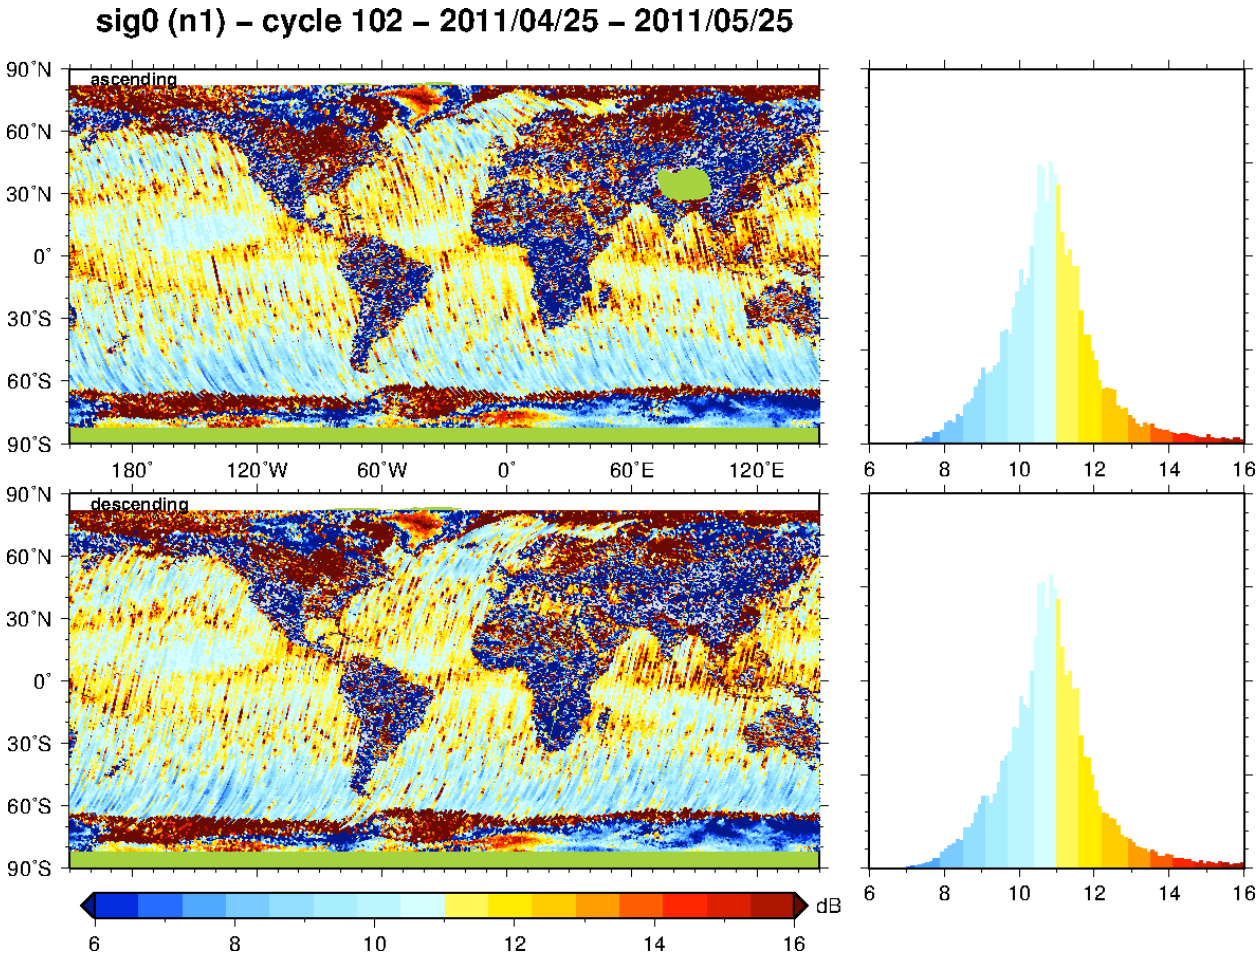

Figure 14. EnviSat σ_0 in dB.

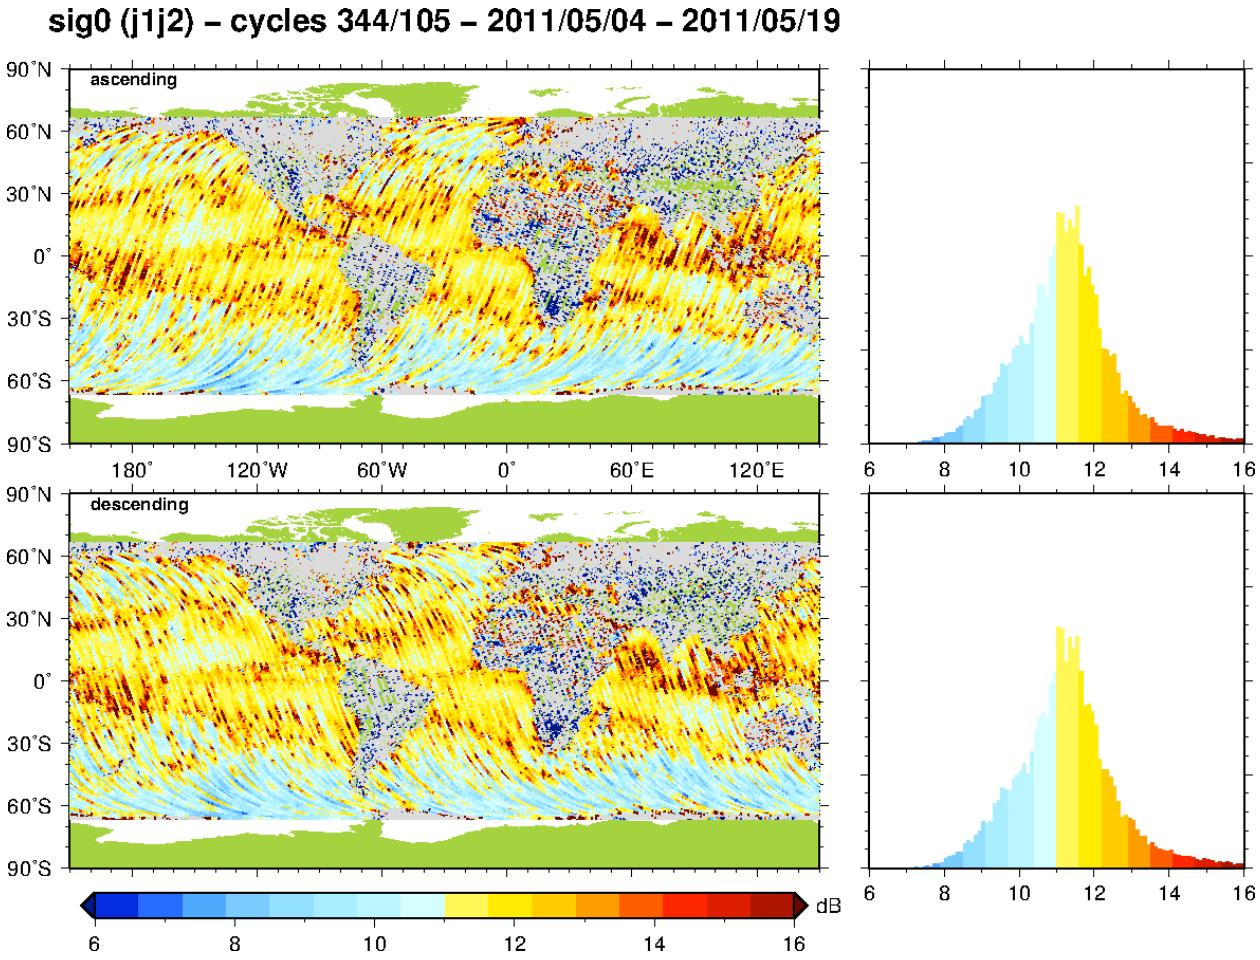

Figure 15. Jason-1 and -2 σ_0 in dB.

Variations in 20 Hz data over one second

The data are retracked at 20 Hz, and then 1-Hz averages are formed. We examined the standard deviation of the 20 Hz values around the 1-Hz average, for SWH, σ_0, and raw sea surface height.

Significant wave height

Histograms look similar for all three satellites. In map view it appears that EnviSat's standard deviation of SWH grows with SWH by more than it does for CryoSat or Jason-1 and -2. The CryoSat performance appears similar to the Jason performance.

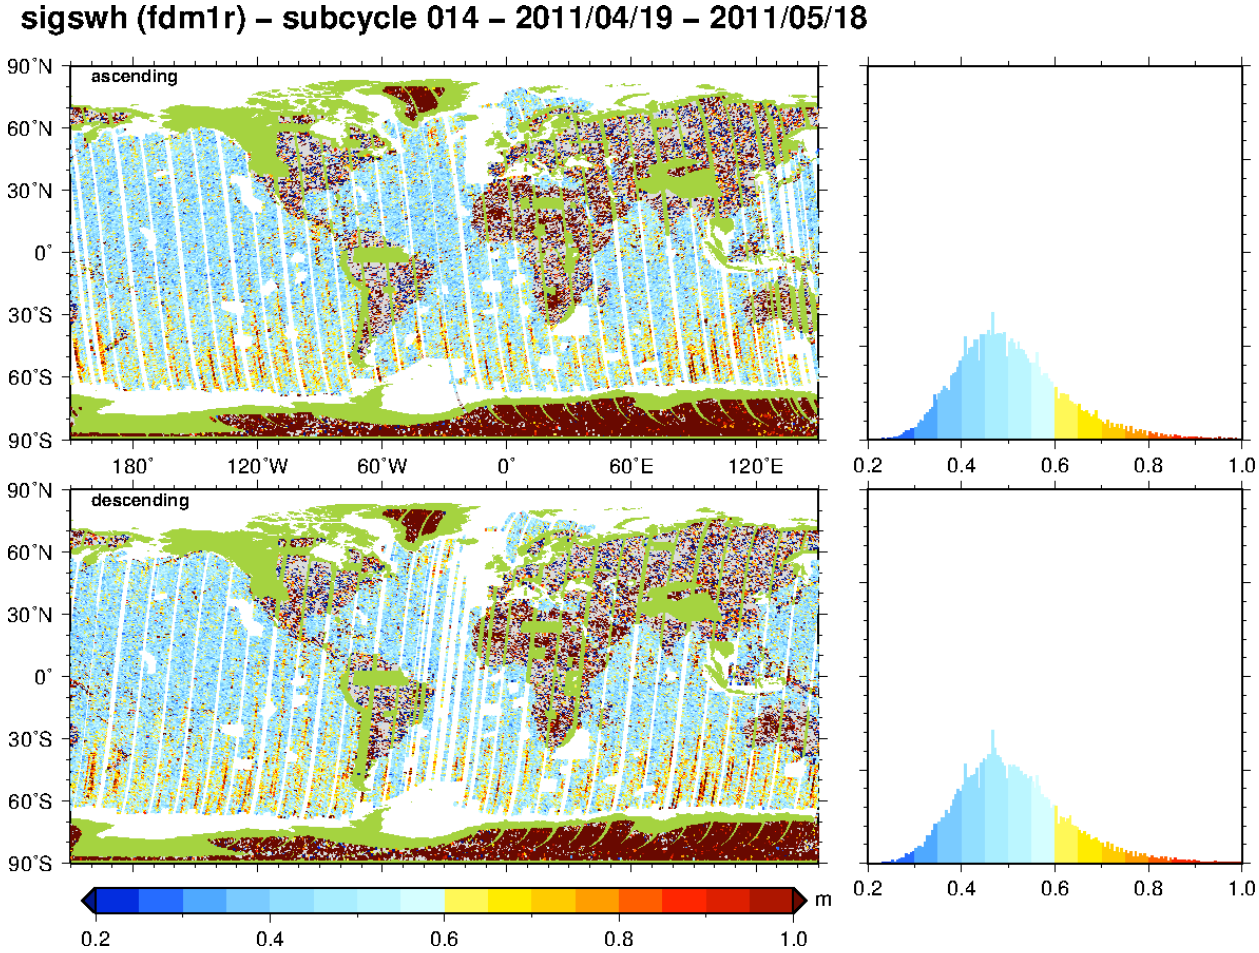

Figure 16. Standard deviation in Cryosat 20 Hz SWH estimates over one second.

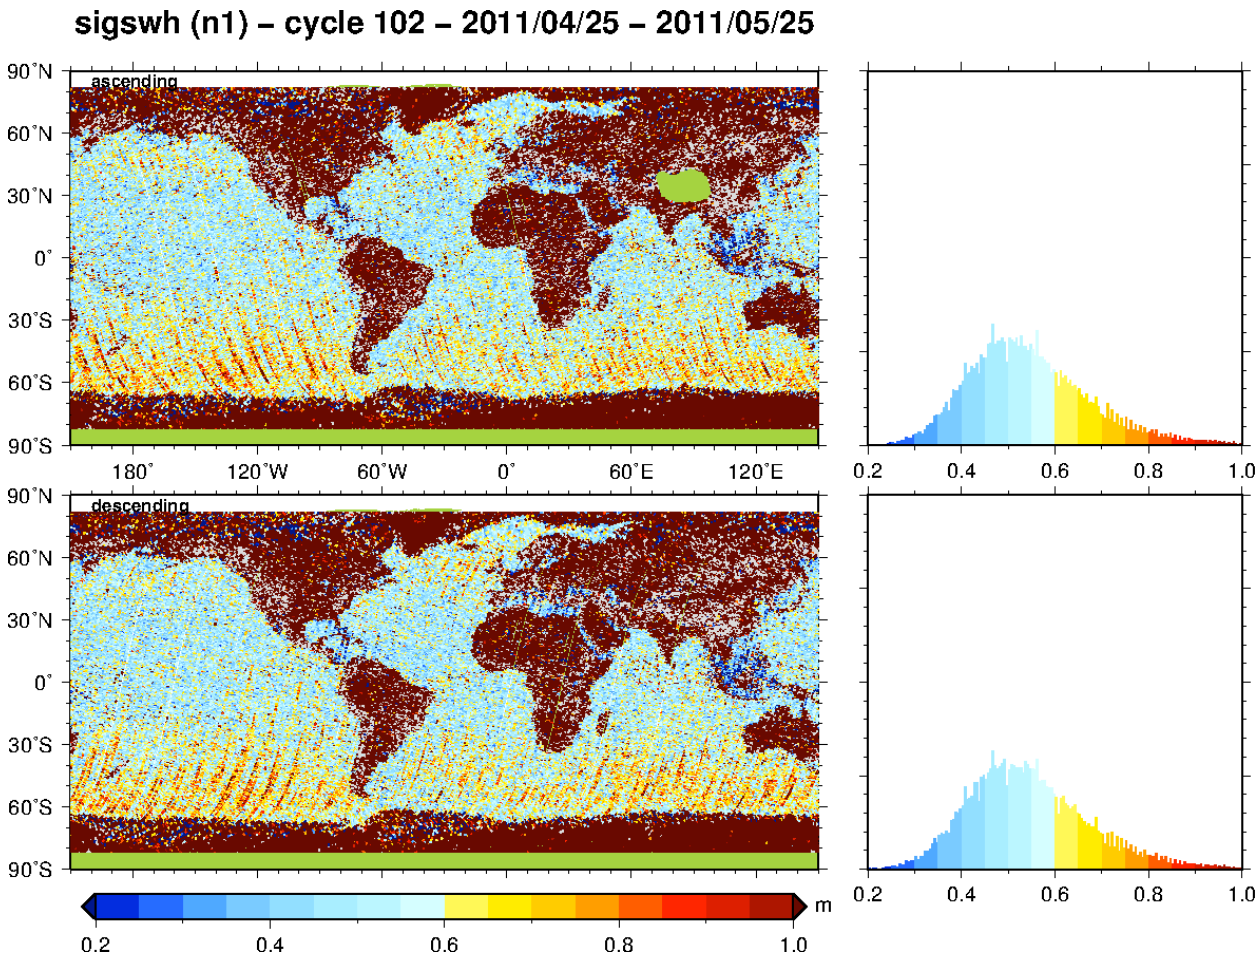

Figure 17. Standard deviation in EnviSat 20 Hz SWH values over one second.

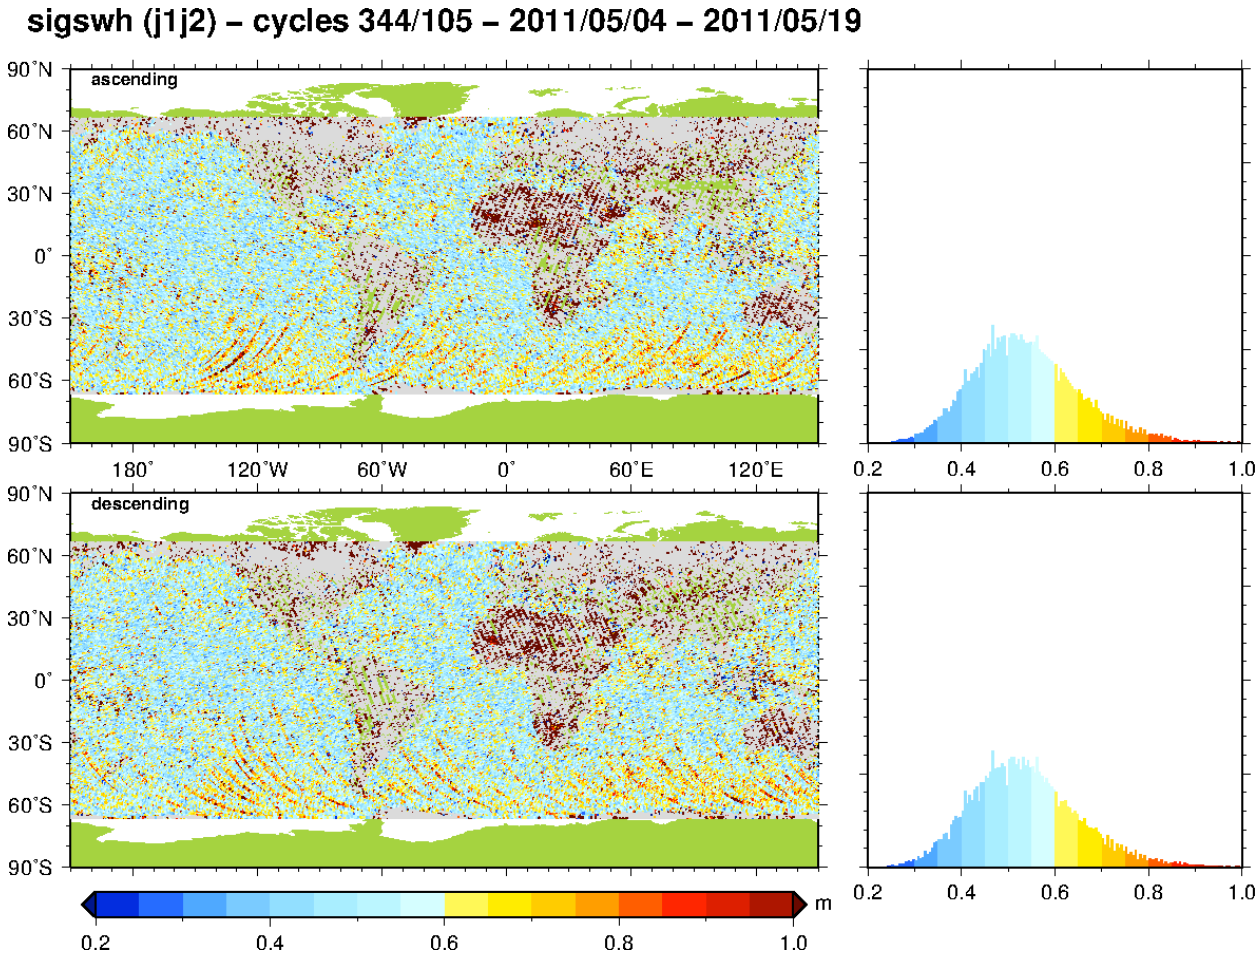

Figure 18. Standard deviation in Jason-1 and -2 20 Hz SWH values over one second.

Sea surface height anomaly

The 1 Hz standard deviation is the root-mean-square of the residuals after a least-squares straight line is fit over one second to the 20 Hz values for orbit height minus retracked range. Dividing the values shown here by the square root of 20, one obtains the typical one-second range precision quoted for altimeters.

CryoSat shows the best performance of all the altimeters compared here, with only a small increase in the standard deviation where SWH is large, and a peak in this histogram between 6 and 7 cm. This is equivalent to a one-second range precision around 1.5 cm. Jason-1 and -2 look similar, though the histogram peak may be slightly higher, around 7 cm. Envisat shows a longer tail on the distribution and a larger increase of standard deviation with SWH.

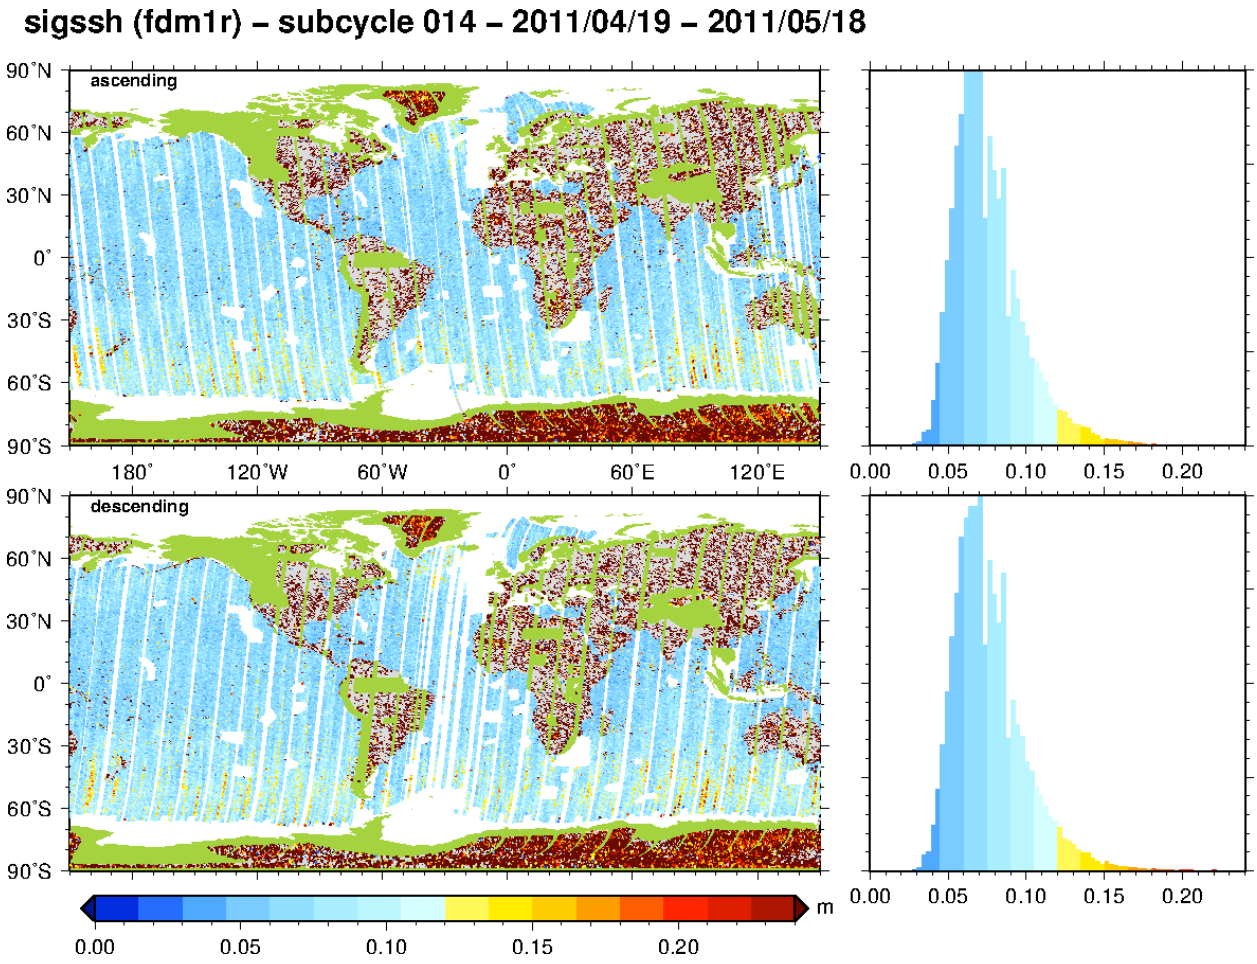

Figure 19. Standard deviation in CryoSat 20 Hz retracked sea surface height during 1 second.

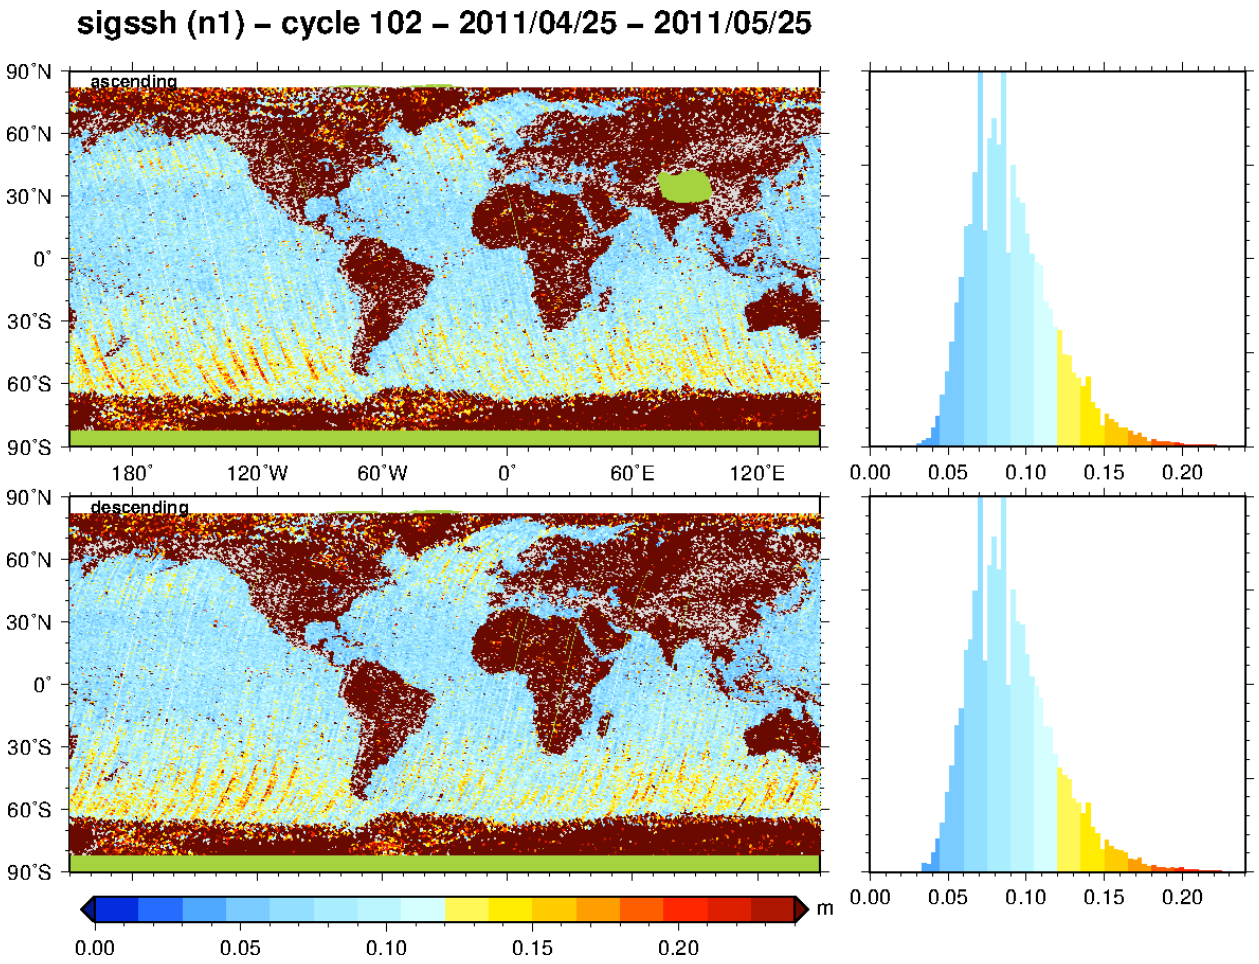

Figure 20. Standard deviation in EnviSat 20 Hz sea surface height during 1 second.

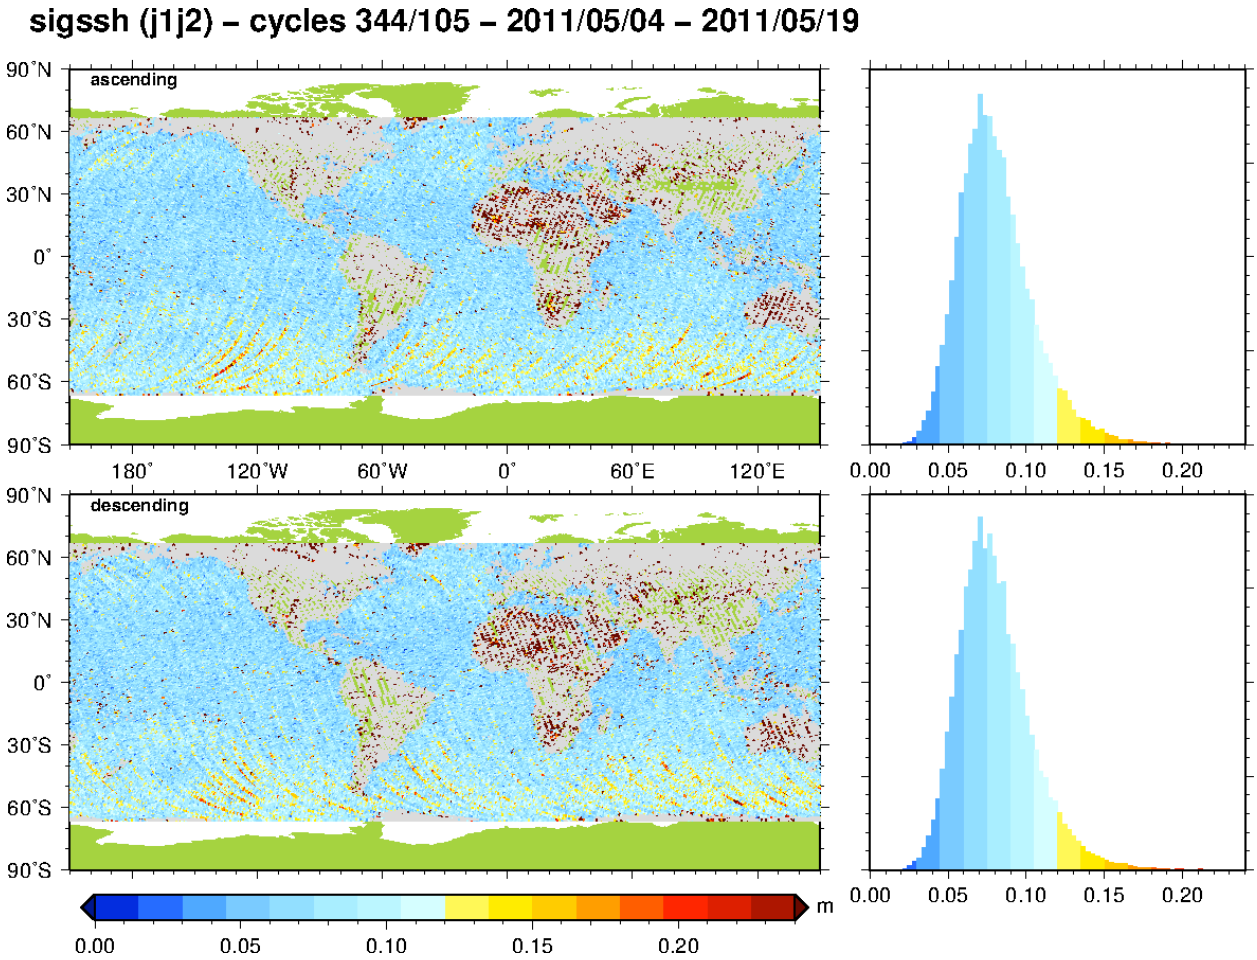

Figure 21. Standard deviation in Jason-1 and -2 20 Hz sea surface height during 1 second.

Sigma-nought

The variation in the 20 Hz σ_0 values taken over one second is quite different among the satellites. CryoSat shows the largest variations. Jason-1 and -2 are somewhat similar. EnviSat shows very small variations. The geographical pattern appears similar in all satellites, with increases in the variation perhaps correlated with areas of rainfall.

We do not understand why the distribution for CryoSat is so much wider than the others, or why that for EnviSat is so much narrower. Perhaps the damping employed in the Automatic Gain Control loops is different, or the loops target different waveform gates in the different satellites. Retracking should compensate this, unless there is clipping of saturated waveforms. Perhaps we increase the noise in σ_0 by applying corrections to AGC for amplitude, SWH and mispointing parameters fit during retracking. However, the map pattern of CryoSat σ_0 standard deviations does not show the pattern of the off-nadir angle or SWH. The distribution here remains unexplained.

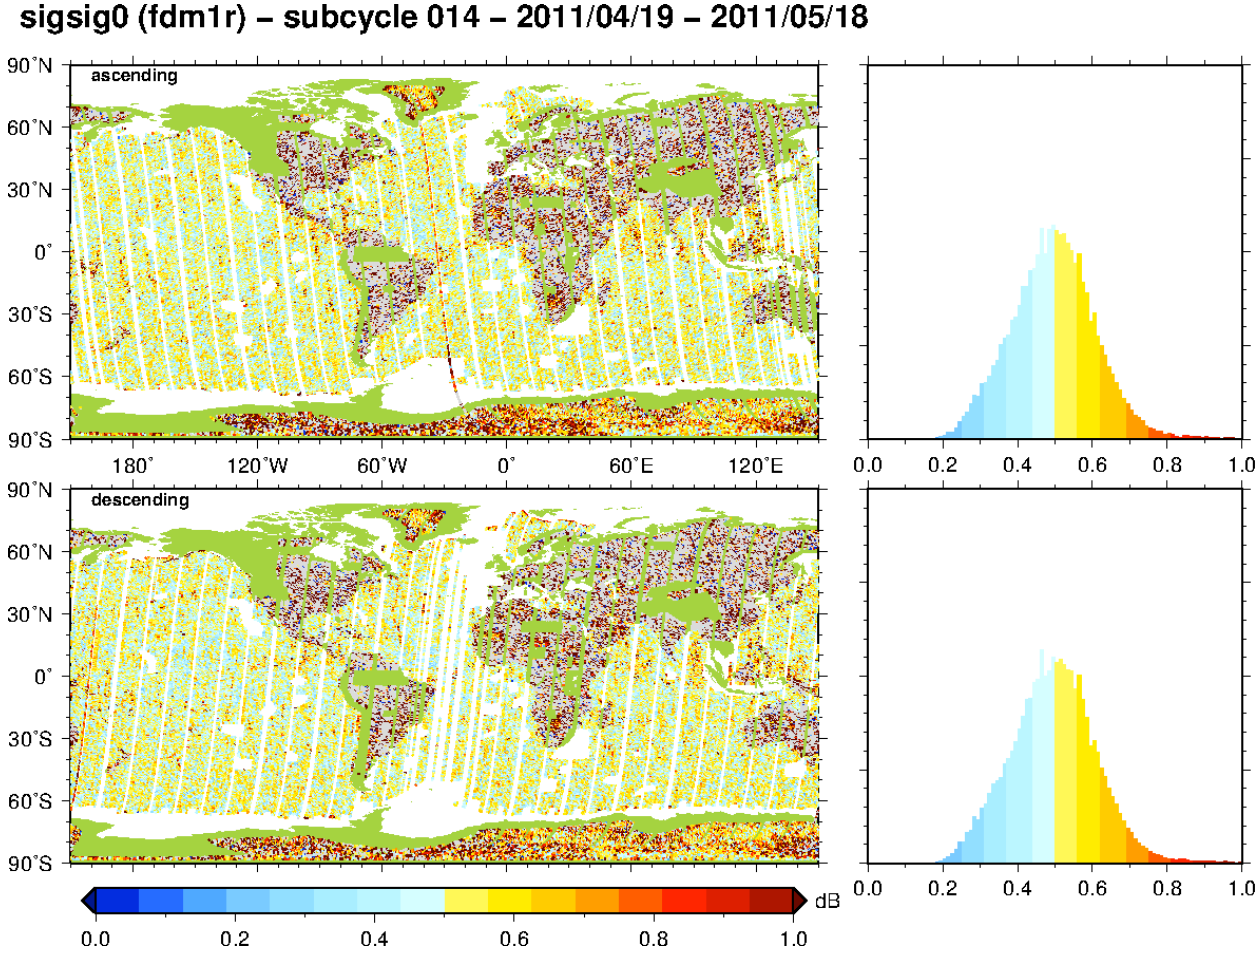

Figure 22. Standard deviation in CryoSat 20 Hz retracked σ_0 estimates during 1 second.

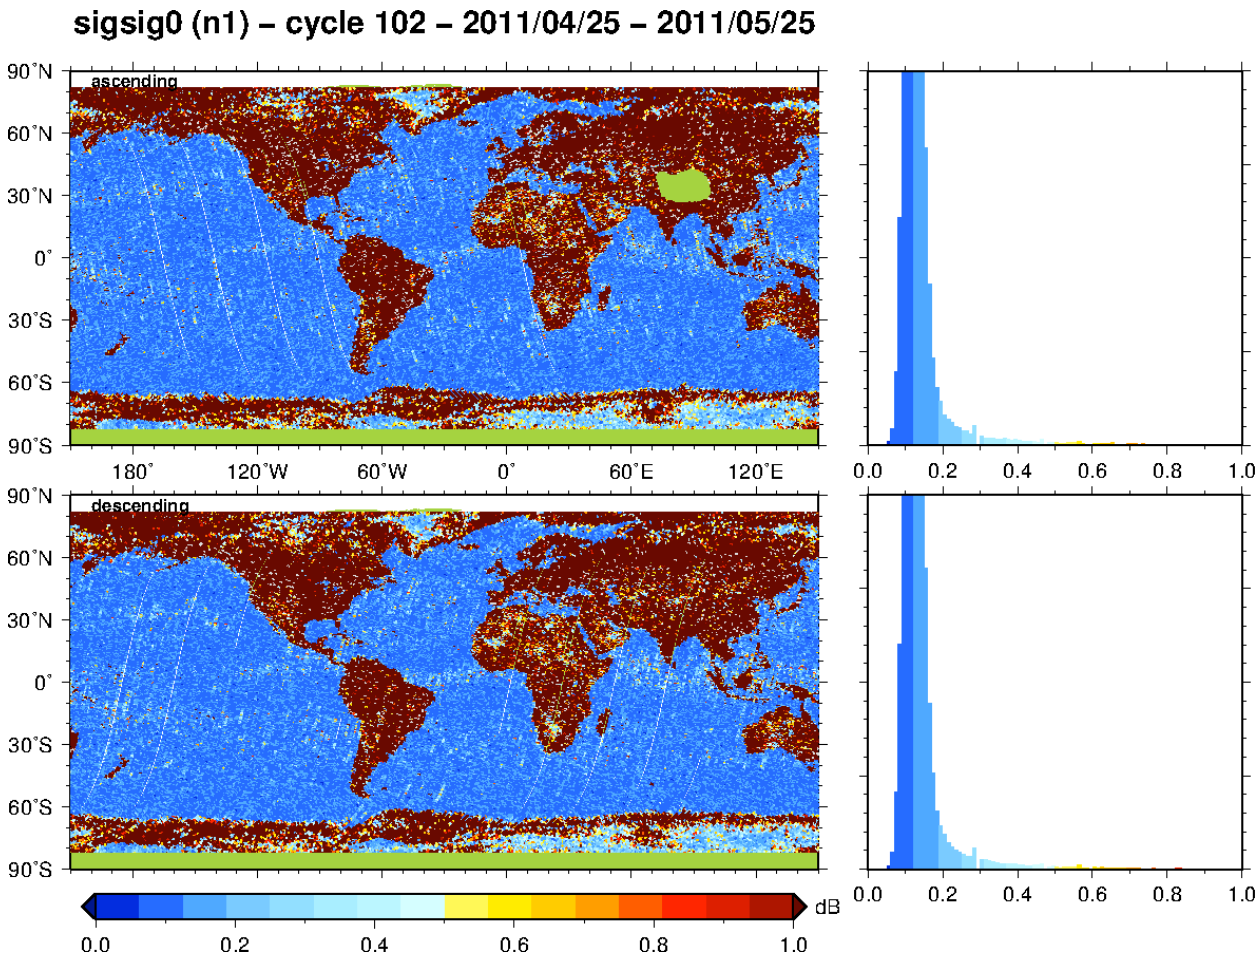

Figure 23. Standard deviation in EnviSat 20 Hz σ_0 values during 1 second.

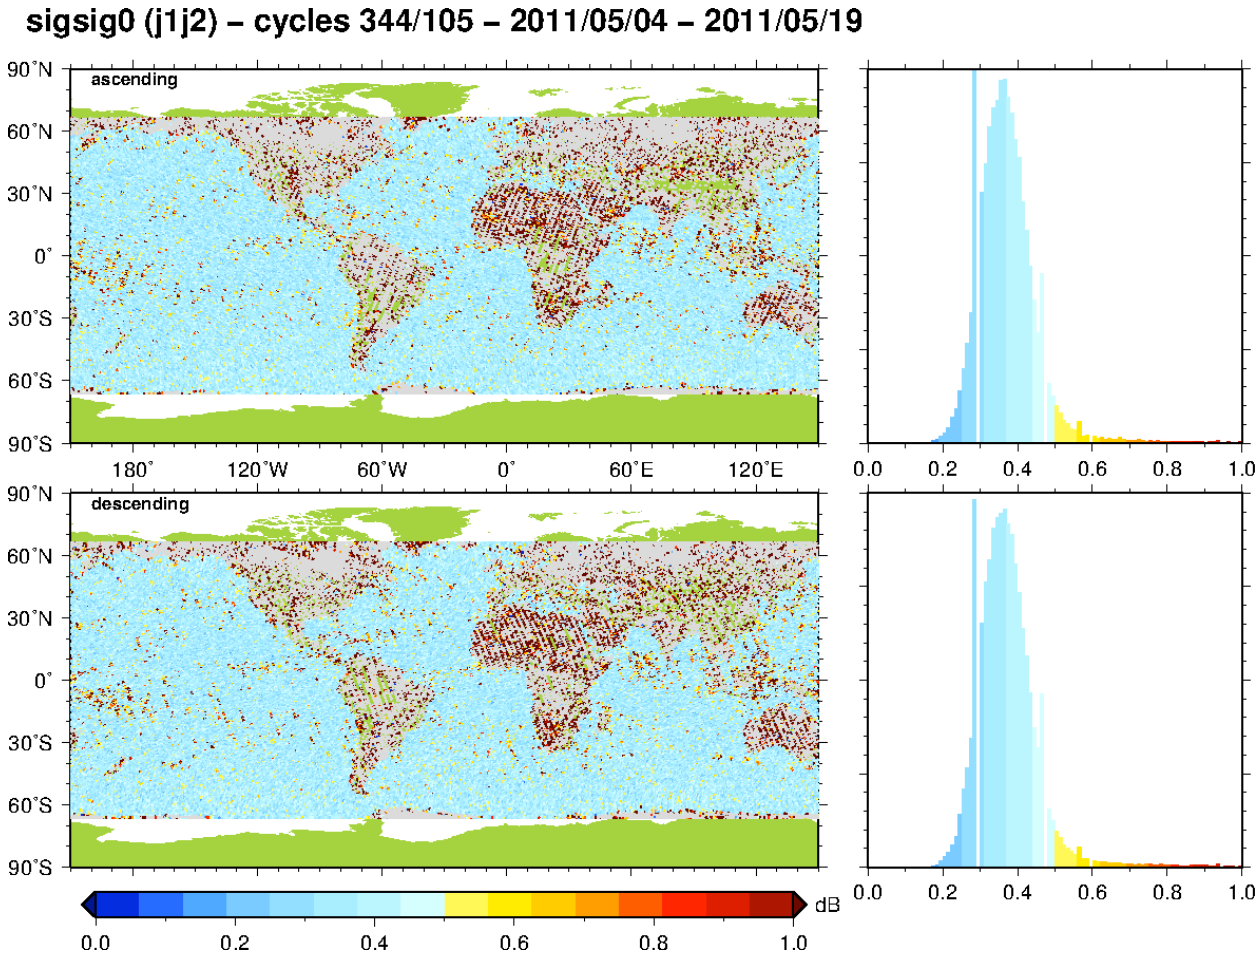

Figure 24. Standard deviation in Jason-1 and -2 20 Hz σ_0 values during 1 second.

The number of valid 20-Hz values in a 1-Hz average

We also checked that the above problem in standard deviation was not due to an anomaly in the number of retrackable 20-Hz samples per second. We found this quantity to be substantially similar in CryoSat and EnviSat, being nearly always 20, and when not 20, nearly always 19. Jason-1 and -2 had values less than 19 occurring more frequently, perhaps due to additional editing criteria applied after retracking. Our only criterion for CryoSat was that the retracker converged normally.

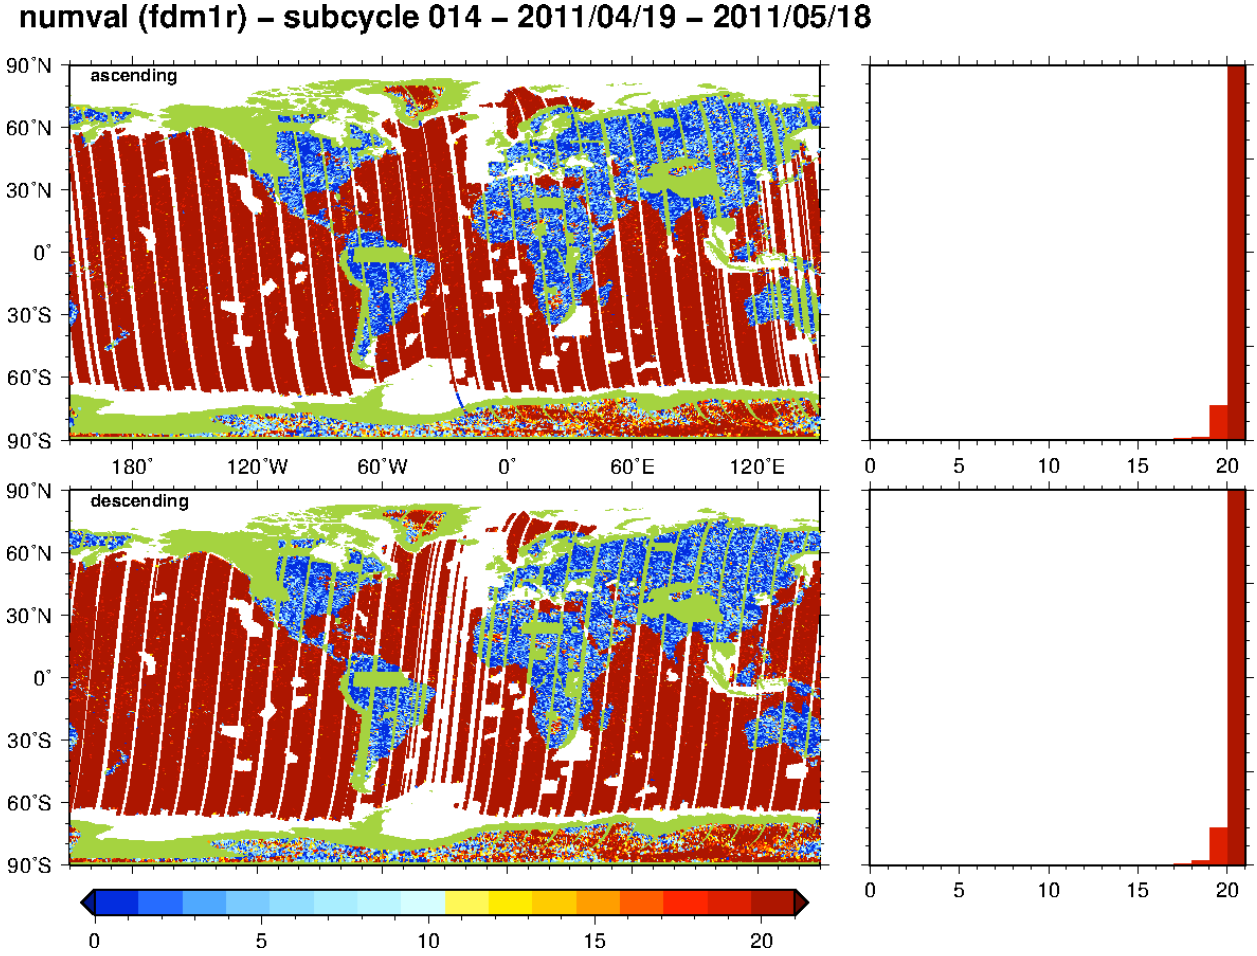

Figure 25. Number of retrackable 20-Hz CryoSat waveforms per second.

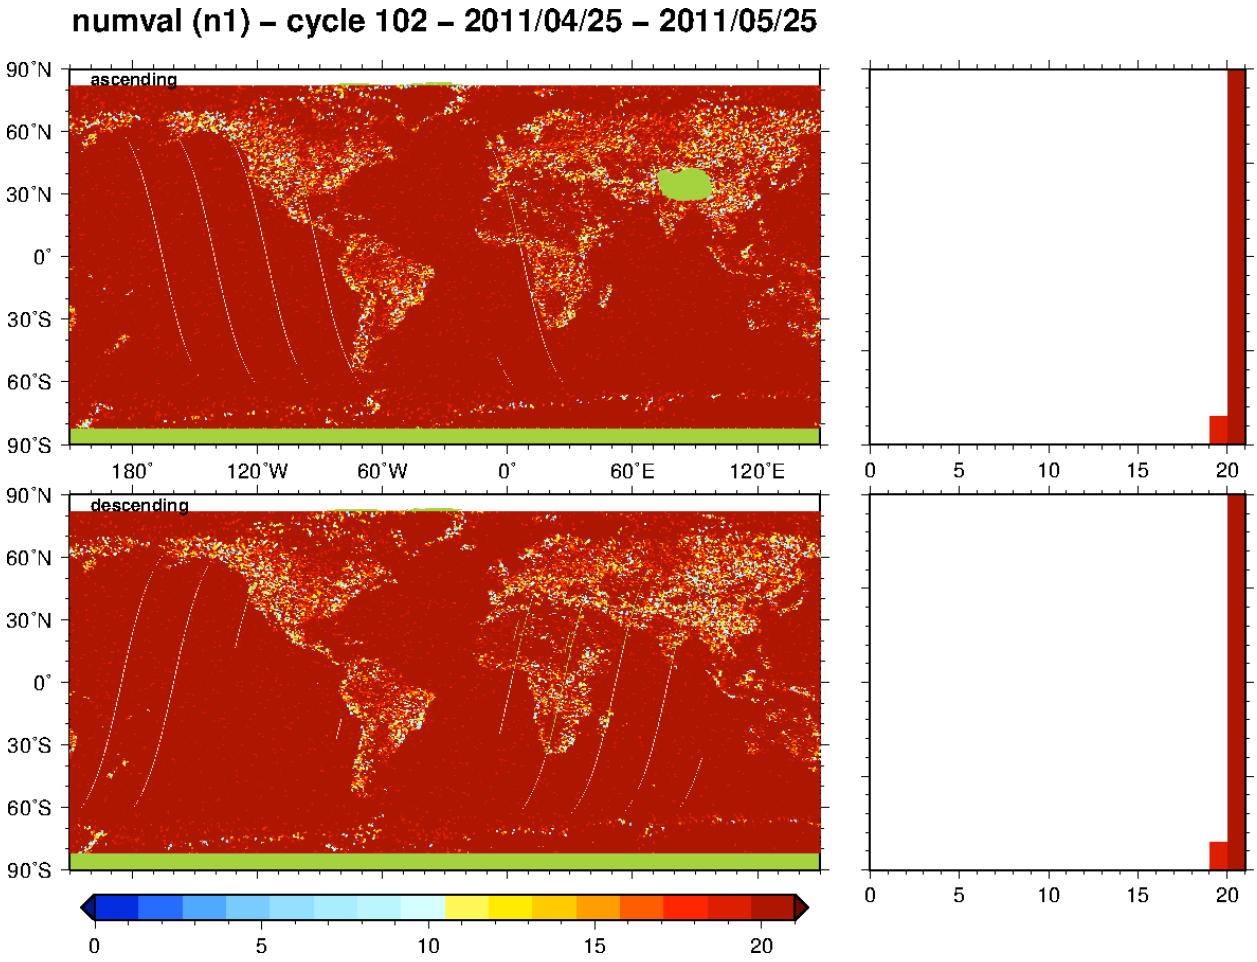

Figure 26. Number of useable 20-Hz EnviSat values per second.

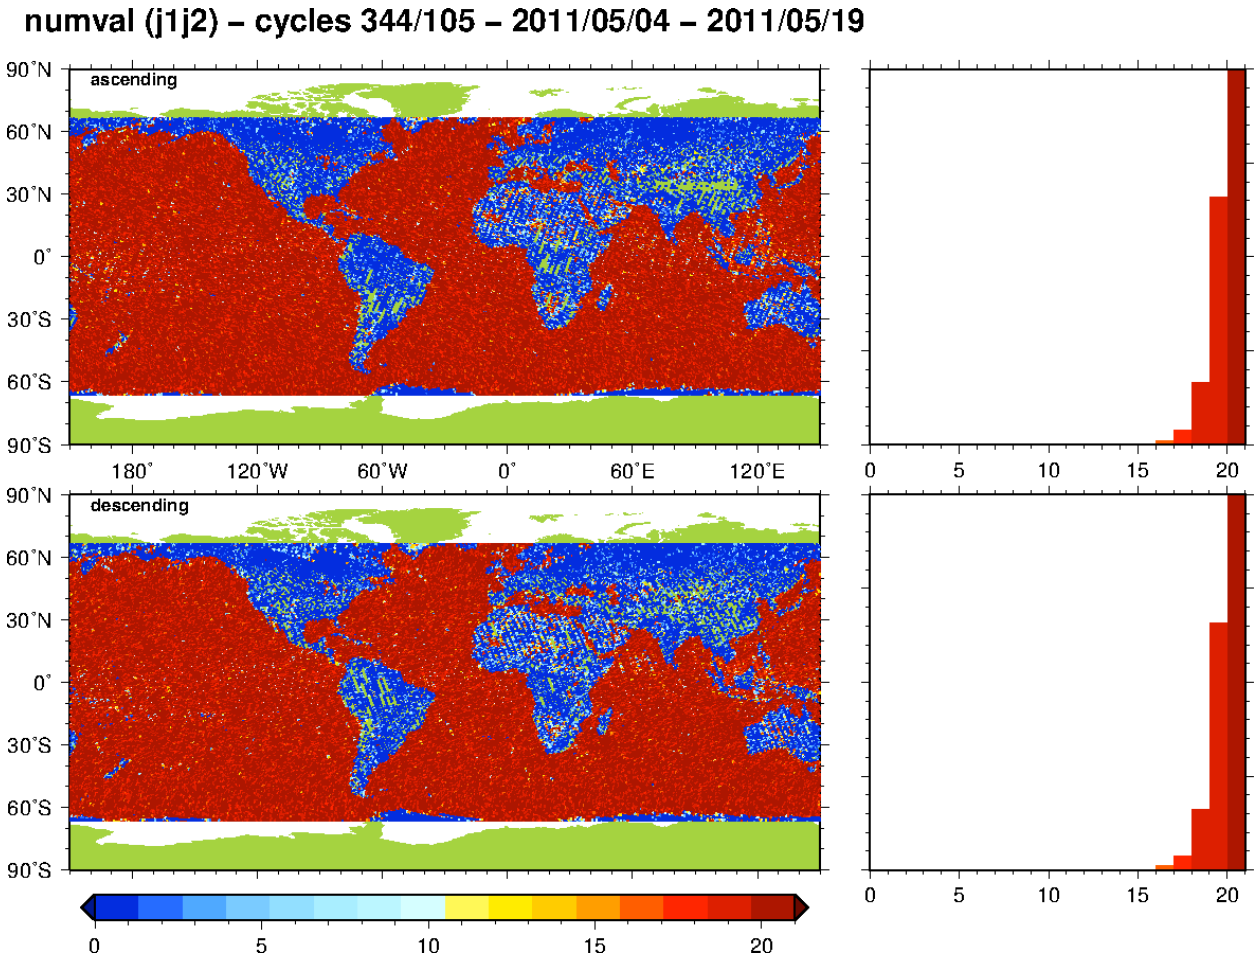

Figure 27. Number of useable 20-Hz Jason-1 and -2 values per second.

Additional illustrations from CryoSat

Automatic Gain Control (AGC)

The sigma-naught values are derived by applying adjustments to the AGC values supplied in the FDM or L1b data for receiver channels 1 and 2. We inspect the Instrument_Mode_ID to obtain the receiver channel to use. The adjustments we apply derive from the variation of orbital altitude relative to a nominal altitude, and from the retracking, which determines adjustments for SWH, mispointing, and the actual fitted amplitude relative to a nominal amplitude. All adjustments are expressed relative to nominal conditions, thus the mean value of the difference between corrected AGC and true σ_0 should reflect 10 $\log_{10}$ of the nominal system gains and losses when the AGC control loop is following nominal ocean returns at nominal satellite altitude. Shown below is the one-second average of CryoSat AGC *without* the adjustments.

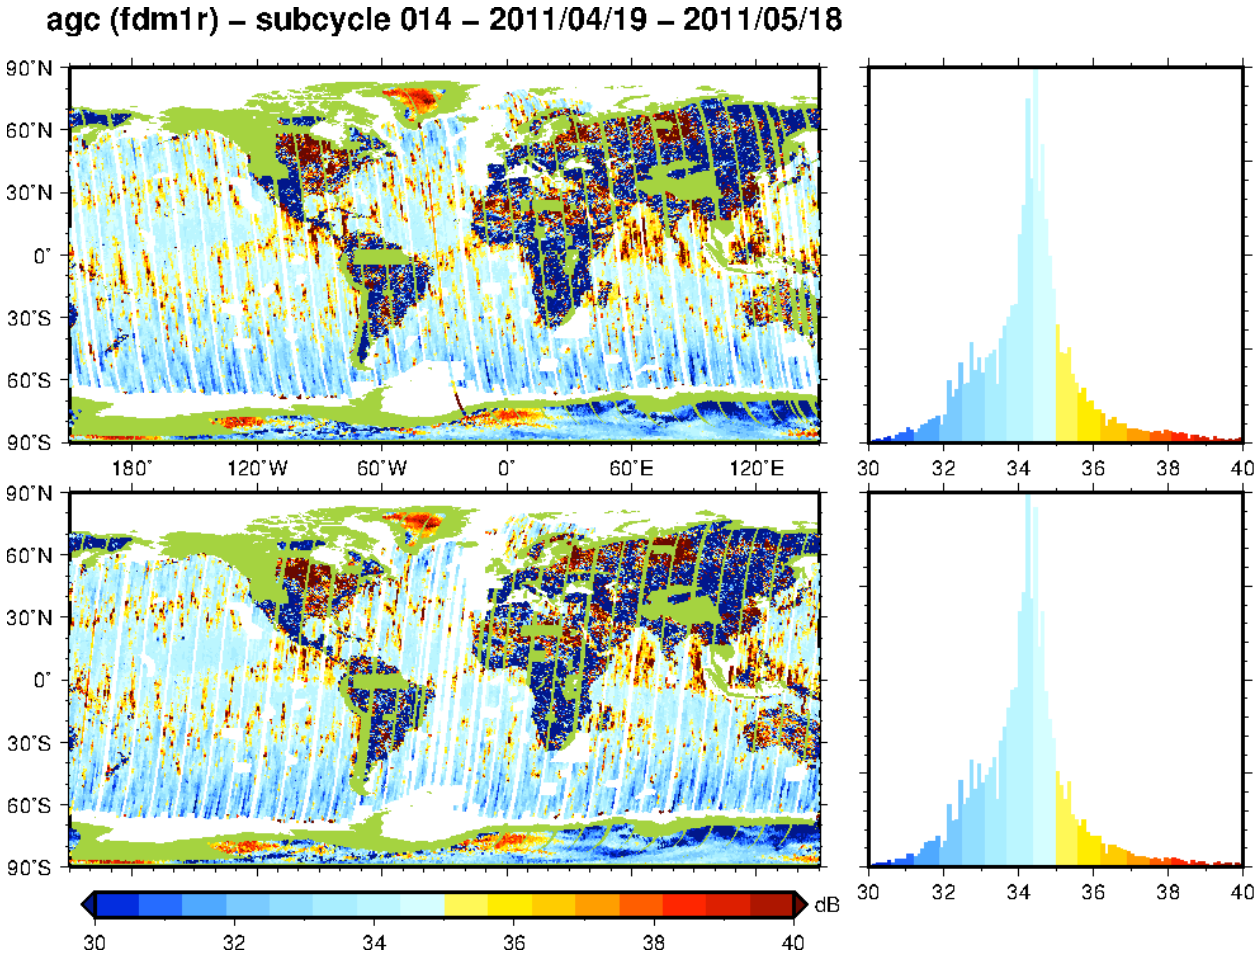

Figure 28. CryoSat AGC before adjustment for altitude and retracking effects.

Mean Quadratic Error

The summed squared misfit of the fitted waveform model to the waveform data, normalized by the fitted waveform power, is known as the mean quadratic error (MQE) in the MLE3 and MLE4 retrackers used in Jason-1 and -2. We have computed the same parameter here, for our retracking of CryoSat. We do not see any systematic or geographical signal in the MQE. Rather, it looks random, as it should. There may be a slight increase with rainfall events.

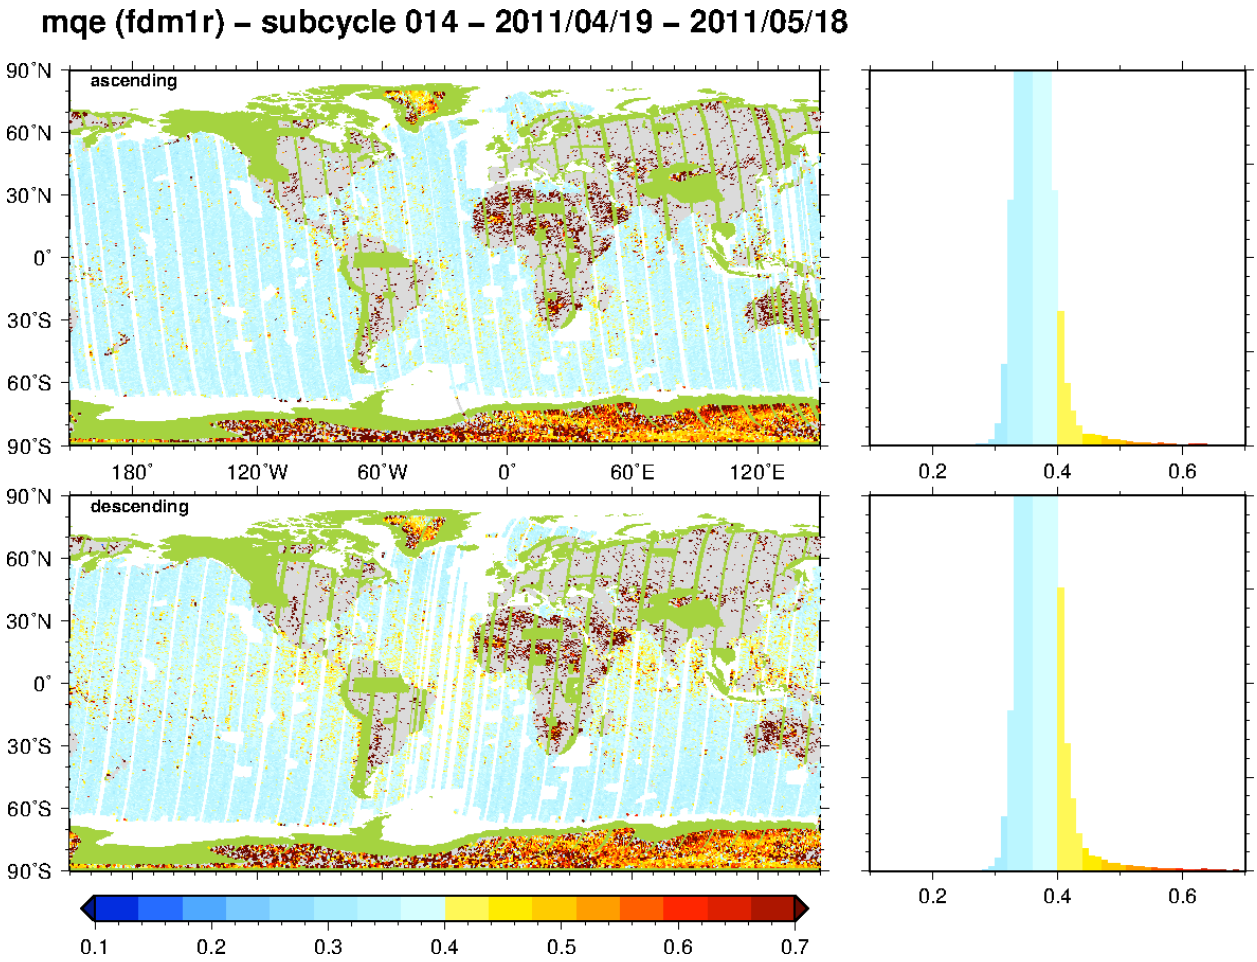

Figure 29. Misfit (MQE) of the retracking model to the CryoSat waveform data.

Convergence

We follow standard MLE retracker practice by testing for convergence only after 3 or more iterations, and limiting the total number of iterations to 25. If the retracker's model fitting scheme reaches 25 iterations we consider this a failure to converge. If the change in MQE over 3 consecutive iterations is less than 5×10^{-4}, then the model is said to have converged; this is also MLE practice. Shown below are the number of iterations needed to converge the waveform model fitting, for the first 20 Hz waveform in the 20-per-second group.

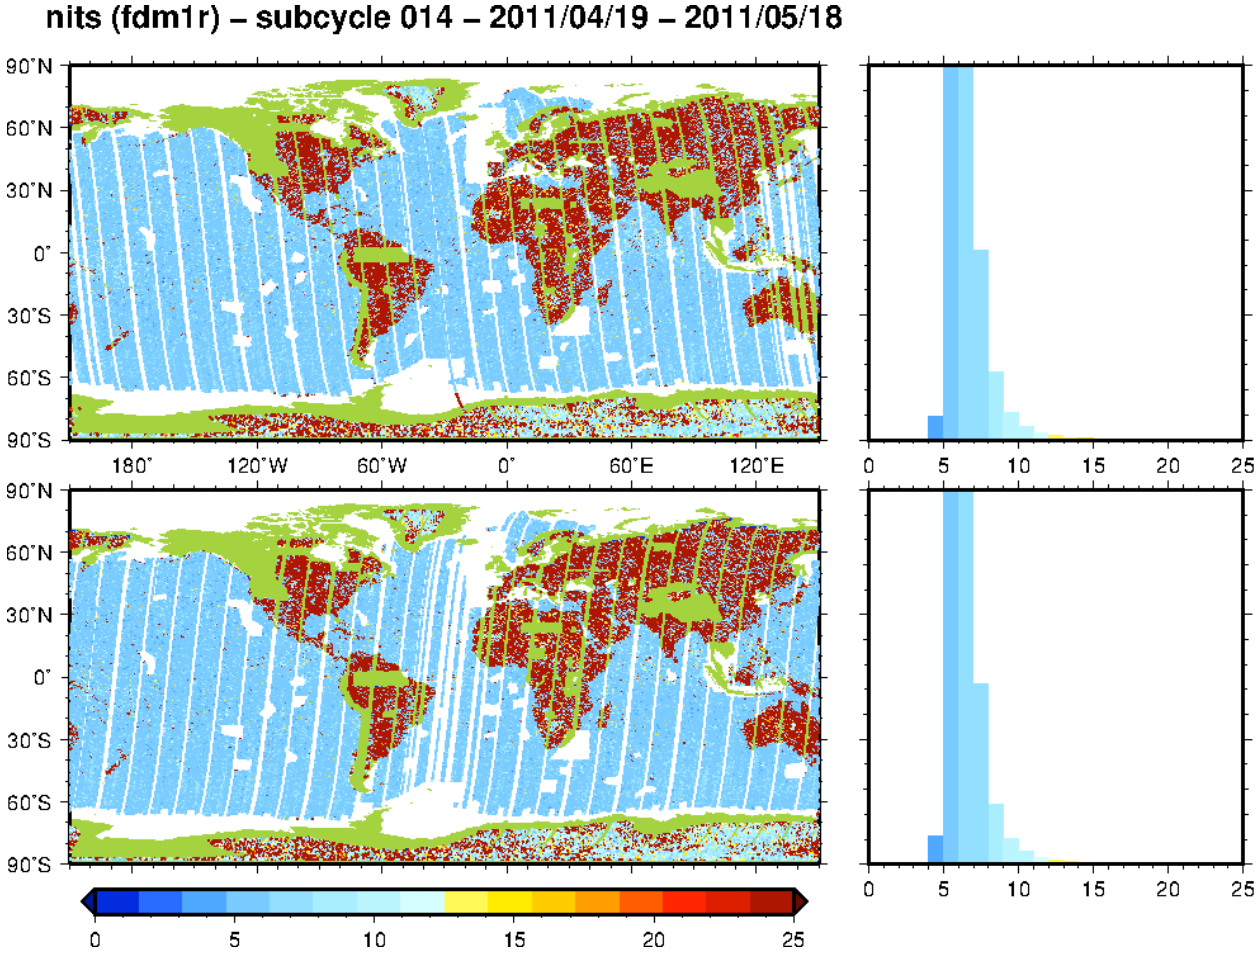

Figure 30. Number of iterations needed for convergence of our CryoSat retracker.