ATSR Images: Frequently Asked Questions

Dr Chris Mutlow

A guide to the features users often report finding in their ATSR-1 and -2 data products

1.0 Introduction

The purpose of this guide is to help users of the ATSR gridded brightness temperature products diagnose the cause of the most frequently encountered "features" and artefacts observed in displayed ATSR images. It is supplementary to the information found in the technical note describing the ATSR data formats (ER-TN-RAL-2164 "SADIST-2 v100 Products by Paul Bailey") available from http://www.atsr.rl.ac.uk.

The queries that reach the project team fall into three broad categories: 1) those resulting from incorrect interpretation of the ATSR data formats, 2) those that result from users not understanding the "features of the ATSR instruments and their data sets, and 3) those that result from instrument anomalies. Each of these types of query is dealt with in a different section of the document. An explanation of the "feature" and its cause is provided, and, if appropriate, a suggested course of action to solve the problem is recommended.

2.0 Problems with the interpretation of ATSR data product formats

2.1 Byte Swapping (or why is does my image look like random noise?)

The user usually reports something like "my image contains rubbish", or "the image on my screen just looks like noise". This is by far and away the most frequent problem encountered by users of ATSR image data, despite the copious information supplied in the product documentation. An image affected by this problem usually looks something like the scene shown in Figure 1, and contains no discernable features. The image has this appearance not because the fundamental data contained in the image file are really corrupted, but, in most cases, because the data elements within the image have been interpreted incorrectly by the user.

FIGURE 1. Typical example of an ATSR image displayed in the wrong byte order

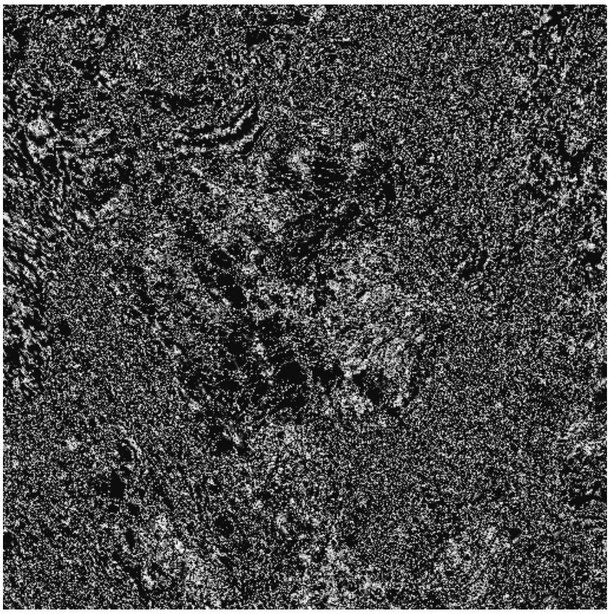

The apparent corruption of the image results from the pixel values in the product being used in a byte-order that is inappropriate for the computer system on which the data are being analysed. Some computers interpret the bytes within integers such that the bytes are given increasing significance, whilst others interpret them such that the bytes are given decreasing significance. These systems are known as little-endian and big-endian, respectively. Thus, when representing the same integer value, these two classes of systems use a different internal byte-ordering; so when given data with the same byte-ordering, each class of system returns different numerical values. Therefore, using the data in a byte-order inappropriate for the particular system you are using results in the most significance being given to what is intended to the least significant byte of the data value. Images interpreted in this incorrect way usually contain no obviously discernable features, and have the appearance of random noise. This is not surprising as they reflect the random variations in the least significant byte of the image pixel value.

An example of an image displayed with the incorrect byte-ordering is shown in Figure 1. The “real” image contained within the data can be recovered by the user applying the correct byte swapping to the data. In the case of 2-byte integers this is straightforward as it is only necessary to reverse the order of the byte in each integer word, for a 4-byte integer it is more complex and the user is referred to the example C programs in the SADIST-2 format document.

SADIST-2, the software that generates ATSR-1/2 image products, is a OpenVMS application. Therefore, as VMS and the hardware on which it runs uses little-endian integer formats, the bytes within two-byte words and four-byte words are stored in increasing significance. If SADIST-2 products are to be read on big-endian systems (such as UNIX running on a Sun or SGI machine) the byte-ordering must be reversed. That is, the internal representation must be changed so that the intended value will be retrieved correctly on a big-endian system.

A mechanism is provided in all ATSR products whereby the correct byte-ordering of the data can be established automatically from the product header. The first two bytes within each SADIST-2 product header are fixed (i.e. they are the same in all products), and can be used to determine the byte-ordering on the local system. Like the rest of the product headers, the first two bytes are ASCII: they are the ASCII codes 65 and 66, which represent the characters A and B, respectively. If these bytes are interpreted by the reading software as a two-byte integer, the result will differ on little-endian and big-endian systems, and can be used to test whether byte-swapping is necessary.

On a little-endian system, the two-byte integer evaluates as:

$$65 + (66 * 2^8) = 16961$$

whilst on a big-endian system, the two-byte integer evaluates as:

$$(65 * 2^8) + 66 = 16706$$

The the correct value of the 2-byte integer is 16961, if the value 16706 is returned the data need to be byte-swapped before use.

If the image you are displaying at the moment looks like the example in the figure, then before you contact the ATSR Project Team for help please use the test described above to check that you have selected the correct internal representation of the data.

2.2 Blanking Pulses (or why is my 0.87/12 µm image covered in noise?)

The second most frequent query from users is when they report that a strange noise pattern is affecting their 0.87 µm and/or 12.0 µm images, and that all the other channels look perfect. A typical example of an affected image is shown in Figure 2.

In this case the “noise” pattern on the 0.87 µm and 12.0 µm images results from the user ignoring the flag data that are carried on the pixel values for these channels. The 0.87 µm and 12.0 µm pixel data include quality flags that apply to all the image planes within the products, but which corrupt the images from the channels that carry them if these flags are not interpreted correctly. Using the pixel data to carry the flags in this way does complicate the data format, but a considerable amount of storage space is saved by using this simple flagging method rather than employing a separate array of flag data in the product.

These flags are needed because some of the other instrument payload carried by the ERS satellites are “active” sensors that transmit radio frequency pulses. Such transmissions can be “picked-up” by other electrical circuits on the satellite, and therefore have the potential to disrupt the data collected by ATSR and other instruments. To guard against this eventuality the satellite’s data system informs each instrument whenever an RF transmission is being made. This information is passed through a series of “event flags” known in ESA’s jargon as “blanking pulses”.

Within the satellite data system there are separate “blanking pulses” produced by each of the “active” instruments to indicate their RF transmission status, and to ensure that any potential corruption of ATSR data by these RF transmissions can be readily identified. The ATSR data processing software (SADIST-2) annotates all the images it produces with “blanking pulse” flags. To save space in the ATSR products a “composite blanking pulse” is used based on a “logical OR” of the individual instrument “blanking pulses” supplied in the platform telemetry. This “composite blanking pulse” does not identify which instrument was transmitting, merely that an instrument was transmitting during the collection of an image pixel. It would be very wasteful of disk space to provide this 1-bit “blanking pulse” information as a separate flag image plane within each ATSR GBT product. Thus, to convey this information in a compact form the values of the 0.87 and 12.0 µm pixels within the product are set negative to indicate whenever the “composite blanking pulse” is set. Put more simply, -ve pixel values (less than -8) in the 0.87 and 12.0 µm channels indicate that an “active” instrument transmitted during the collection of the ATSR pixel data at that location.

These “blanking pulse flags” apply to the pixel data from ALL the ATSR channels collected at that time not just those channels which carry the flags.

In practice no anomalous signals that can be attributed to the “pick-up” of RF transmissions on the satellite have been detected in either the ATSR-1 or -2 data sets, so most users can safely ignore these flags. However, the flags are retained in the data format in case this situation should change due to the satellite and instruments ageing.

FIGURE 2. A typical example of a displayed image showing the “blanking pulses”

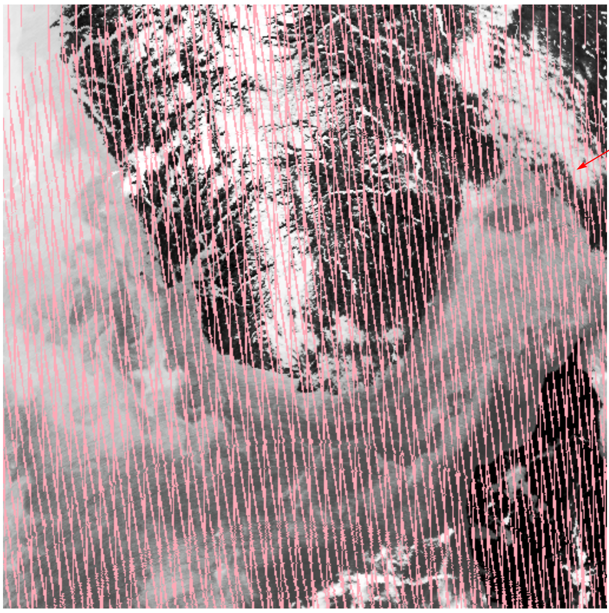

To eliminate the interference caused by “blanking pulses” users should take the absolute value of the pixel values before they are used, or otherwise scaled for display. In IDL this means typing “TV, BYTSCL(ABS(IMAGE))” rather than “TV, BYTSCL(IMAGE)”. To be strictly correct users should also remember to take account of the exceptional pixel flag values (i.e. pixel values in the range (-1 to -8)) instead of just ignoring them as these flags provide essential data quality information.

Users with this “blanking pulse” problem are also usually affected by a similar problem present in the 0.65 µm and 10.8 µm images, but often don’t know as is much harder to see the effects of ignoring the flags contained on these channels than it is in the case of the 0.87 µm and 12.0 µm images (see Section 2.3 on page 5 for further explanation of this).

2.3 Cosmetic Fill (or why is my 0.65µm/10.8 µm image covered in noise?)

Questions about “cosmetic fill” are less frequent than those about “byte ordering” and “blanking pulses” because its effects are fairly difficult to see in most nadir view images, although on inspection of most forward view data they are immediately apparent.

Most users who notice the effects of “cosmetic” fill report unusual noise in forward 0.87 µm and 12.0 µm view images, but are not aware of the problems also present in the nadir data. The reason for this becomes obvious when you look at the following sample images illustrating the effects of “cosmetic” fill. Figure 3 shows the effect on the nadir view, and Figure 4 similarly shows a forward view image. Figure 4 clearly shows the undersampling that takes place in the forward view, and the effects of the cosmetic fill.

FIGURE 3. Example of the effects on cosmetic fill on a typical nadir view image

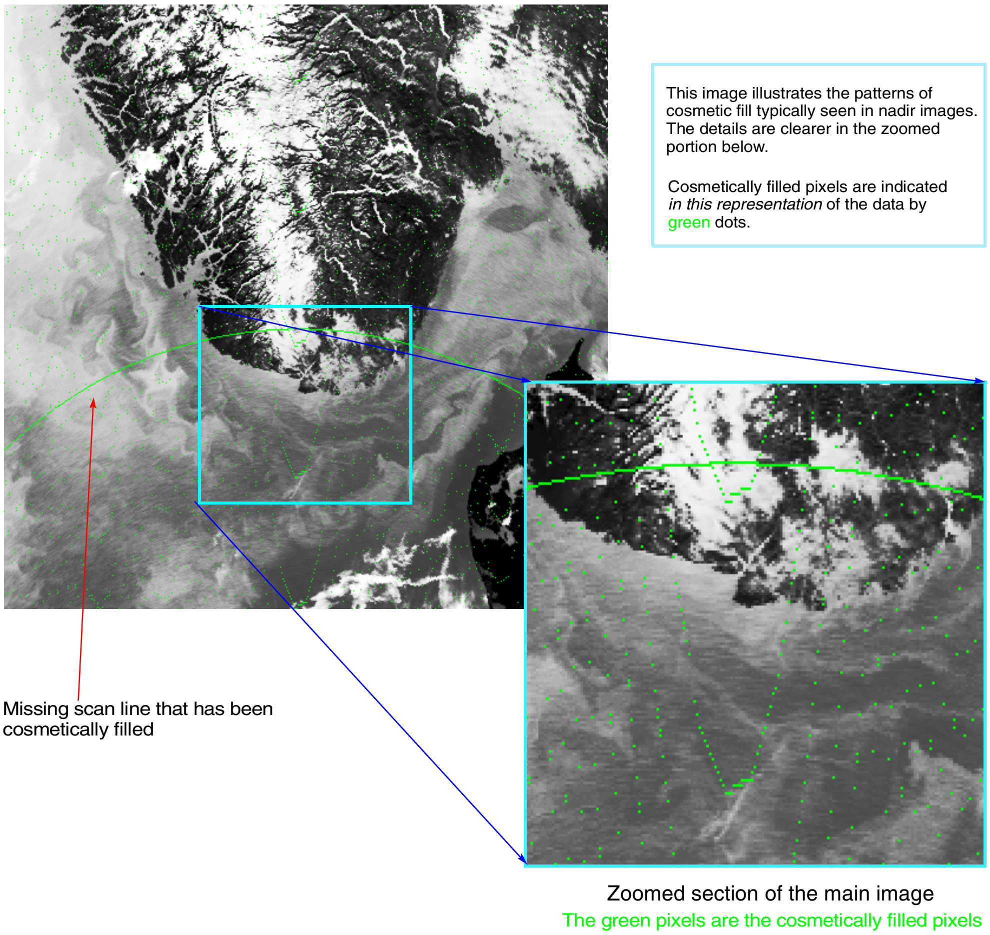

FIGURE 4. Example of the effects on cosmetic fill on a typical forward view image

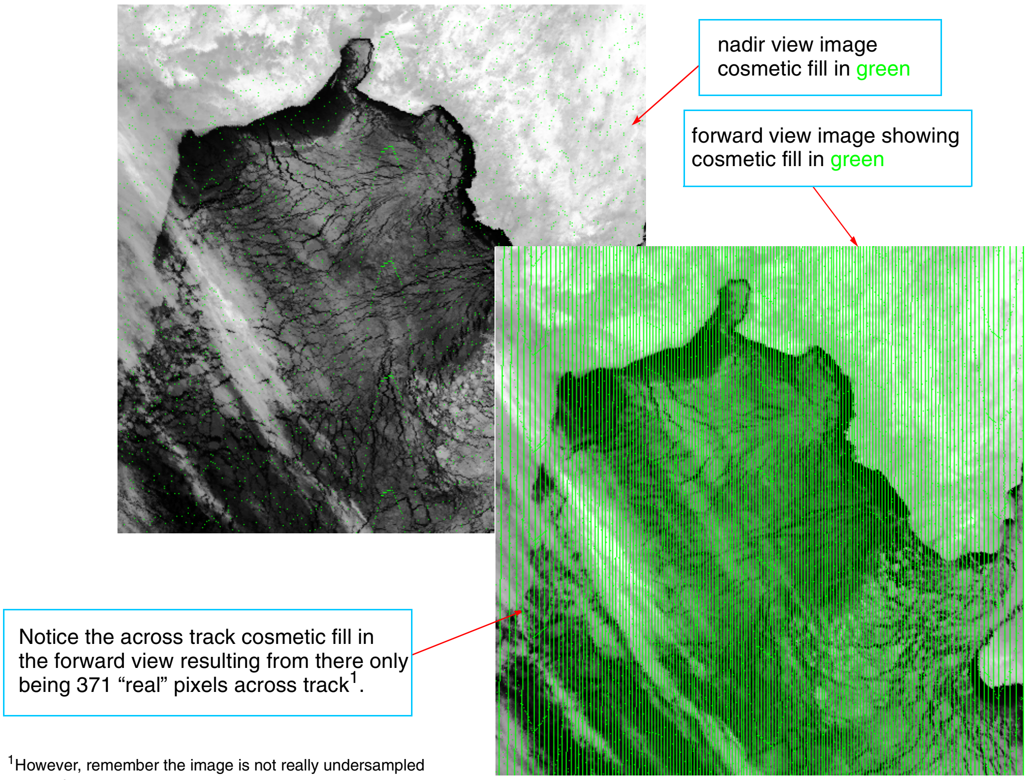

[1]However, remember the image is not really undersampled as the forward view pixels are 1.5km x 2km, and have been put into a 1km grid. Thus due to forward view pixel's larger size the duplication of pixel's by cosmetic fill is an approximation to the forward view pixels "real" size in the image. The along track dimension is oversampled wrt. to forward view pixel size as it is collected a 1km spatial interval.

These "cosmetic" fill pixels arise out of the re-gridding process where the curved instrument swaths of data are mapped onto a regular grid. The business of regridding places each ATSR pixel into a 1.0km $\times$ 1.0km box, and assumes all pixels have the same size. Regridding has two interesting effects. Pixels which are small, and whose Earth-locations are therefore very small, may be placed within the same 1.0km $\times$ 1.0km box (in which case the first would be overwritten). Also, some pixels in the re-gridded image may remain unfilled. This unfilling occurs when pixels are large, and consequently further apart than 1km. After regridding, each pixel image is "cosmetically" filled: pixels which remain unfilled by real data are filled by copying the nearest (filled) neighbour.

It can be seen that this process of cosmetic-filling has the effect of (approximately) reconstituting original pixel sizes. Filling occurs only where actual pixels are large, and therefore widely-spaced, but have been squeezed into 1.0km $\times$ 1.0km boxes. Nearest-neighbour copying reverses the pixel squeezing, and allows pixels to expand to a more representative size.

Pixels filled in this fashion are flagged in the ATSR products by negating the 0.65µm and 10.8µm pixel values to indicate that the data at that location obtained from the "cosmetic" fill procedure. The flags derived from these data apply to the image pixels for all the ATSR channels in the products at the same location.

To eliminate the interference caused by "cosmetic" fill users should take the absolute value of the pixel values before they are used, or scaled for display. In IDL this means typing "TV, BYTSCL(ABS(IMAGE))" rather than "TV, BYTSCL(IMAGE)". To be strictly correct users should also remember to take account of the exceptional pixel flag values (i.e. pixel values in the range (-1 to -8)) instead of just ignoring them by taking the absolute value since these flags provide essential data quality information. However, remember that these "cosmetic" fill flags need to be used properly to eliminate duplicate pixels if meaningful statistics are to be derived from ATSR images, i.e the flags are provided for a reason!

Also please remember this "cosmetic" fill flags apply to the pixel data in ALL the ATSR channels not just those channels which carry the flags.

3.0 Features of the ATSR instruments and their data sets

3.1 Missing data at the beginning or end of an orbit (or why is the forward/nadir view missing when the other one is there?)

These missing data are a consequence of the measurement strategy employed by the instrument, the partitioning of the data into orbits, the way the raw data is distributed on tape, and the limitations of the automatic data processing software to cope with splicing orbits of data together.

FIGURE 5. Example of a nadir view/forward view pair collected near the start of an orbit

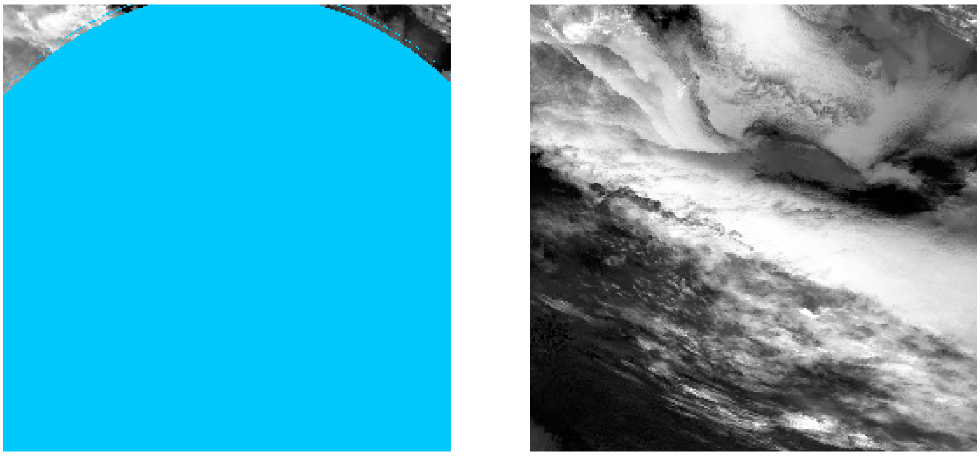

Forward view image

The blue area contains no data

Nadir view image

Such missing data always arises at the beginning or ends of orbits, the reasons for this are as follows. The ATSR instruments collect their forward and nadir view data swaths using a conical scan across track and the motion of the satellite along the orbit to provide a continuous swath of pixels along track. This sampling strategy results in the forward (along) track data at a given location being collected some 150 seconds before the corresponding data in the nadir view. Or put another way, the nadir view data from a given scan corresponds to forward view data collected 150 seconds earlier.

This time separation of the data from the two collocated views causes the loss of either the forward or nadir data sets whenever there is a break in the data flow. At the start of each orbit the forward data required to match the current nadir view scans will be missing as it was collected in the previous orbit. And similarly at the end of each orbit the nadir view data will be missing as it will be collected during the next orbit. This poses a real problem in the automatic data processing because of the additional complication that the tape archive containing the raw data set cannot assumed to be in time-order, so we cannot search the previous or next orbit on a tape to find the missing forward or nadir view data.

Figure 5 shows an example of a collocated forward/nadir view image pair collected near the beginning of an orbit. It is actually the third image in the sequence from the start of the orbit so it contains a partially filled forward view image and a complete nadir view. In the previous two images in the sequence the forward views were completely empty. In the first image the nadir image was also only partially filled because the start of the requested product was at a time before the first raw data was available in the orbit, and the ATSR data processing software set the absent section of the image with the value -2 to indicate the data was missing in the raw ATSR telemetry.

“Missing” data caused by this problem occurs at predictable locations through each orbit cycle and can be overcome by ordering extra products to cover the orbit joins. One product will carry the forward view data and the other the collocated nadir view. The other option is to request RAL generate a spliced orbit product, but this requires manual intervention and can only be undertaken infrequently.

3.2 Pixel selection maps and transitions (or why are portions of my image missing?)

Figure 6 shows the kind of features often seen in the visible channel data from ATSR-2. These features are never seen in ATSR-1 data or the infrared data from ATSR-2, and will not be present in any AATSR data. They arise from the data compression methods used by ATSR-2 to maximise the data return from its visible channels as the telemetry bandwidth available to the instrument varies with the use of the ERS-2 scatterometer wind-wave mode.

During the day time ATSR-2 is guaranteed access to a minimum data rate, known as L-rate, but is only able to use the high data rate, known as H-rate, whenever the scatterometer is not in wind-wave mode (i.e. over most large land masses). H-rate allows all the ATSR-2 infrared and visible channel data to be sent at full 12-bit resolution over the full 512 km swath. However, much of the time only the L-rate is available, which is only sufficient to send all the infrared data in ATSR-1 format and a limited amount of visible channel data. To maximise the visible data available to users two basic L-rate data collection strategies are used: a) used for 20 days each month provides a 180 km visible channel swath at full 12-bit resolution, and b) for 10 days each month provides the full 512 km swath but at 8-bit resolution. H-rate is used whenever it is available, and the instrument reverts to the scheduled L-rate mode for the day when the scatterometer operates in wind-wave mode.

The selection of these alterative instrument telemetry modes is achieved by using different pixel selection maps in Digital Electronic Unit that controls the instrument. These pixel selection maps instruct the instrument's data formatter which pixels are to be inserted into the telemetry stream from the instrument, how many bits of data are to be included, and how they are to be encoded. A more complete explanation of the ATSR pixel maps and their purpose can be found in the ATSR User Guide, so look there if you need to know all the nitty gritty details.

Figure 6 shows an image collected partly in one of the ATSR-2's wide swath modes which also includes a transition to a 180 km narrow swath mode. Images that actually show one of these mode transitions like this occur fairly infrequently, more often an image will contain data from only one pixel map. Check the pixel map information in the header as it tells you about which map is selected. Don't be alarmed if the visible channel data is only present over a narrow swath and the infrared data is available over the whole 512 km swath, this is normal if the instrument is in its narrow swath L-rate visible channel mode.

FIGURE 6. Example of a pixel map transition between L-rate narrow swath and wide swath L-rate or H-rate

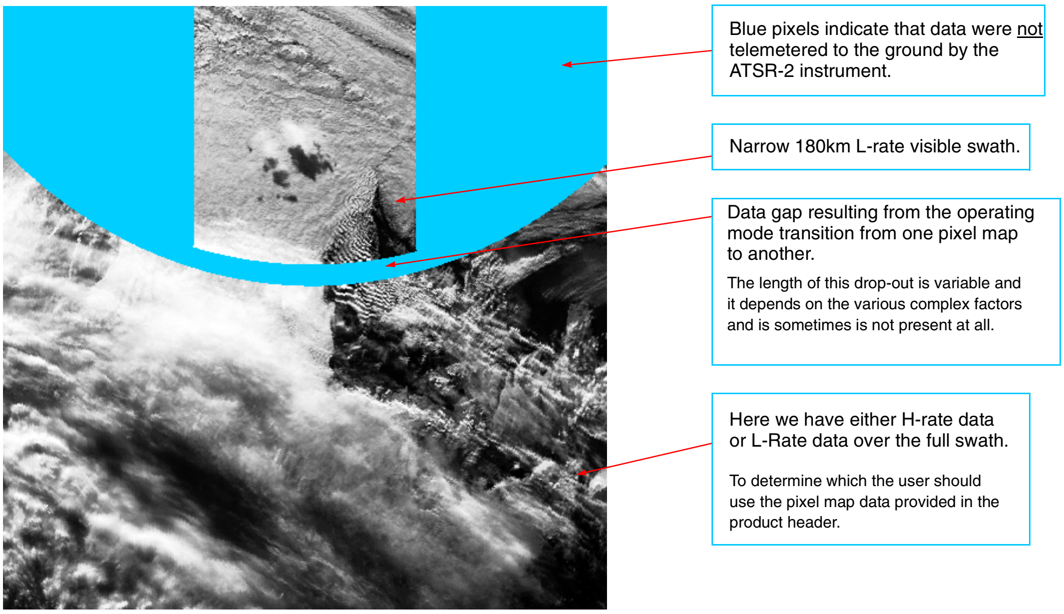

3.3 Saturated pixels

The ATSR instruments were originally built to make accurate and precise observations of sea surface temperature so their design was optimised to measure temperatures in the of expected range ocean temperatures (i.e. between 270 and approx. 310K). As a result the instrument's channels saturate over hot land surfaces, such as deserts. Saturation of the data in a particular channel is indicted by a -5 value in an image pixel.

Figure 7 shows a more unusual set of ATSR images containing saturated pixels, where the scene was collected over a cold cloud not over a hot target. The three images in the figure illustrate that the 3.7μm channel can saturate even over a cold target when there are objects in the scene that scatter solar radiation strongly. In the case shown in the figure the thermal signal at 3.7µm is dominated by the solar signal and causes the channel to saturate, this also happens over bright desert scenes.

FIGURE 7. Example of an ATSR-2 image containing saturated pixels

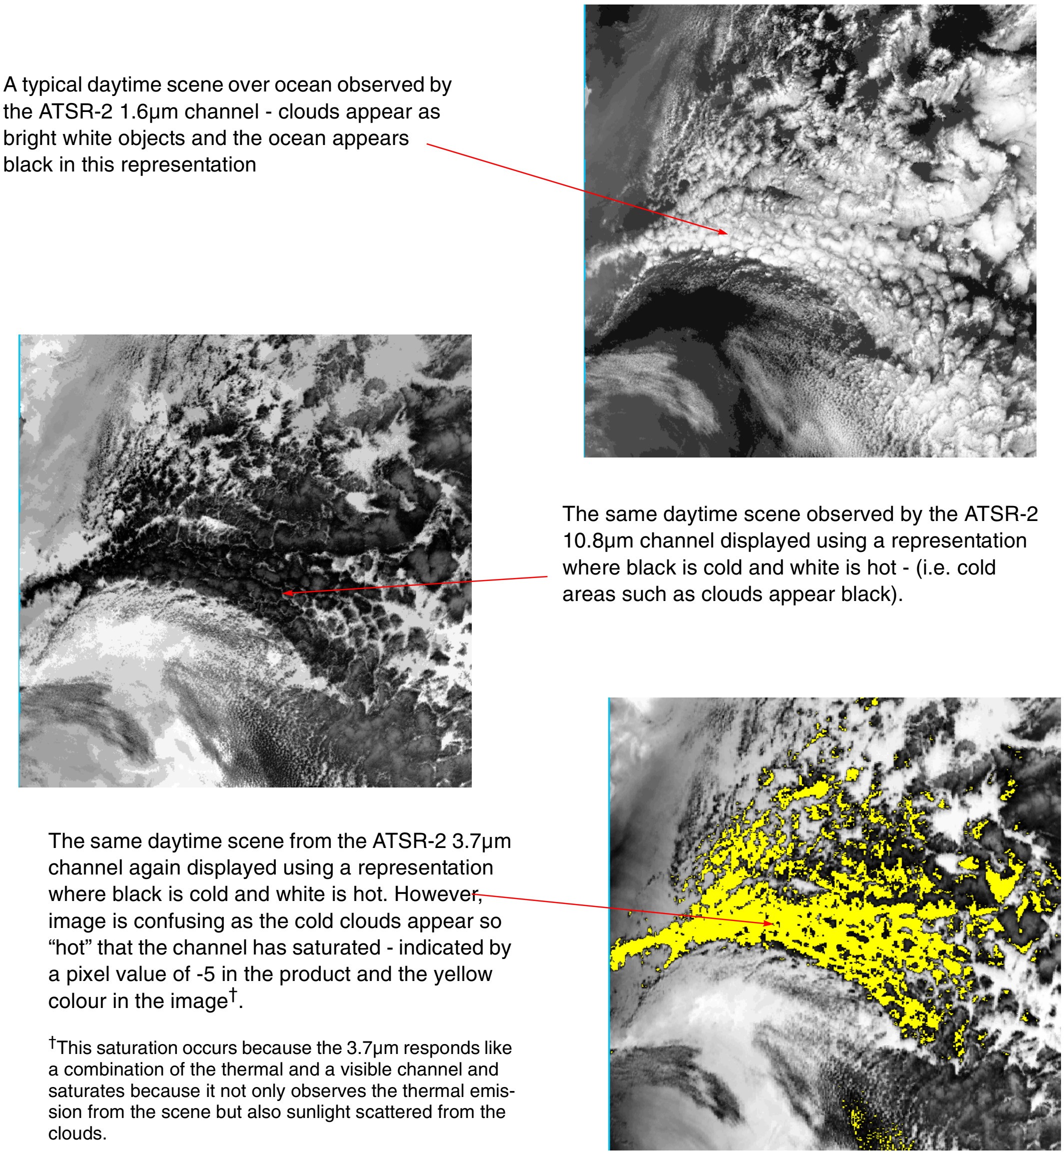

3.4 Sunglint

During certain parts of daylight orbits the visible, 1.6µm and 3.7µm data from the ATSR instruments are affected by sunglint (i.e. the specular reflection of the sun from the ocean surface into the instrument field of view).

Figure 8 shows an ATSR image pair where the nadir view of which is affected by sunglint, and shows the bright region down one side which is characteristic of sunglint. Only the nadir view is affected because the relationship between the solar angles and the viewing geometry is different for the forward view and specular reflection does not occur. It is always the case that only one of the views will contain sunglint, but it can be either of the views that is affected.

FIGURE 8. Example of a collocated forward and nadir view pair illustrating that sunglint only affects one of the image pair

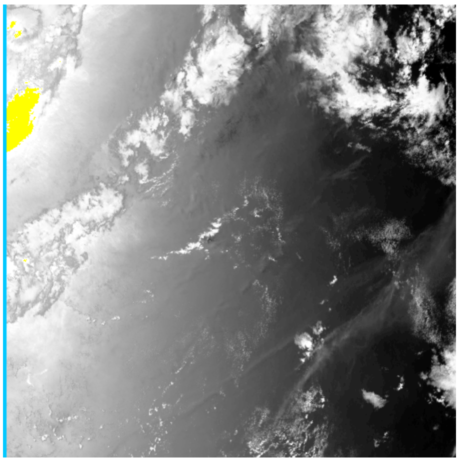

Nadir view 1.6 µm image

Illustrates strong across image variation in the brightness of clouds and sea due to sunglint - i.e. the reflection of the sun from the surface

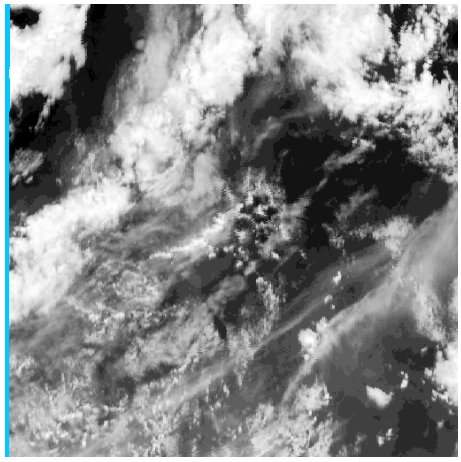

Forward view 1.6 µm image

Shows uniform distribution of brightness in the sea and cloud across image seen under normal illumination conditions.

4.0 Instrument Anomalies

4.1 ATSR-2 scan mirror jitter

The scan mechanism on the ATSR-2 instrument does not always perform as well as it should because of a hardware problem with the scan drive system. Most images include a few “jittered” scans like those illustrated in Figure 9 which are caused by minor speed variations in the mirror drive. Mirror speed control is achieved through an active control loop that adjusts the mirror drive at the start of each scan. The effect of the speed variation is cumulative around the scan which explains why the forward view data, that always come at the end of each scan, are more strongly affected by these minor variations in mirror speed than the nadir view data that are collected in the early part of the scan. These “jitters” are often very hard to detect even in forward view data, in the example shown in Figure 9 it has been necessary to zoom the image to make the “jitters” visible.

In most cases these minor "jitters" do not have a serious impact on utilisation of the data, but in some applications these "jitters" can be a problem when matching up forward/nadir view image pairs. Users need to be aware that these "jitters" do occur and decide for themselves whether they are likely to be a problem for their work. It is hoped that a future version of the ATSR processing scheme will include a "jitter" correction scheme

FIGURE 9. . An ATSR-2 image illustrating the effect of minor "jitters" on the data

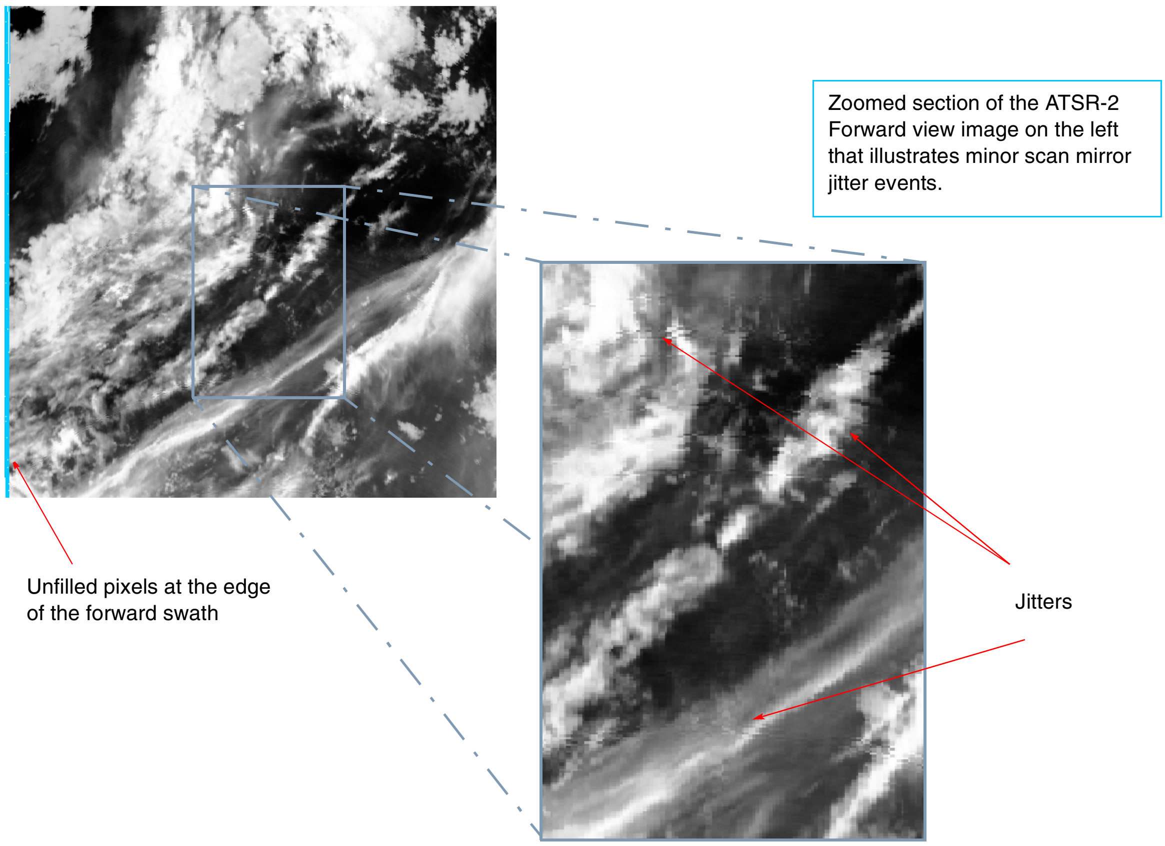

More serious periods of "jitter" also occur that can corrupt whole portions of an ATSR image. An example of one of these events is shown in Figure 10 below, fortunately these events are occur fairly infrequently. There is no way of recovering the missing data during these events as the instrument's data formatter loses track of where in the scan the real pixel position is. Most users spot these "jitters" very easily, but automated ATSR data processing systems may not correctly handle data during these events. Users should ensure that all data they are using was processed with SADIST-2 version 301 (or a later), as improvements in "jitter" flagging were made in v301.

FIGURE 10. Example of an ATSR-2 image affected by a major scan mirror “jitter” event

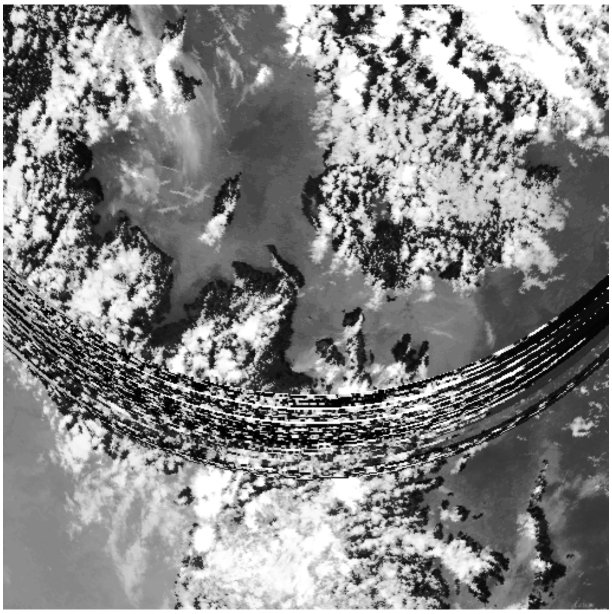

4.2 ATSR-2 3.7µm anomaly

Since the shutdown of the instrument for the Leonids shower on the 17th August 1999 partial drop-outs have been observed in the 3.7µm channel data. Figure 11 below shows a typical 3.7µm images affected by these drop-outs. There are no signs of these features in any of the other channels, and at present there is no explanation for what is causing this corruption of the data.

FIGURE 11. Example of the unexplained drop-out features in the 3.7µm data collected since 17th August 1999

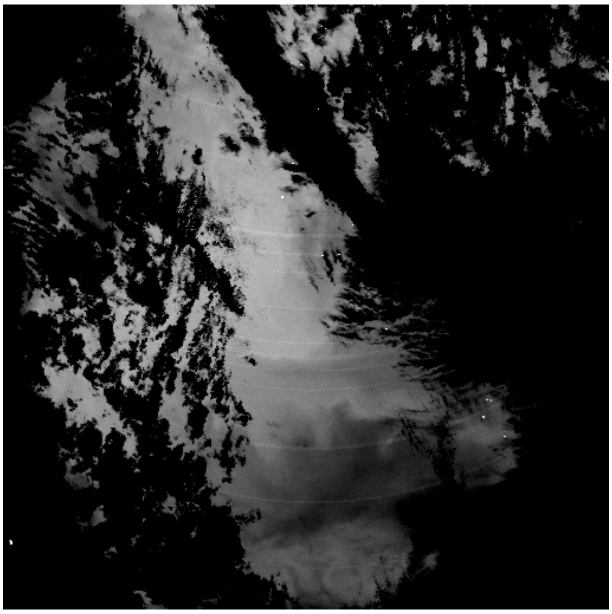

4.3 Why are my visible channel data values -ve?

During a period of a few weeks from February 11th 2000 data until March 30th 2000 ATSR-2 data processed with versions of SADIST-2 prior to v321 contained spurious -ve pixels in the visible channel data.

This incorrect processing occurred after the ATSR-2 instrument was restarted following a planned shutdown of ERS-2 and ATSR-2's auto gain/offset loops failed to initialise and operate correctly. The processing error ocurred because an incorrect offset value from the ATSR-2 housekeeping telemetry was accidental being used within SADIST-2 to calibrate all the visible channel data. There is nothing wrong with the raw ATSR-2 data collected during this period, and correct visible channel data can be regenerated by using SADIST-2 v321or later to re-process the required data.

Prior to v321 data from the 0.87µm channel in SADIST-2 products was irrecoverably corrupted, and the 1.6µm and the other visible channels are difficult to interpret. Most of the faulty data sets were distributed came from the Tromsø NRT system, very few off-line products processed incorrectly reached users. Any user requiring visible channel data collected during the affected period should contact the Project Team at RAL.