Data Quality Report - 2012

Hyperspectral

ARSF - Data Analysis Node

Updated on: April 11, 2013

Contents

1 Overview **2**

2 Geo-referencing accuracy **2**

3 Timing Errors **2**

4 Sensor calibration **3**

4.1 Spring 2012 calibration 4

4.1.1 Wavelength calibration accuracy 4

4.1.2 Radiometric calibration 6

4.2 Winter 2012 Calibration 9

5 Overflowed Pixels **10**

6 Smear Correction **11**

7 Bad CCD Pixels **15**

7.1 Detection method A - Constant input variable output 16

7.2 Detection method B - Constant input constant invalid output (CICO) 16

7.3 Detection method C - Linear input non-linear output (LINO) 16

7.4 Detection method D - Rapid saturation 17

7.5 Detection method E - Visual inspection 17

1 Overview

This report describes issues that should be considered when further processing any of the 2012 Airborne Research Survey Facility (ARSF) datasets. The document may be updated over the course of the year, with the latest version available at:

http://arsf-dan.nerc.ac.uk/trac/wiki/Reports

2 Geo-referencing accuracy

ARSF currently deliver data at level 1 (calibrated sensor data) and for 2012 are also trialing delivering mapped level 1 data (level 3). This allows users to get quick access to georeferenced data but also maintains the capability of being able to applying user-developed algorithms and generating level 2 products (e.g. atmospherically corrected radiances) prior to mapping to a projection or datum that suits.

The quality of the geocorrection for each project is described in the documentation supplied with the project and is normally of the order of a couple of metres (approximately 1 pixel). Where a vector overlay or other ground truth information is available, ARSF provide an indication of the average error, included in the screenshot images. If you need higher accuracy, please contact us at: arsf-processing@pml.ac.uk. It may be possible to tune specific flight lines for higher accuracy or we can provide instructions on how to make your own alignments.

3 Timing Errors

Due to an error in the handling of synchronisation between the navigation system and the Eagle and Hawk sensors, small timing errors (order of 0.05s) may occur. The consequence of timing errors is to cause scan lines to be positioned incorrectly and manifest visually as "wobbles" in the imagery. The wobbles are correlated to, but out of sync with, movements of the aircraft. An example is shown in Figure 1 and Figure 2 below.

This issue has been extensively investigated and demonstrated to be a fault in the Specim systems. Specim are working with ARSF to provide upgrades and improvements to correct this issue, but have not yet succeeded.

Therefore we endeavour to correct all timing errors prior to delivery. As this is a manual process and relies on finding suitable visible features in the

imagery, some errors may still remain. If any are found, please contact us at arsf-processing@pml.ac.uk.

Figure 1: *timing error in an Eagle line*

Figure 2: *corrected version of above (0.13 seconds difference)*

4 Sensor calibration

Prior to the start of the 2012 flying season, the Eagle and Hawk instruments were calibrated at a NERC ARSF facility in collaboration with the Field

Spectroscopy Facility. A post season calibration was also performed at the end of the 2012 flying season. A summary of the calibrations follow.

4.1 Spring 2012 calibration

4.1.1 Wavelength calibration accuracy

Wavelength calibration was undertaken by viewing a number of spectral lamps using both the Eagle and Hawk. The lamps viewed were: Hg-Ar, He, H, Kr, O, Ne and CO2. These provide a number of spectral emission features at known wavelengths that can be seen with the sensors, this allows specific pixel numbers to be confirmed as viewing particular wavelengths. It is normal for there to be a very small change in these results from one calibration to the next.

Wavelength calibration was undertaken in January 2012. Differences from the previous calibrations can be found in tables 1 and 2 - note the single Fail for Eagle probably denotes a mismatched spectral feature. The threshold for a "Pass" for Eagle is lower than that for Hawk because of Eagle's higher spectral resolution.

Spectral Line (nm)	Measured Wavelength (nm)	FWHM (nm)	Error (nm)	Pass/Fail (<4nm Err.)
1083.0	1082.78	7.76	0.22	Pass
1181.9	1178.93	3.03	2.97	Pass
1363.4	1361.87	3.06	1.53	Pass
1442.7	1443.88	3.04	1.18	Pass
1816.7	1816.06	3.05	0.64	Pass
1875.1	1873.06	7.92	2.04	Pass
2058.7	2059.10	8.04	0.40	Pass
2190.3	2188.25	3.04	2.05	Pass
Mean:		4.87	0.98	Pass

Table 1: Wavelength calibration offsets for Hawk, Spring 2012 calibration

The FWHM values obtained in these tables have been obtained by fitting gaussian curves to the spectral features and then measuring the best-fit FWHM. While it is believed that this gives a reasonable estimate of the

Spectral Line (nm)	Measured Wavelength (nm)	FWHM (nm)	Error (nm)	Pass/Fail (<2nm Err.)
1083.0	1082.02	3.18	0.98	Pass
965.8	965.19	2.73	0.61	Pass
922.4	921.88	2.88	0.52	Pass
912.3	911.74	2.77	0.56	Pass
892.9	892.37	2.75	0.53	Pass
877.6	877.19	2.76	0.41	Pass
866.8	866.29	3.15	0.51	Pass
852.1	851.68	2.95	0.42	Pass
841.8	841.32	3.46	0.48	Pass
826.5	825.98	3.00	0.52	Pass
801.1	805.40	2.93	-4.30	FAIL
794.8	794.23	2.87	0.57	Pass
777.4	776.64	2.80	0.76	Pass
763.5	762.74	2.77	0.76	Pass
750.9	750.04	3.08	0.86	Pass
738.4	737.61	2.85	0.79	Pass
727.3	726.60	3.14	0.70	Pass
696.5	696.12	2.96	0.38	Pass
667.6	667.65	2.90	-0.05	Pass
656.3	656.29	2.79	0.01	Pass
578.1	578.01	3.84	0.09	Pass
556.9	556.71	2.88	0.19	Pass
546.1	545.81	2.83	0.29	Pass
501.6	502.04	3.20	-0.44	Pass
486.1	486.61	2.91	-0.51	Pass
435.8	435.56	2.80	0.25	Pass
Mean:		2.97	0.23	Pass

Table 2: Wavelength calibration offsets for Eagle, Spring 2012 calibration

FWHM values of the sensors, this has not been conclusively demonstrated - use with caution. Note also that the FWHM as labelled in the data header (hdr) files is actually bandwidth of each band.

4.1.2 Radiometric calibration

Following the wavelength calibration, radiometric calibration is undertaken by viewing an integrating sphere provided by the NERC Field Spectroscopy Facility using both the Eagle and Hawk. The sphere is calibrated to NPL standards. This provides a light source of known radiance at each wavelength in use, allowing a calibration curve for the instruments to be calculated.

Figures 3 and 4 show the percentage change in the Eagle and Hawk calibration multiplier for each pixel on the CCD since the 2011 calibration that was in use. It can be seen that the percentage change for Eagle is quite high in some regions. This was caused by a shift in components within the instrument during work to change the shutter of the Eagle. There are no implications for data collected from 2012 onwards (this is the purpose of the calibration), but as a result of this ARSF have opted to primarily use a gain setting of 2 for Eagle (previously 1 was used) to increase the sensitivity of the instrument.

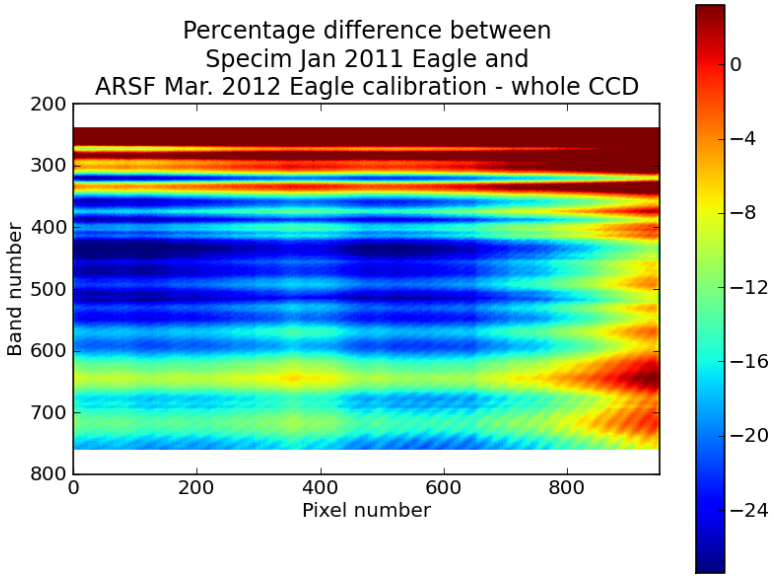

Figure 3: *Eagle calibration multiplier percentage differences from previous calibration*

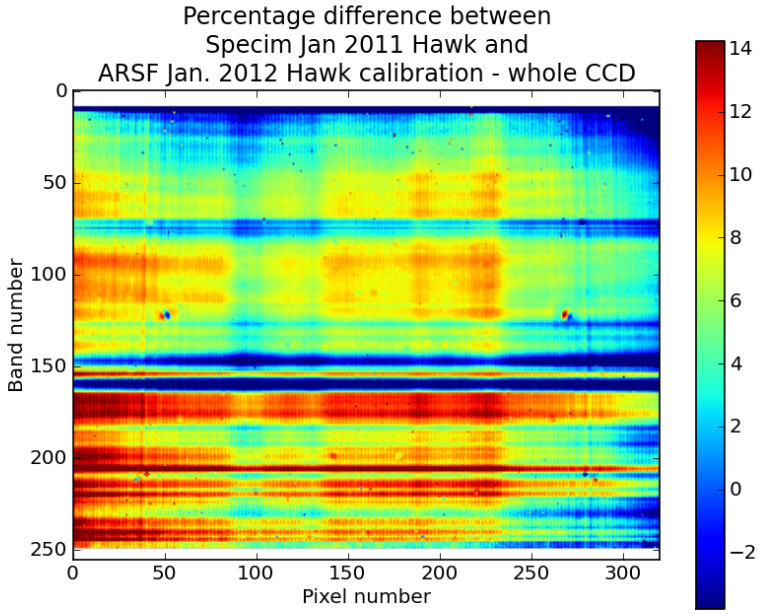

Figure 4: *Hawk calibration multiplier percentage differences from previous calibration*

4.2 Winter 2012 Calibration

During the down period for the aircraft, over the winter, ARSF undertakes a calibration of the hyperspectral instruments. During the course of our normal calibration in December 2012, we discovered that the Eagle was suffering from a loss of sensitivity as against the previous year's calibration. This sensitivity loss is wavelength-dependent and ranges from around -3% at the red end of the spectrum to around -18% at the blue end of the spectrum. This is illustrated in Figures 5 and 6.

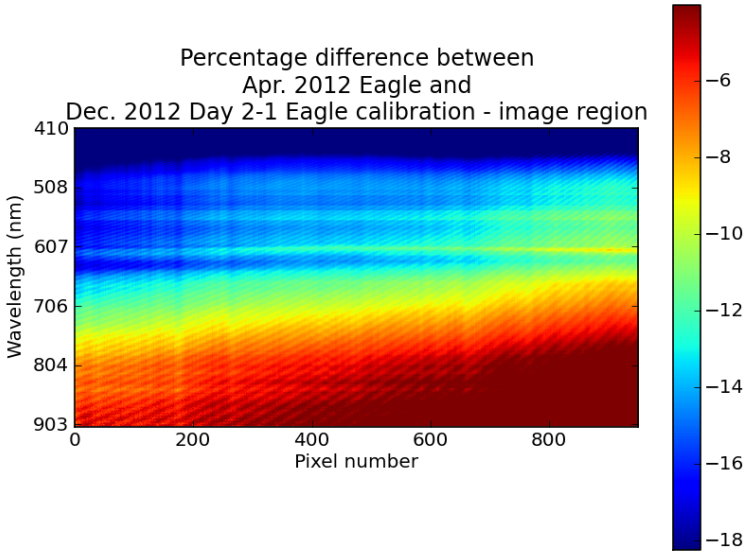

Figure 5: *Percentage change in Eagle calibration multiplier - whole CCD*

In order to maximise available time for data collection, we do not carry out calibration during the flying season and therefore we are not able to characterise the progression of the loss of sensitivity during 2012. The sensor manufacturer has suggested that this is a gradual and expected deterioration of the sensor due to age, this is a possibility but seems a very large loss of sensitivity for this cause, and we have not seen such deterioration during previous flying seasons. If this is the case we do not know whether the loss was linear or non-linear over the course of the flying season. Another possibility is that there was a step change in sensitivity at some point during the season. We have no available evidence for or against any of these possibilities.

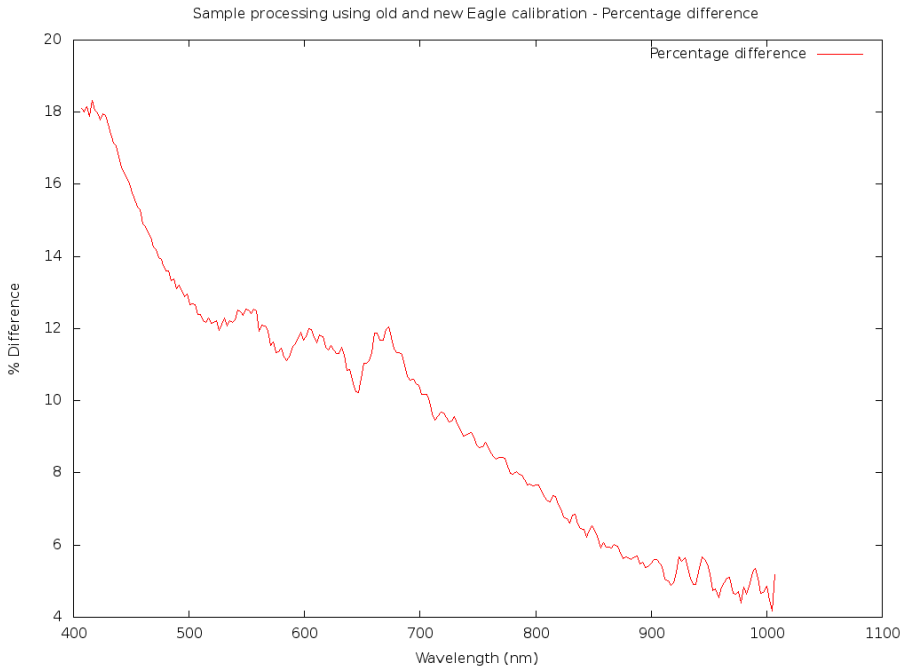

Figure 6: *Percentage change in Eagle calibration multiplier - single pixel*

We are looking at options for mitigating this sensitivity loss during the next flying season, including replacement of the sensor, but we cannot guarantee that data collected during 2012 were processed with the most appropriate calibration applied. All delivered data have initially been processed using the calibration file from the beginning of the 2012 flying season - this is our standard practice and is done to maintain consistency in the instrument calibration used over the course of a flying season, but it may not be appropriate here given the large loss of sensitivity.

It is worth noting that this problem affects only the AISA Eagle. Data from the other ARSF instruments (Hawk, LiDAR, RCD camera) are unaffected.

5 Overflowed Pixels

The instruments have a limited dynamic range and must be set to capture data over the appropriate range of signal strength. For example, if the area of interest is dark, then the instrument will be configured to capture as much low light detail as possible. This configuration is set based on operator

experience, the principal investigator's indication of the areas of importance and the prevailing conditions. Inevitably, some pixels are unexpectedly bright - e.g. sunglint over water or part of a cloud. These pixels may exceed the maximum capture level and overflow. Typically they are not in areas of interest, but should be accounted for. The accompanying mask file will contain an overflow flag value in the level 1 equivalent pixel.

In Hawk, overflows are marked for just the pixel/band in question. However, Eagle uses a frame transfer CCD, where data are read out in rows. Incoming light continues to accumulate in unread rows during the transfer and is removed by "smear correction" software, which relies on data from one row to correct the next. If a pixel overflows, information is lost and all subsequent pixels in that column cannot be fully corrected. In Eagle, the net effect is that an overflow at 600nm will cause all bluer bands (600nm -> 400nm) to be under-corrected for that spatial pixel. In this case, the mask file will contain a "smear affected" flag value for the equivalent pixel position. When Eagle data with overflows are delivered, we mask all bands (in the mask file) following an overflow as they will incorporate some unknown additional light. If you would prefer your actual level 1 files to be masked out rather than use the separate mask file please contact arsf-processing@pml.ac.uk.

Figures 7 and 8 show an Eagle level 1 band and equivalent mask band in the blue part of the spectrum (450nm). Note that although these data may appear good in the level 1 image, the bright mask values mark pixels that have been adversely affected due to overflows occurring in a higher band.

6 Smear Correction

The Eagle uses a CCD that shifts data out line by line at the end of a frame. While this readout process is quick, additional light still falls onto the detector during the readout period. Currently this is corrected for by subtracting a small amount of light measured in the previous line(s) as they are read-out. This procedure assumes the light input is unchanged during the integration and read-out, but this is a good approximation. However, the sensor is often run with a bandset that doesn't record all of the lines.

The Eagle CCD is 1024x1024, with nominal sensitivity from ~200nm to 1200nm, with readout progressing from red (1200nm) to blue (200nm). In operation, only the middle ~500 bands are recorded (~450-950nm), partly due to low sensitivity in the other regions, but also because there are significant internal reflections/second order effects (which is normal). Figure 9

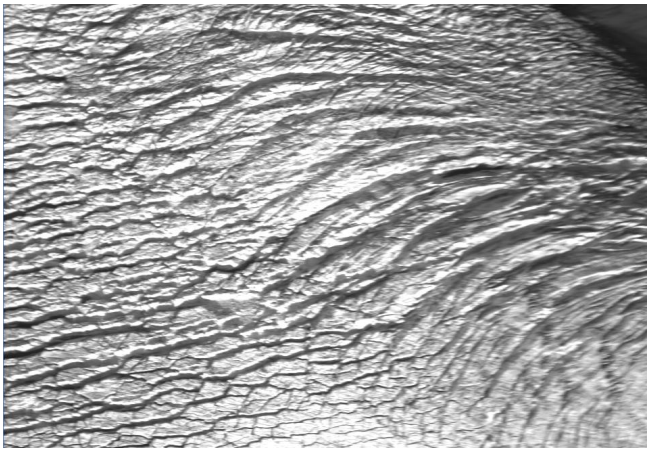

Figure 7: *level 1 image of bright ice at 450nm*

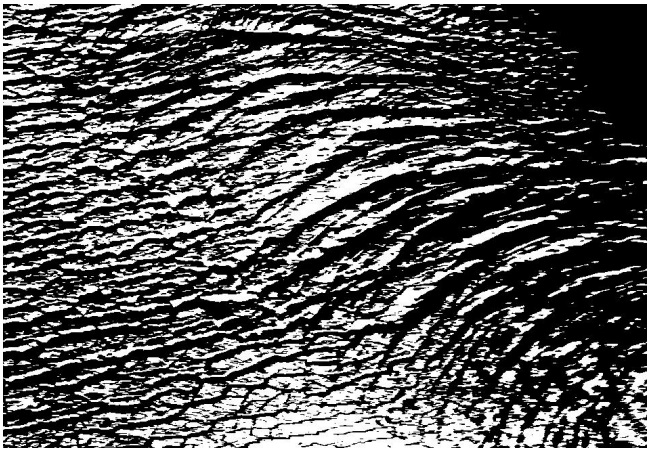

Figure 8: *mask image of same band as Figure 7. White pixels are adversely affected.*

shows a view of the amount of light falling on different parts of the detector. The red box shows the approximate area that is recorded in normal conditions. Internal reflections can be clearly seen, although they have been highly enhanced to make them visible. The amount of light in the central region greatly overwhelms that of the reflection, although their contribution to error can still be significant in weakly illuminated bands.

Consequently, any light falling in the 950-1200nm area cannot be corrected for and smear resulting from that light will remain in the final image.

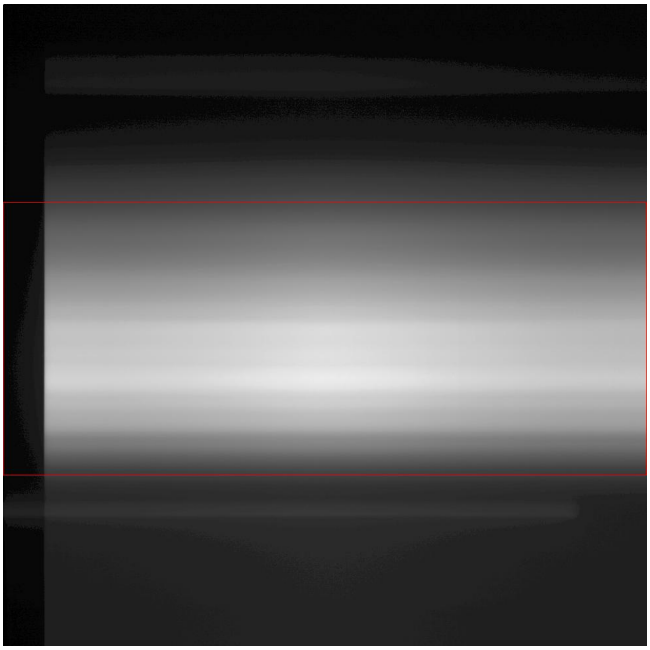

Figure 9: *highly enhanced view of the Eagle CCD, showing recorded bands (inside the red box) and light outside this region.*

This erroneous light will have the greatest effect where the "true" signal is lowest (e.g. absorption bands). We have run some simulations and the error is naturally worse in bands where there is little light (the red and blue ends of the spectral range, where signal versus noise is lowest, which possibly contributes to why these are poorer quality) and worse the shorter the integration time (below 10ms integration times, errors rapidly increase).

Figure 10 shows estimates of the error introduced by smear into a real-world dataset. The values for unrecorded smear were taken from a calibration experiment and applied to real data from Ethiopia (day 299a/2008) at a variety of integration times. This will likely cause overestimates of the error, as the calibration lamp is brighter than real-world collection conditions, but it is indicative of the relative magnitudes of the error to be expected. As can be seen, the error is dependent on the signal at a particular band, with higher error (2-3%) at the edges of the spectral range, but also with spikes of 1% error in the absorption bands.

The estimated (probably overestimated) error ranges from <0.1% (long integration times, high signal bands) to ∼70% error (worst possible case of short integration time, low signal bands). The following table 3 shows

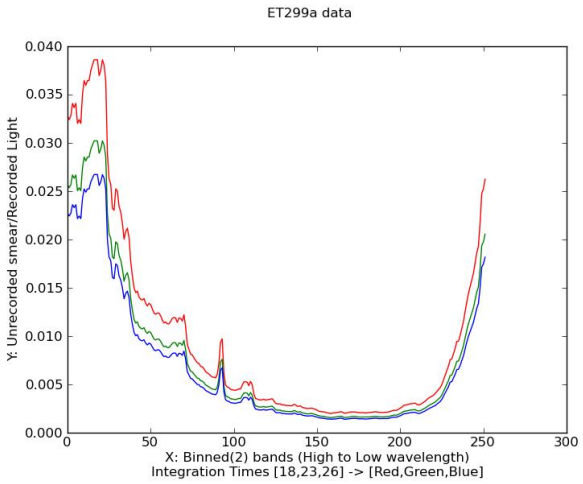

Figure 10: *estimated real-world error due to uncorrected smear, using calibration lamp data to provide estimates of the smear and real Ethiopia data (day 299a/2008). Note that the x-axis is band-number, not wavelength (band 0 = ∼1000nm, band 250 = ∼400nm).*

the best and worst-case error estimates for real-world data at a variety of integration times. Figure 11 shows the distribution of integration times over several years of data collection, to give an indication of the likely impact.

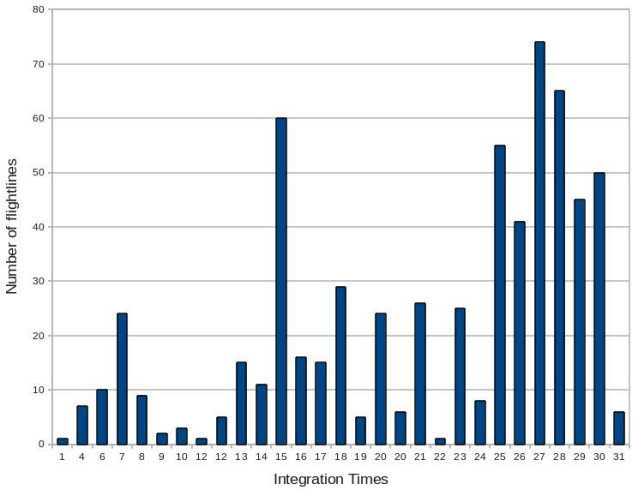

Figure 11: *distribution of integration times for flights in 2008-2010.*

Integration Time (ms)	**Best case error (% @ peak signal strength)**	**Worst-case error (% @ weakest signal)**
1	5%	70%
4	∼2%	20%
6	2%	10%
8	0.6%	9%
12	0.5%	6%
16	0.3%	4%
18	0.2%	4%
23	0.2%	3%
26	0.1%	2.5%
28	0.1%	2.5%
31	0.1%	2%

Table 3: *Best and worst-case error estimates for real-world data*

7 Bad CCD Pixels

The Hawk instrument has a number of bad pixels that give inaccurate values. There are many different types of error (e.g. constant pixel values, uncorrected offset, duplicating neighbouring pixels, etc), and ∼1% (about 600) of pixels are to be expected to be bad on the type of CCD used in the Hawk instrument. A list of known bad pixels has been included within the mask files. The bad pixels will appear in level 1 datasets as straight lines along the direction of flight and as undulating lines in level 3 following the motion of the aircraft (e.g. Figure 12). Typically, they will only affect a single band and are difficult to detect. A complete solution for detecting and removing these started in 2010 and has been finalised for 2012 data.

The final list of bad pixels uses 5 methods of detection, run on a set of test data. If a pixel fails any method then it is marked as bad. The methods are tested sequentially, and if a pixel fails an earlier method then it is still tested in later methods, so that it is marked bad for every test that it fails. The method used to flag the bad pixels can be identified using the mask files, which in turn allows masking of only certain bad pixel types.

The methods used are summarised below:

- Constant input - variable output

- Constant input - constant invalid output
- Linear input - non-linear output
- Rapid saturation
- Visual inspection

7.1 Detection method A - Constant input variable output

This method marks pixels bad when they vary significantly given a constant light as input. This is tested by considering the median raw value (DN) for the pixel over time, and testing individual epochs against a percentage threshold. If the value exceeds this threshold it is flagged as bad.

7.2 Detection method B - Constant input constant invalid output (CICO)

This method marks pixels bad when their response to a constant light greatly varies from the response of their spatial and/or spectral neighbours. Again, this is done using the raw values. To determine when a pixel's response varies it has to be compared to its close neighbours. A moving window is used to detect responses that differ significantly.

The mean and standard deviation of pixel values are calculated for the moving window and then used within the formula below.

$$\frac{\mu_{ofCCDpixel} - \mu_{windowcentredat(sample,band)}}{\sigma_{windowcentredat(sample,band)}} \quad (1)$$

where the μ are means and σ the standard deviation. When this exceeds the theshold the bad-counter for this pixel is incremented. When the counter reaches a maximum allowed value the pixel being tested is marked bad.

7.3 Detection method C - Linear input non-linear output (LINO)

This method takes the average DN for each pixel over time, for multiple data captured at increasing integration times. The increasing integration times should correspond to increasing sensor response in a linear fashion. Using regression over the average values versus integration time it is possible to get a measurement of linearity from the Pearson Product-Moment Correlation

Coefficient - the closer the value to 1.0 the better the fit. The formula is given as:

$$\sum_{i=1}^{N} \frac{(X_i - \mu_X) * (Y_i - \mu_Y)}{(N-1) * \sigma_X * \sigma_Y} \tag{2}$$

with X as the integration times and Y as the mean value of the CCD pixel, μ and σ are the sample mean and standard deviation, N the number of integration times tested (sample size). A pixel is marked bad if the pearson product-moment correlation coefficient is less than a threshold value.

7.4 Detection method D - Rapid saturation

The rapid saturation detection works similarly to CICO in order to detect linear but invalid responses to different integration times. This method uses the gradients of the regression fit of mean pixel values versus integration times. Once again, a moving window of certain spatial and spectral width will iterate over each pixel of each band. If the slope at the centre of the window, which is being tested, varies greatly in relation to its neighbours, then the pixel in that position will be called bad. The function used to scan over each pixel is:

$$\frac{slope_{ofCCDpixel} - \mu_{windowcentredat(sample,band)}}{\sigma_{windowcentredat(sample,band)}} \tag{3}$$

when the above value exceeds the threshold then the pixel will be called bad. In other words, if a pixel is saturating more rapidly than its neighbours then it is detected as bad.

7.5 Detection method E - Visual inspection

This method is purely a visual examination of the test data and marking any pixels as bad that are clearly giving erroneous results.

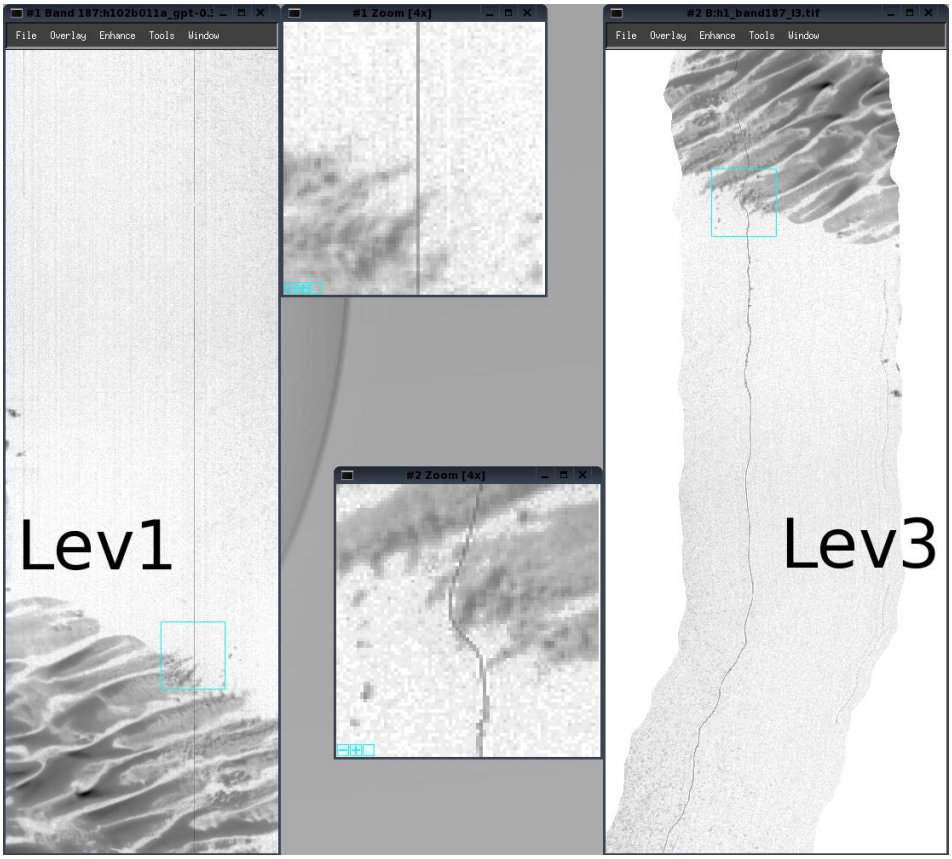

Figure 12: *A bad pixel on Hawk band 187, in a scene over water (images inverted to improve contrast on paper)*