

The calibration of AMIE images

Björn Grieger

2009-Mar-23

Document number	S1-AMIE-SGS-TN-013
Issue/Revision	1/-

Contents

1	Introduction	2
2	Dark correction	2
2.1	The theory of dark correction	2
2.1.1	The dark current contribution	2
2.1.2	Dependence on exposure time	3
2.1.3	Dependence on CCD temperature	3
2.2	Master dark frames from in flight dark sky observations	4
2.2.1	The set of dark sky images	4
2.2.2	Estimation of master darks	11
2.2.3	Explained variance	13
2.3	Applying the dark correction	13
2.3.1	Test: Dark correcting a dark image	13
2.3.2	Example: Dark correcting a lunar image	16
3	Flat fielding	20
3.1	The theory of flat fielding	20
3.2	A flat field from in flight data	20
3.3	Comparing space and lab flat	24
3.4	Applying the flat fielding	27
4	Performing the calibration	27
4.1	Cookbook for home grown calibration	27
4.1.1	Reading an image	27
4.1.2	Dark correction	27
4.1.3	Optional filtering	31
4.1.4	Flat fielding	31
4.2	Calibration using the provided IDL routines	31
4.2.1	Calibrating single images	31
4.2.2	Running the calibration pipeline	32

5	Open calibration issues	32
5.1	Missing calibration information	32
5.2	The corrupted band adjacent to the CCD readout area	32
5.3	Vertical stripes with 8 pixel spacing	33
5.4	Scattered light inside the camera	36
5.5	Absolute calibration	45
5.6	Distortion	45

1 Introduction

During its Earth escape phase, the SMART-1 spacecraft crossed the Earth's radiation belt various times. This caused a considerable increase of the dark current of the CCD of the AMIE instrument. Therefore, the dark images which had been acquired in the laboratory prior to the mission are not adequate to perform a dark correction of images taken during the lunar phase of the mission and the extended mission. In order to dark correct the AMIE images, new master dark frames have been computed from dark sky observation which were carried out during lunar phase and extended mission. The dark correction is described in section 2.

Also a new flat field has been estimated from in flight data. The camera's flat field was less affected by radiation damage than the dark current behaviour. However, there are some considerable differences from the laboratory flat field. Moreover, because of the fact that the spectrum of the lamp used in the laboratory is very different from an average lunar spectrum, we obtain strong biases close to the filter boundaries. These biases are avoided by using the in flight flat field. The flat fielding is described in section 3.

A guide how to perform the calibration of AMIE images is provided in section 4. Open issues of the calibration are discussed in section 5.

2 Dark correction

2.1 The theory of dark correction

2.1.1 The dark current contribution

The data number D measured for a particular CCD pixel is

$$D = d_0 + B \cdot f(T) + S \cdot f(T) \cdot t_e + C \cdot I \cdot t_e, \quad (1)$$

where

- d_0 is the fixed offset (which is 8 data numbers for AMIE),
- B is the bias which is built up during readout,
- $f(T)$ describes the temperature dependence,
- S describes the dark current during exposure,

- t_e is the exposure time,
- C is the radiance conversion factor, and
- I is the incoming radiance.

The quantities D , B , S , C , and I depend on the pixel, the others not. For a dark image, the incoming radiance is $I = 0$, and we get

$$\begin{aligned} D_{\text{dark}} &= d_0 + B \cdot f(T) + S \cdot f(T) \cdot t_e \\ &= d_0 + (B + S \cdot t_e) \cdot f(T). \end{aligned} \quad (2)$$

To obtain the dark corrected image, we obviously have to subtract the dark current contribution from the measured data number, i. e.,

$$D_{\text{corr}} = D - D_{\text{dark}}. \quad (3)$$

If we have a dark image (or several) taken with the same temperature and exposure time as the image we want to correct, that is quite straight forward. However, if we only have dark images available which have been taken with a different exposure time or at a different temperature, things get a little bit more complicated.

2.1.2 Dependence on exposure time

Besides the known offset d_0 , there are two contributions from the dark current. One, the bias B , is independent of the exposure time, while the other, $S \cdot t_e$, grows linearly with the exposure time. If possible, one can take a dark image with zero exposure time, which gives directly the value of B for each pixel. If we subtract this image from another dark image taken with non-zero exposure time and divide by the exposure time, we obtain the value of S for each pixel.

In fact, we do not necessarily need one image with zero exposure time. If we just have two images with different exposure times, a simple line (for each pixel) through the two points given by the exposure times and the respective measured data numbers provides the value at zero exposure time $B + d_0$ (including the offset) and the slope S . Because of the unknown dark noise, using just two images may give results which are not very accurate. Ideally, we have many images taken with a wide range of exposure times. A line fit through many such points is less affected by noise errors of individual points.

2.1.3 Dependence on CCD temperature

The dark current increases in a non linear way with the temperature of the CCD. The theoretically expected temperature dependence is

$$f(T) = \left(\frac{T}{T_0}\right)^{\frac{3}{2}} \cdot \exp\left(\frac{E_g(T_0)}{2 \cdot k \cdot T_0} - \frac{E_g(T)}{2 \cdot k \cdot T}\right), \quad (4)$$

where T_0 is a reference temperature (we shall choose below $T_0 = 273.15$ K), $k_0 = 8.6171 \cdot 10^{-5}$ eV/K is the Boltzmann constant, and

$$E_g = 1.11557 \text{ eV} - \frac{7.021 \cdot 10^{-4} \cdot T^2}{1108 \text{ K} + T} \cdot \frac{\text{eV}}{\text{K}} \quad (5)$$

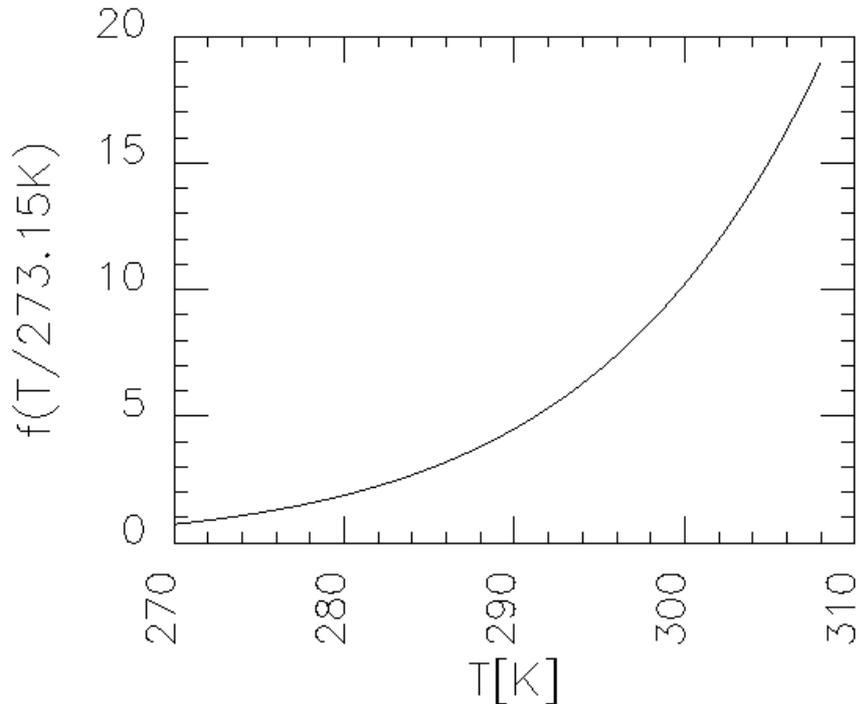


Figure 1: Temperature dependence of the dark current on the CCD temperature as theoretically expected.

is the silicon band gap energy. The temperature dependence is illustrated in Fig. 1.

2.2 Master dark frames from in flight dark sky observations

2.2.1 The set of dark sky images

The original set of dark sky images comprised 547 (full frame) images. The images had been selected by the PI team automatically. The selection took place at a time point when not all images had been unpacked from the DDS. About 14,000 images have been unpacked later, however, a fraction of these (probably around 30 %) may have been empty anyway (cf. the remark on empty images in the `DATASET.CATALOG` of the AMIE dataset).

For many of the dark sky images, no temperatures were initially available, therefore these were not used to estimate the dark current behaviour. Later, temperatures for almost all images could be modelled, see the corresponding document in the ‘DOCUMENT’ directory of the AMIE dataset.

After excluding all empty images, images with telemetry errors, images with visible straylight, and images with lacking temperature information (at that time), we were left with 259 dark sky images. If the dark correction should be reassessed, we recommend to start out with a new search for dark sky images.

The location of these 259 dark sky images in the parameter space spanned

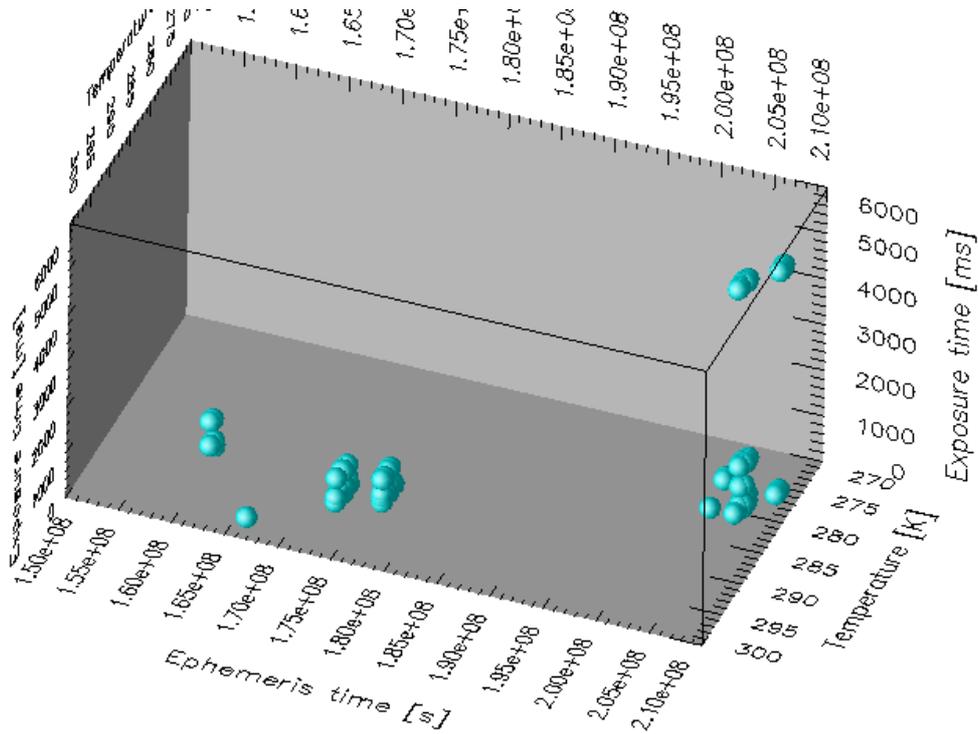


Figure 2: Location of 259 selected dark sky images in the parameter space spanned by mission time, exposure time, and focal plane temperature.

by mission time, exposure time, and focal plane temperature is shown in Fig. 2. Projections showing the location in plots of temperature versus mission time, exposure time versus mission time, and exposure time versus temperature are shown in Figs. 3, 4, and 5, respectively. We can distinguish two sets of dark sky images:

- Early lunar set:
 - 127 dark sky images
 - Orbits 113–977
 - High temperatures, around 290 K
- Late lunar set:
 - 132 dark sky images
 - Orbits 2472–2764
 - Low temperatures, around 280 K

The mean “brightness” (data number averaged over all pixels excluding the 64 lines closest to the CCD readout area; these excluded lines are discussed in section 5.2) of each of the dark images is illustrated in Fig. 6. The dark current

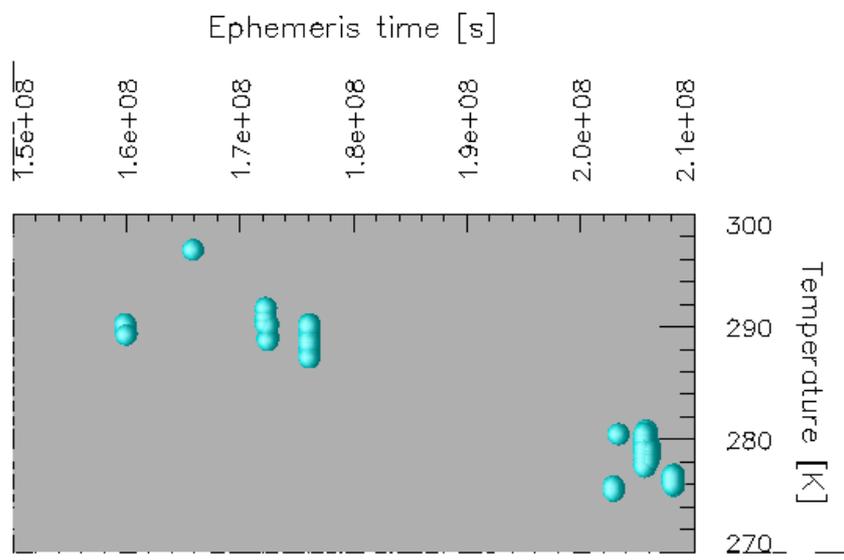


Figure 3: Location of 259 selected dark sky images in the parameter plane spanned by mission time and focal plane temperature.

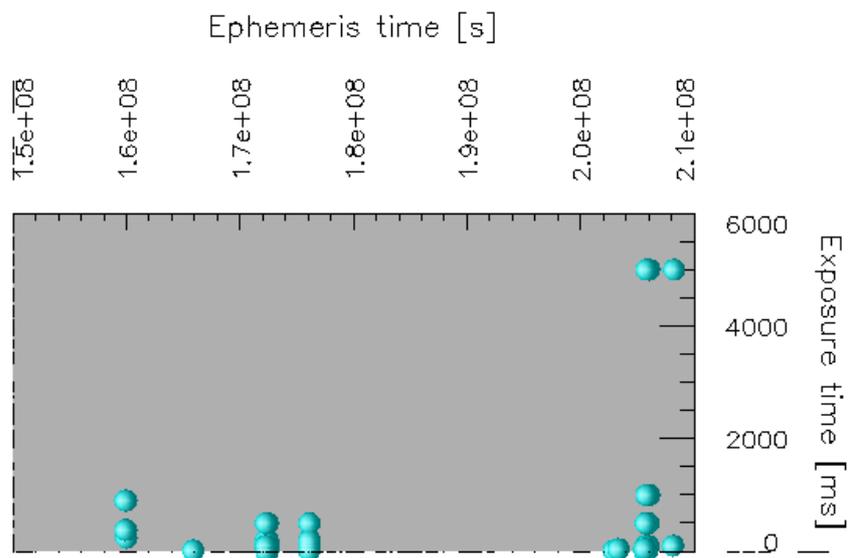


Figure 4: Location of 259 selected dark sky images in the parameter plane spanned by mission time and exposure time.

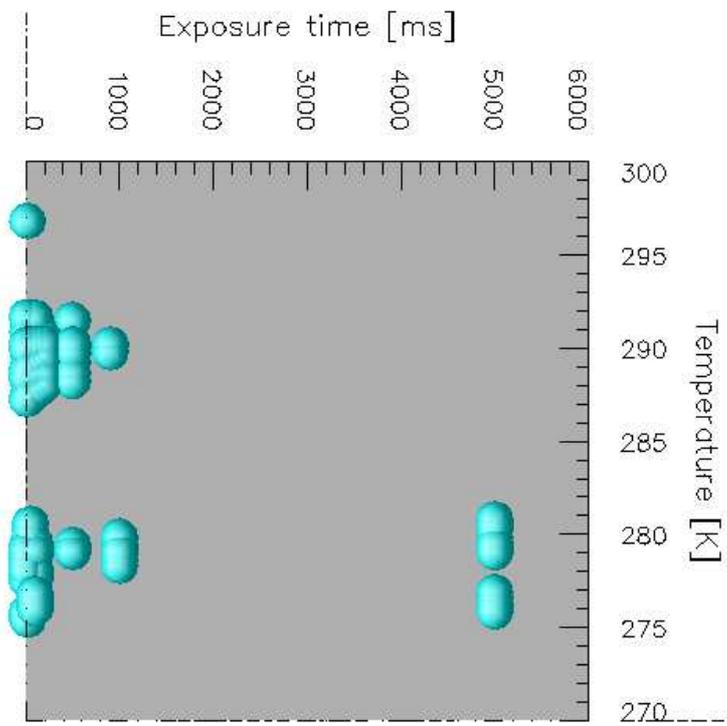


Figure 5: Location of 259 selected dark sky images in the parameter plane spanned by exposure time and focal plane temperature.

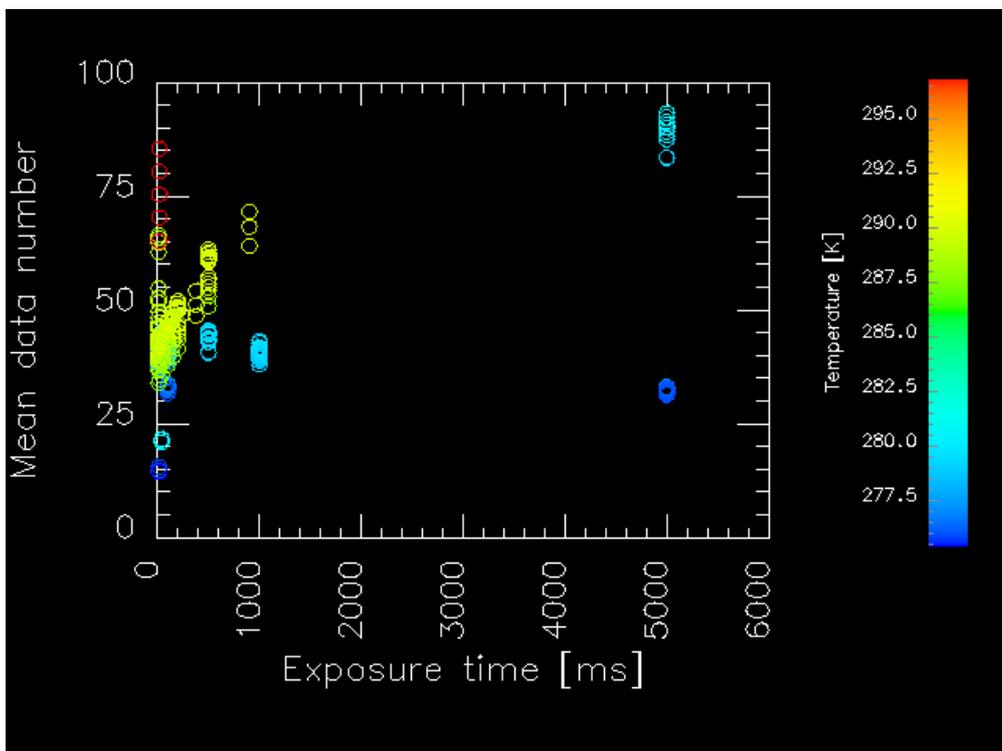


Figure 6: The mean “brightness” of the 259 dark sky images in dependence on exposure time. The color of the marks illustrates the focal plane temperature.

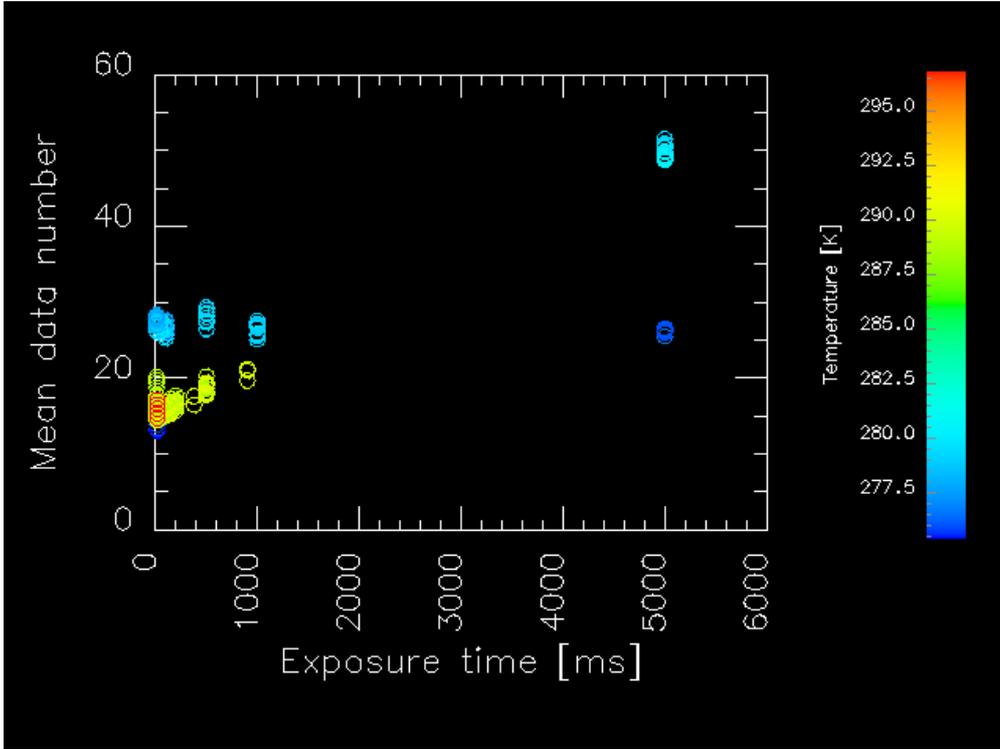


Figure 7: The mean “brightness” of the 259 dark sky images in dependence on exposure time, scaled to the same focal plane temperature. The color of the marks illustrates the original temperature.

contribution given for an individual pixel in Eq. (2) can similarly be formulated for the image mean:

$$m(D_{\text{dark}}) = d_0 + (m(B) + m(S) \cdot t_e) \cdot f(T), \quad (6)$$

where $m(\dots)$ denotes averaging over the image. Assuming that the theoretical temperature dependence as given by Eq. (4) holds, we can scale the images according to their respective focal plane temperature:

$$\begin{aligned} m(D_{\text{scaled}}) &= \frac{m(D_{\text{dark}}) - d_0}{f(T)} + d_0 \\ &= d_0 + m(B) + m(S) \cdot t_e. \end{aligned} \quad (7)$$

Now we have eliminated the temperature dependence. The scaled image means $m(D_{\text{scaled}})$ should depend linearly on the exposure time t_e . These values are plotted in Fig. 7. While most images are close to a line ranging from about (0,15) to (5000,50), there are various images with data numbers about 25–30 which do *not* fit on this line. These are 105 images in the late lunar set which have been commanded with exposure times 10, 50, 100, 500, 1000, and 5000 ms, but which all look as if they have been acquired with an exposure time of about 1500 ms. We drop these 105 suspicious images and keep only 27 from the late lunar set, leaving us (together with the 127 from the early lunar set) with 154 images in total.

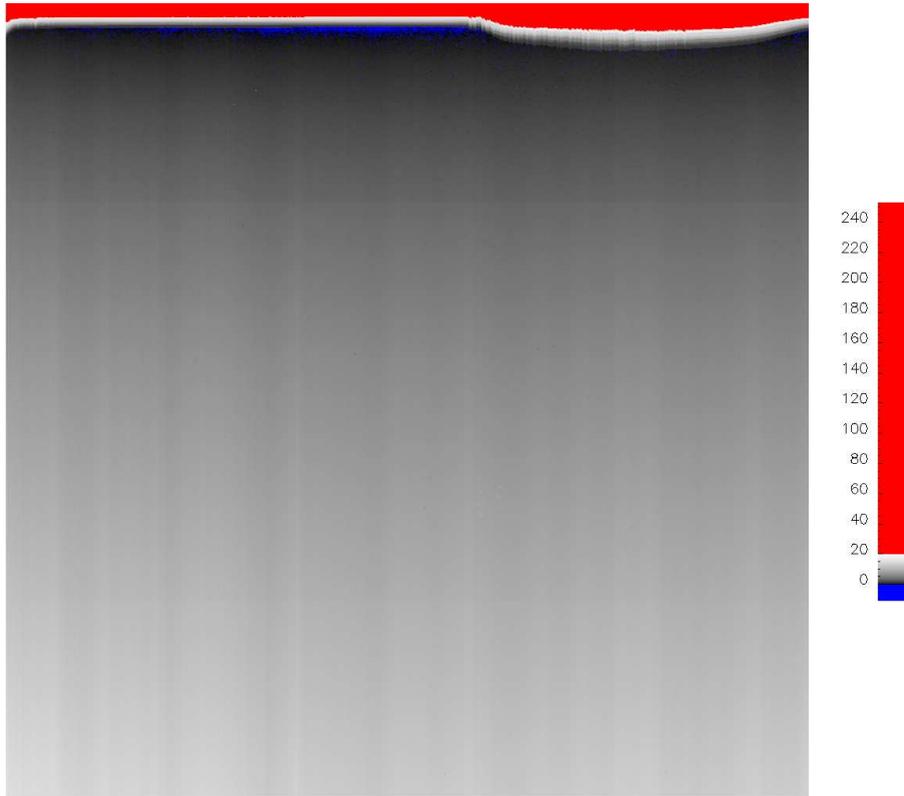


Figure 8: The master frame for the bias which is built up by dark current during the read out of the CCD. The bias frame is independent of exposure time. The units are data numbers scaled for a temperature of 273.15 K.

2.2.2 Estimation of master darks

The master dark frames are estimated pixel by pixel. For each pixel (i, j) , we have 154 values from the 154 dark sky images. The values from all the images $k = 1 \dots 154$ are scaled according to their temperature (similarly as in Eq. (7), but now for each pixel individually):

$$\begin{aligned}
 D_s(i, j, k) &= \frac{D(i, j, k) - d_0}{f(T(k))} \\
 &= B(i, j) + S(i, j) \cdot t_e(k).
 \end{aligned} \tag{8}$$

For each pixel (i, j) , we fit a line $B(i, j) + S(i, j) \cdot t_e(k)$ through the points $(t_e(k), D_s(i, j, k))$ for $k = 1 \dots 154$. The resultant frames B and S represent bias and dark current. They are presented in Figs. 8 and 9, respectively. The corrupted band at the top which is adjacent to the CCD read out area is discussed in section 5.2.

The bias frame B is available as image in PDS format as the file

`'AMI_LMA_071101_00001_00000.IMG'`

in the 'CALIB' directory of the AMIE dataset. The rightmost field of digits in the file name is reserved for the exposure time, and it is set to '00000', because

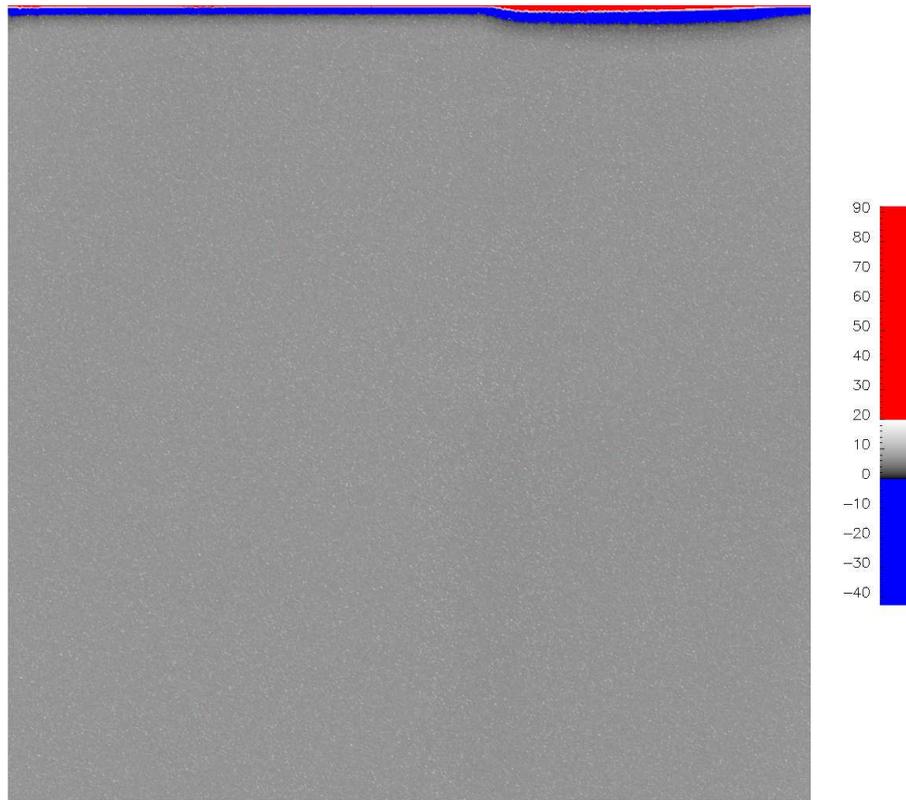


Figure 9: The master frame for the charge which is built up by dark current during the exposure. The dark current frame scales linearly with the exposure time. The units are data numbers scaled for an exposure time of 1000 ms and a temperature of 273.15 K.

the bias frame does only comprise the contribution of the dark current during CCD read out and thus represents an image with zero exposure time.

The dark current frame S is available as image in PDS format as the file

`'AMI_LMA_071101_00002_00001.IMG'`

in the 'CALIB' directory. Here the field reserved for the exposure time is set to '00001', because the frame is scaled to represent the charge built up by the dark current during 1 ms of exposure time.

2.2.3 Explained variance

Our estimation of the master dark frames rests on the assumption that the temperature dependence of the dark current follows the theoretical expectation as given in Eq. (4). To assess the accuracy of this assumption, we compute the explained variance of the 154 dark sky images by the master frames.

For each dark image k , we can compute a "model" image

$$D_{\text{mod}}(k) = d_0 + (B + S \cdot t_e(k)) \cdot f(T(k)) \quad (9)$$

from the master frames. The variance explained by the model images is

$$v_{\text{exp}} = 1 - \frac{\langle (D - D_{\text{mod}})^2 \rangle}{\langle (D - \langle D \rangle)^2 \rangle} = 98 \%, \quad (10)$$

where $\langle \dots \rangle$ denotes averaging over all pixels (i, j) and all dark images $k = 1 \dots 154$. The very high explained variance of 98 % supports the assumption of the theoretical temperature dependence.

The RMS deviation between the observed and the modelled images is

$$\sqrt{\langle (D - D_{\text{mod}})^2 \rangle} = 3.5. \quad (11)$$

This value of 3.5 data numbers represents the average dark noise.

2.3 Applying the dark correction

2.3.1 Test: Dark correcting a dark image

If D represents the data numbers of the pixels of the raw image, the dark corrected image is obtained by

$$D_{\text{corr}} = D - (d_0 + (B + S \cdot t_e) \cdot f(T)). \quad (12)$$

In order to assess the dark correction using the master frames, we apply it to a dark sky image, particularly, to image number 7 from orbit number 976, which resides in the raw data set of the lunar phase in the files

`'AMI_LE?_R00976_00007_00500.IMG'`,

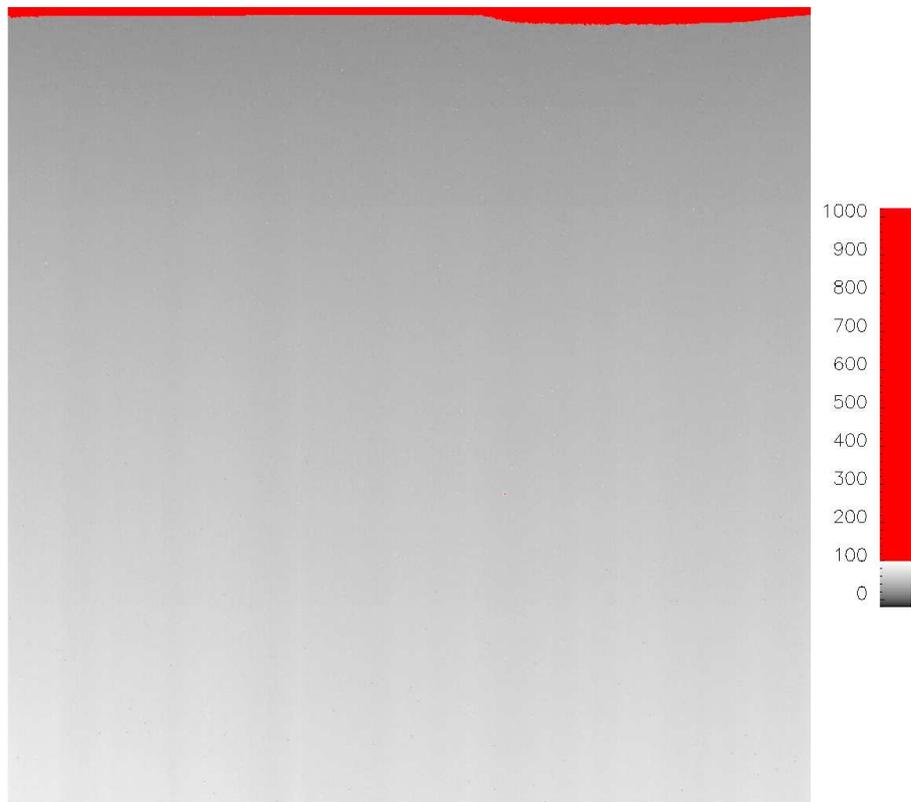


Figure 10: Raw dark sky image from orbit 976, number 7.



Figure 11: Dark corrected dark sky image in the same gray scale as the raw image shown in Fig. 10.

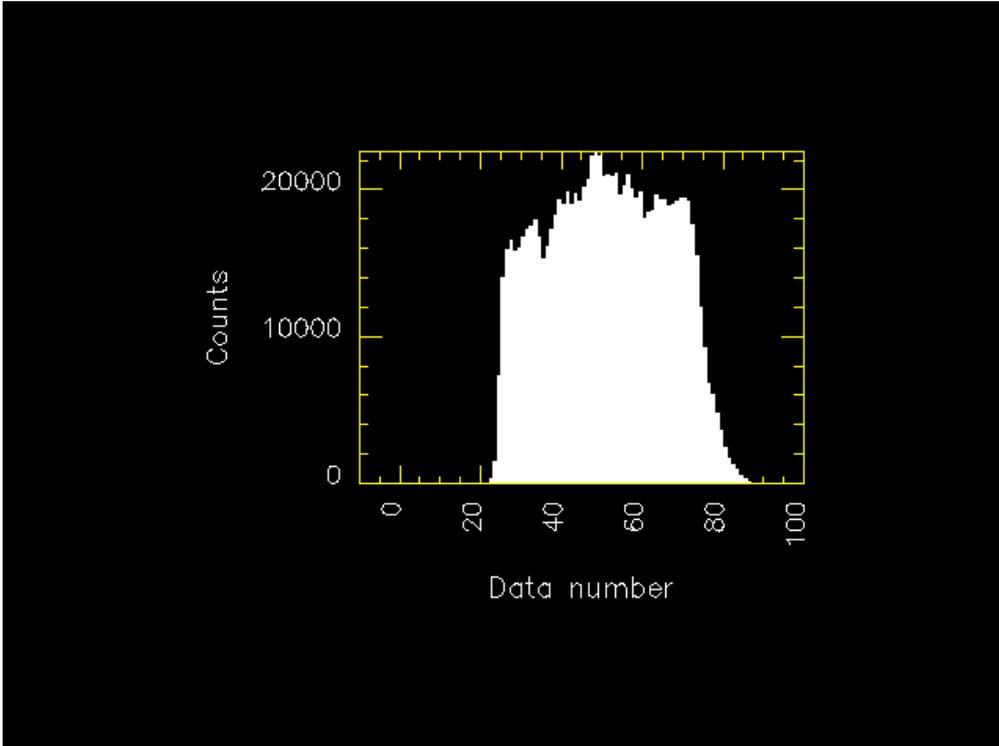


Figure 12: Histogram of the raw image shown in Fig. 10.

where “?” runs over the filter numbers 1–8. The raw image is shown in Fig. 10. The dark corrected image is shown (on the same gray scale) in Fig. 11. The bright dot to the lower right of the image center is a bright star. There are stars visible in many dark sky images, however, when computing the master frames from 154 dark sky images, these are averaged out almost completely. In fact, when looking very closely to the area near the bright dot, we may notice a couple of faint dark dots surrounding the bright one; these are remnants from the bright star, which was at different locations in some of the dark sky images used to compute the master frames.

The histogram of the raw image is shown in Fig. 12. We notice the clear offset from zero which is due to the dark current. The histogram of the dark corrected image is shown in Fig. 13. Now the offset from zero is removed and moreover, the variance is substantially decreased. Fig. 14 shows the histogram of the dark corrected image on an expanded scale. We notice that the histogram is nicely centered around the zero level and that its width is consistent with the estimated average dark noise of 3.5 data numbers, cf. Eq. (11).

2.3.2 Example: Dark correcting a lunar image

As example image, we use image number 33 from orbit 1845, which resides in the raw data set of the extended mission in the files

`‘AMI_LE8_R01845_00033_00030.IMG’`,

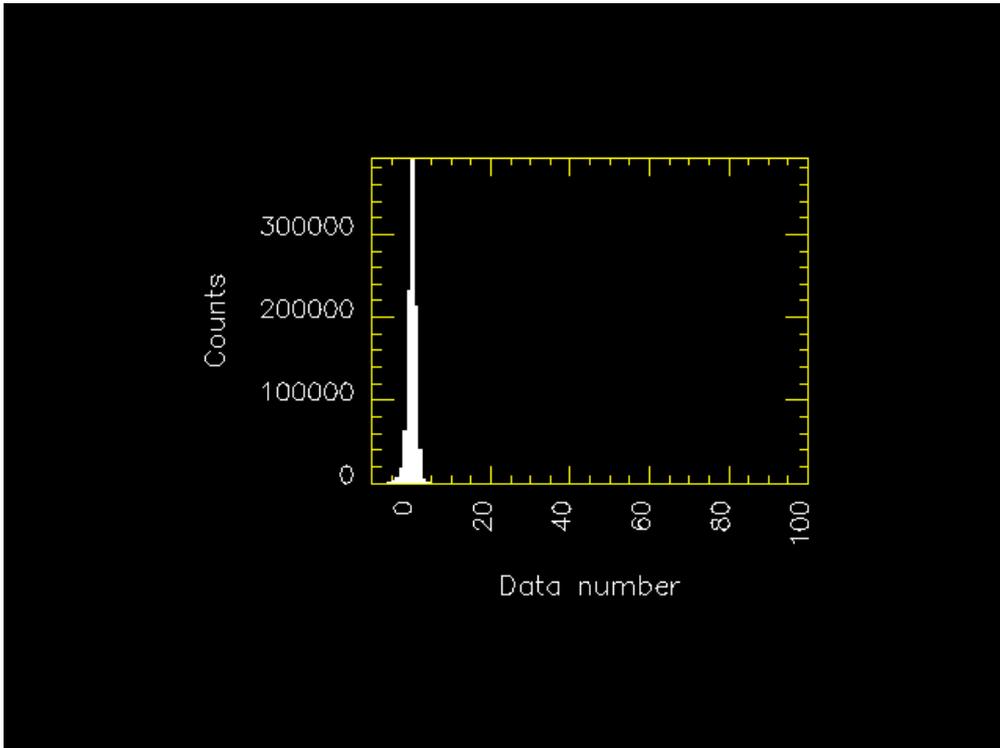


Figure 13: Histogram of the dark corrected image shown in Fig. 11.

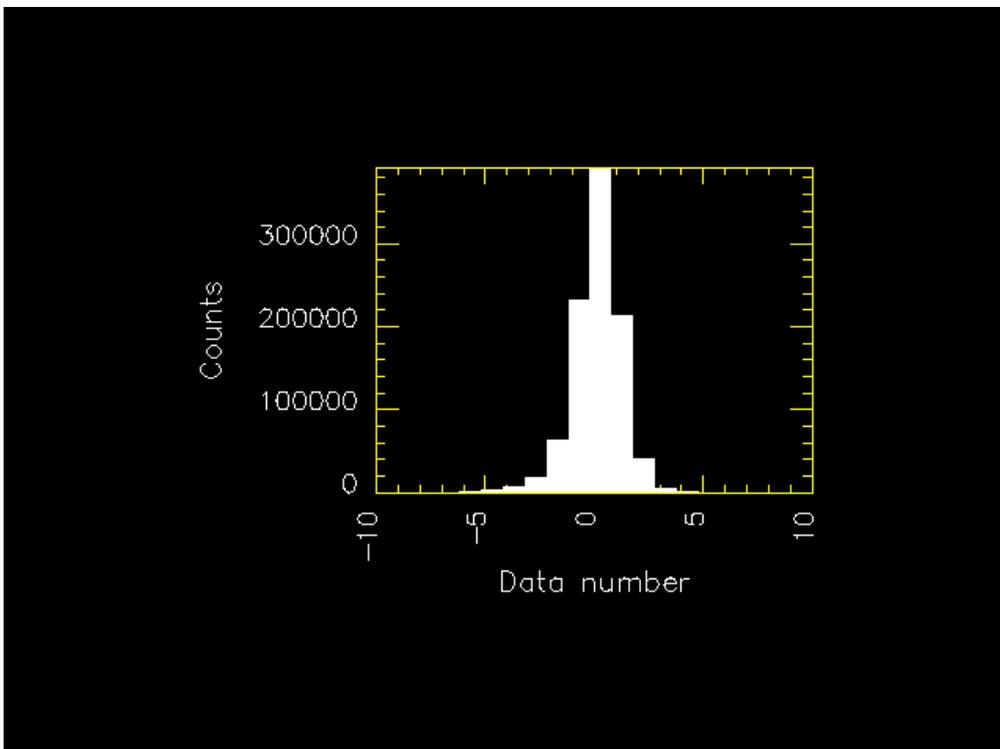


Figure 14: The same histogram as in in Fig. 13, but with an expanded scale.

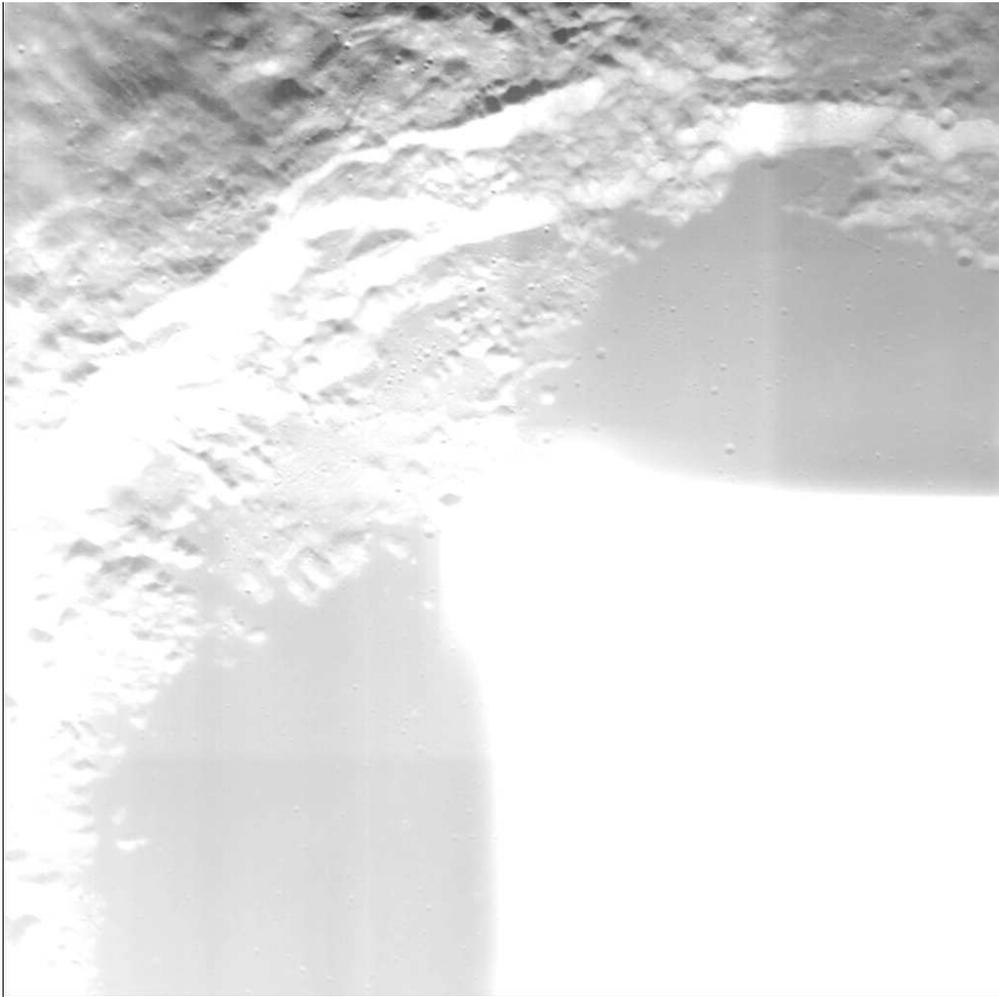


Figure 15: Raw example image from orbit 1845, number 33.

where where “?” runs over the filter numbers 1–8. The raw image is shown in Fig. 15. We notice a brightness gradient from top to bottom, e.g., the filter `VIS_Y` in the lower left corner is considerably brighter than the filter `VIS_X` in the upper right corner, while these are similar filters which should exhibit the same brightness. Additionally, we notice a somewhat smooth irregular vertical pattern, particularly in the lower left of the image. Both are features of the bias frame, cf. Fig. 8.

We do not notice much of the noisy pattern of the dark current slope frame. In Fig. 9, the dark current frame is scaled to an exposure time of 1000 ms, and then its level is comparable to the bias frame. However, the exposure time of our example image is only 30 ms, and therefore the contribution from the slope frame is very small.

The dark corrected example image is shown in Fig. 16. Now the vertical gradient and the smooth irregular vertical pattern are removed. The raw image and the dark corrected image are shown in the same gray scale. The large brightness difference of the two images illustrates the significance of the dark

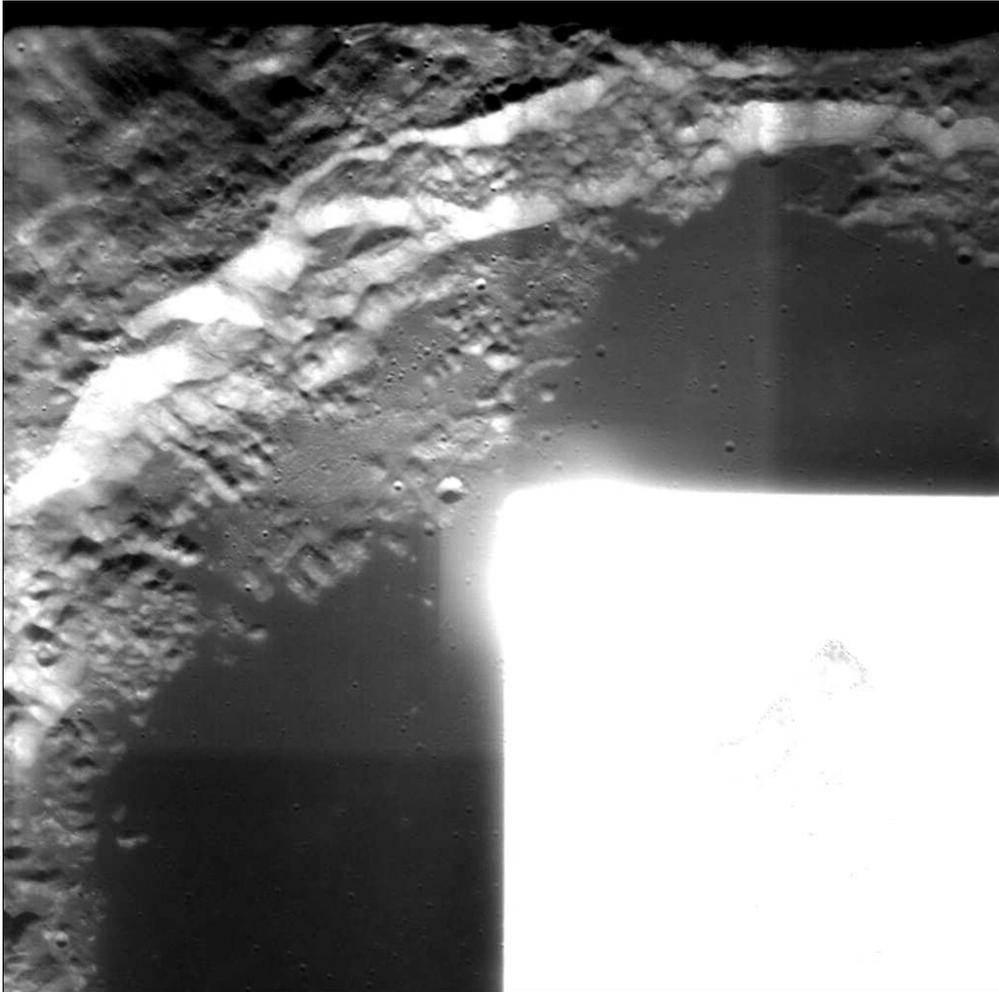


Figure 16: Dark corrected example image in the same gray scale as the raw image shown in Fig. 15.

current contribution.

The dark correction based on the master frames from in flight dark sky observations seems to work well for all images from lunar phase, extended mission, and for images from the late Earth escape phase. It seems *not* to work well for the early Earth escape phase. This has not been investigated systematically.

3 Flat fielding

3.1 The theory of flat fielding

The data number D measured for a particular CCD pixel was given in Eq. (1). After dark correction, cf. Eqs. (2, 3), we are left with

$$D_{\text{corr}} = C \cdot I \cdot t_e, \quad (13)$$

where C is the radiance conversion factor, I is the incoming radiance, and t_e is the exposure time. A laboratory flat field is usually taken in an integrating sphere with a known incoming radiance I_0 , which is the same for all pixels, and a certain exposure time t_{e0} , thus

$$D_{\text{flat}} = C \cdot I_0 \cdot t_{e0}. \quad (14)$$

If we then have an actual (dark corrected) image D_{corr} taken with exposure time t_e , we can recover the incoming radiance by dividing by the flat field:

$$I = \frac{D_{\text{corr}}}{D_{\text{flat}}} \cdot \frac{t_{e0}}{t_e} \cdot I_0. \quad (15)$$

3.2 A flat field from in flight data

Because we know that the dark current behaviour of the CCD changed significantly during the cruise phase of the mission, we want to compute also a flat field from in flight data and compare it to the laboratory flat field. The basic idea is to average so many images that all image features are averaged out and that the average incoming radiance is the same for all pixels.

However, this approach is not as straightforward as one might think. The problem is saturation, of which we have a lot; in every image with an exposure time appropriately set for the filter areas, the NONE filter area is completely saturated. If in some of the images which we use for averaging NONE filter pixels — or, more general, some of the pixels with higher responseivity — are saturated, the mean of all images underestimates highly responsive pixels. This problem could not be solved by just treating the filters separately and doing without equalization of the filters. Towards the boundary of the NONE filter, the pixels of the filter frames get very responsive and these pixels are saturated in many filter images.

It is not an option to completely discard all images where some pixels are saturated. This would leave us with only a very small number of images, which would be preferably dark, i. e., exhibiting a small signal to noise ratio, particularly for the filter areas. Therefore we could not expect to effectively average out image features.

The next idea might be just to discard saturated pixels. This means that for a particular pixel, we average all images where this pixel is not saturated. So we discard a certain fraction of the brightest images. However, for a more responsive pixel, this pixel is saturated in more images, and thus we discard a larger fraction of the brightest images for more responsive pixels. Therefore we underestimate the brightness of highly responsive pixels.

The solution is obviously to first scale all images to the same mean brightness. Then we would not any more use preferably darker images for highly responsive pixels. However, to scale the images, we have to first compute their mean, and here the problem of saturation strikes again. For partly saturated images, the mean underestimates the brightness. Therefore such images would be scaled too bright. While they are discarded for highly responsive pixels which are saturated, they are used for less responsive pixels. This yields an overestimation of the brightness of less responsive pixels.

Therefore we do not use the mean, but the median, to first scale all images. The assumption is that for a large number of data — 1 million pixels in our case — and a not too strange distribution, the median should be a good approximation for the mean. And oppositely to the mean, the median is not affected by saturated pixels until their fraction becomes more than 50 %.

When looping through all (full frame) images of lunar phase and extended mission to compute the flat field, we apply the following steps for each image:

1. Mark saturated pixels invalid. This has to be done before dark correction. We assume pixels with data numbers of 960 or higher as saturated.
2. Apply dark correction.
3. Mark dark pixels invalid. Dark pixels have a low signal to noise ratio and would make the flat field noisy. We discard pixels with data number lower than 8.
4. Discard the image if more than one third of the pixels is saturated or dark. Only if more than half of the pixels are saturated, the median becomes biased, so we leave some margin here. We also discard images with too many dark pixels because these may be polar images with lots of shadows, which would also bias the median, or it may even be dark sky images.
5. Compute the median. Finding the median of 1 million pixels would take ages if done in the naive way of sorting by repeated passes through the data. We adapt a routine provided by Michel Ollagnon which uses the recursive procedure described in Knuth, *The Art of Computer Programming*, vol. 3, 5.3.3. This is pretty fast.
6. Divide the image by its median.
7. Add the valid pixels to the sum of images. In fact, we do not sum the images, but we update directly the mean of all images processed so far using an algorithm described in Knuth, *The Art of Computer Programming*, vol. 2, p. 232. This is equivalent to summing all images and then

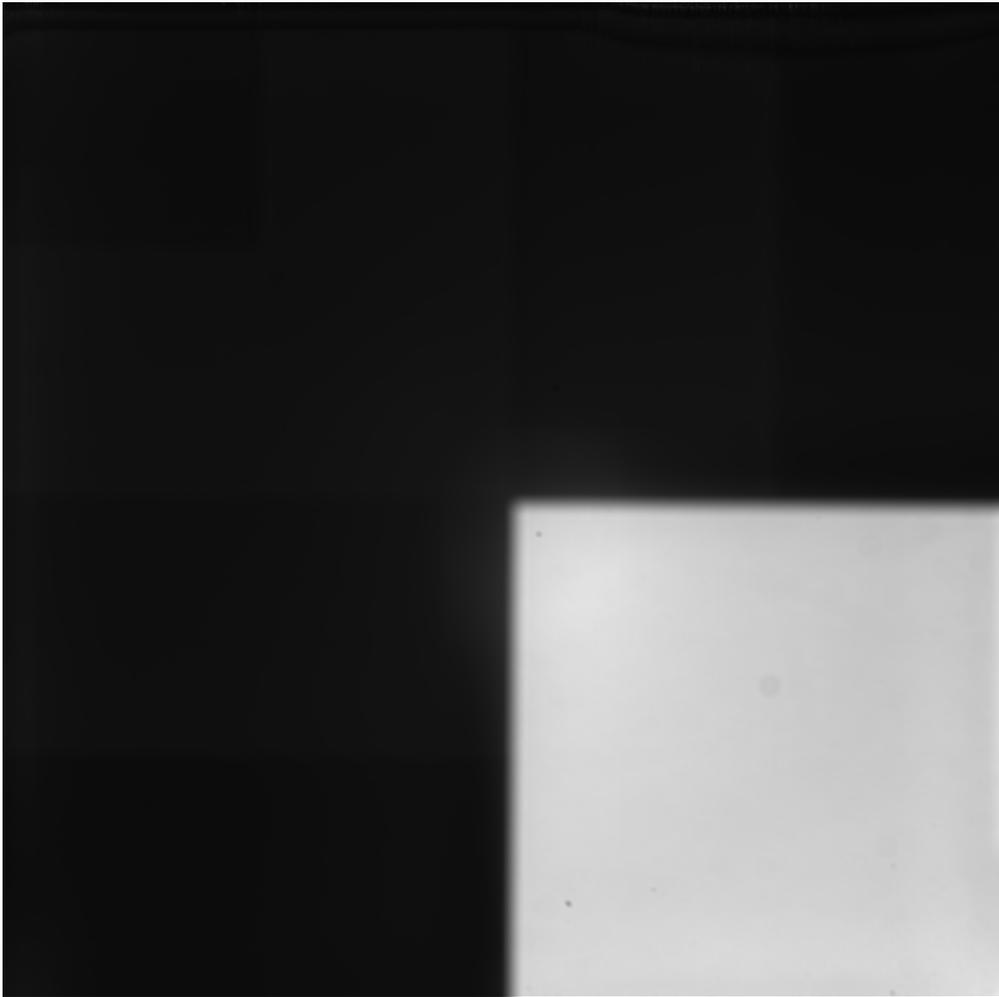


Figure 17: The flat field computed by averaging in flight images from lunar phase and extended mission. The gray scale is adapted to the NONE filter area.

dividing by the number, but it is numerically more stable and does not need a division by the number of images used (for each pixel) at the end.

To summarize the consequences of this algorithm:

- We discard images which have more than one third saturated or very dark pixel.
- Each image is scaled to the same brightness by dividing by its median.
- The flat field is computed as the mean of all images, where saturated or dark pixels are omitted.

The resultant flat field is shown in two different gray scales in Figs. 17 and 18. The corrupted band at the top is discussed in section 5.2.

The flat field estimated from in flight data is available as image in PDS format as the file

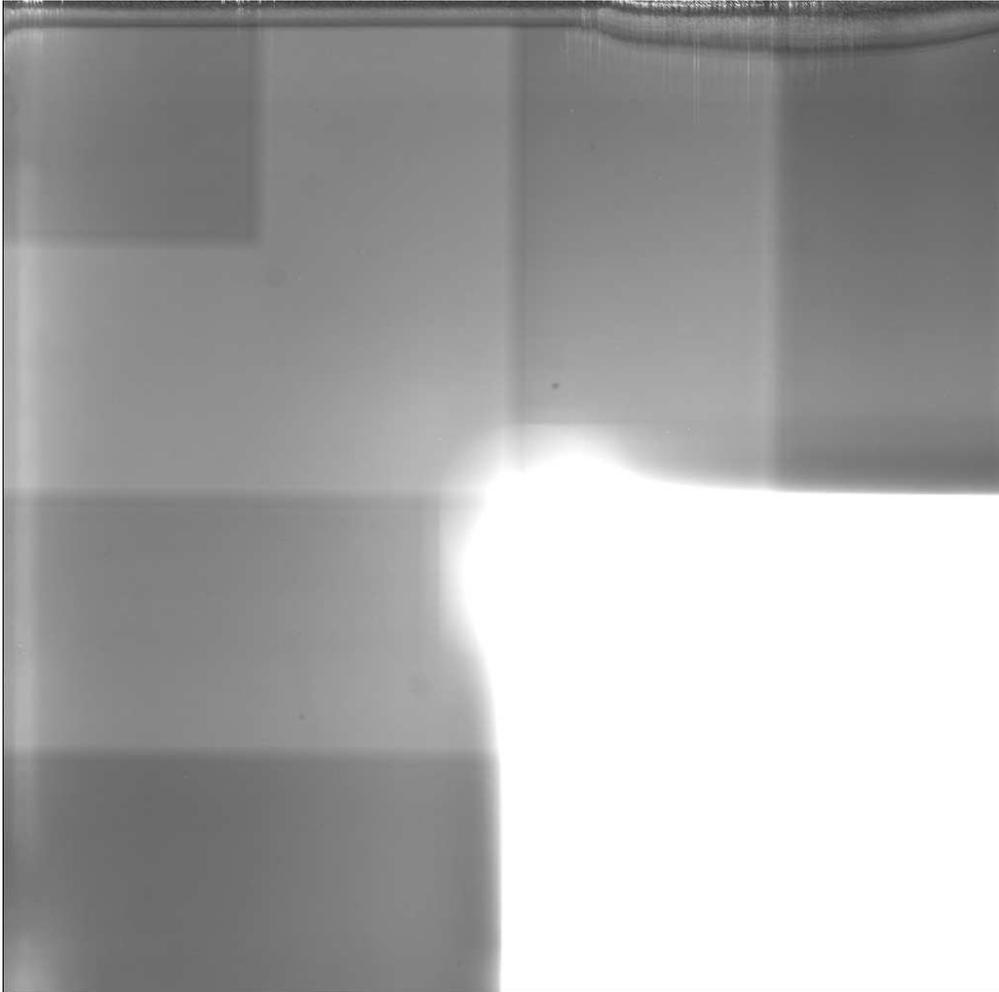


Figure 18: The flat field computed by averaging in flight images from lunar phase and extended mission. The gray scale is adapted to the filter areas.

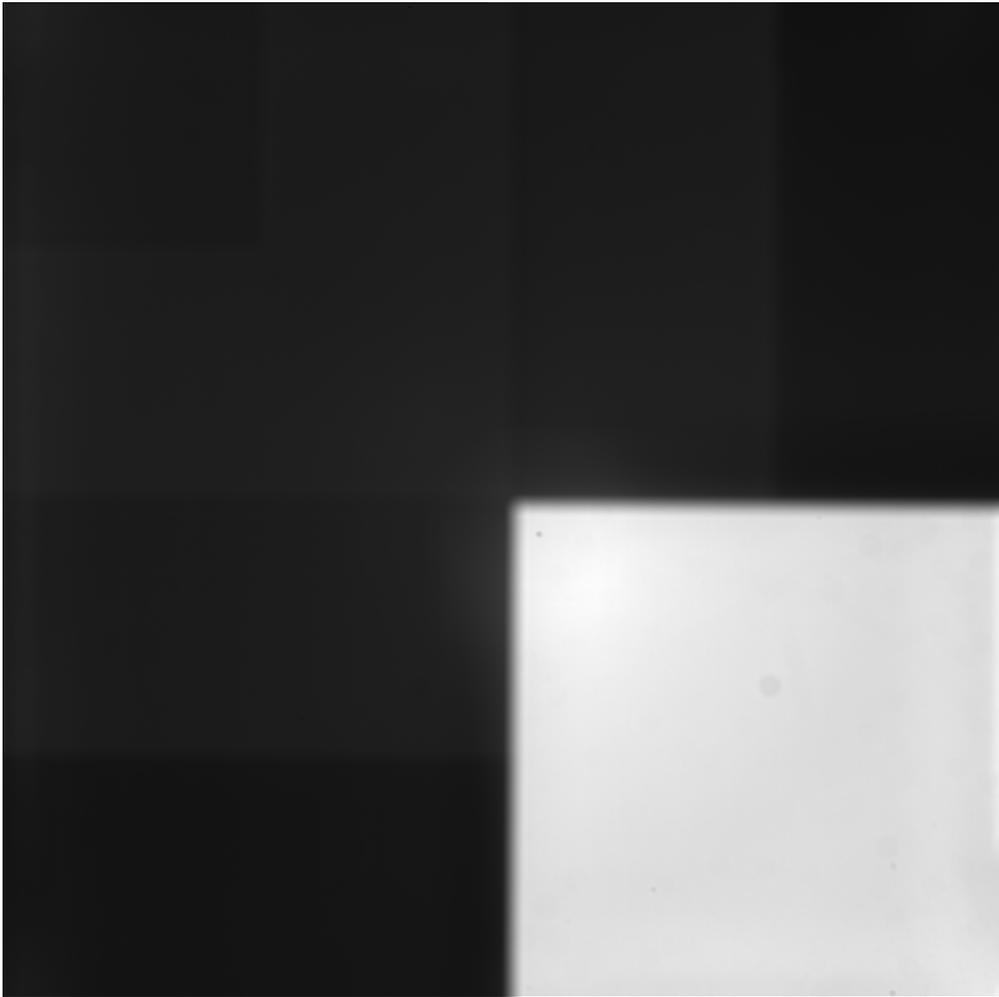


Figure 19: The flat field computed from laboratory images. The gray scale is adapted to the NONE filter area.

```
'AMI_LMA_080319_00001_XXXXX.IMG'
```

in the 'CALIB' directory. A field of five digits in the file name is reserved for the exposure time. Because this flat field does not represent a particular exposure time, this five character field has been padded with 'XXXXX', similarly as for the images with unknown exposure time, cf. section 5.1.

3.3 Comparing space and lab flat

The flat field estimated from laboratory measurements is shown in two different gray scales in Figs. 19 and 20. When comparing the lab flat to the space flat, we notice a few additional defects in the latter. Thus there is some change during the Earth escape phase, but it is not as dramatic as the change in the dark current behaviour.

The lab flat exhibits some smooth vertical pattern which is also visible in

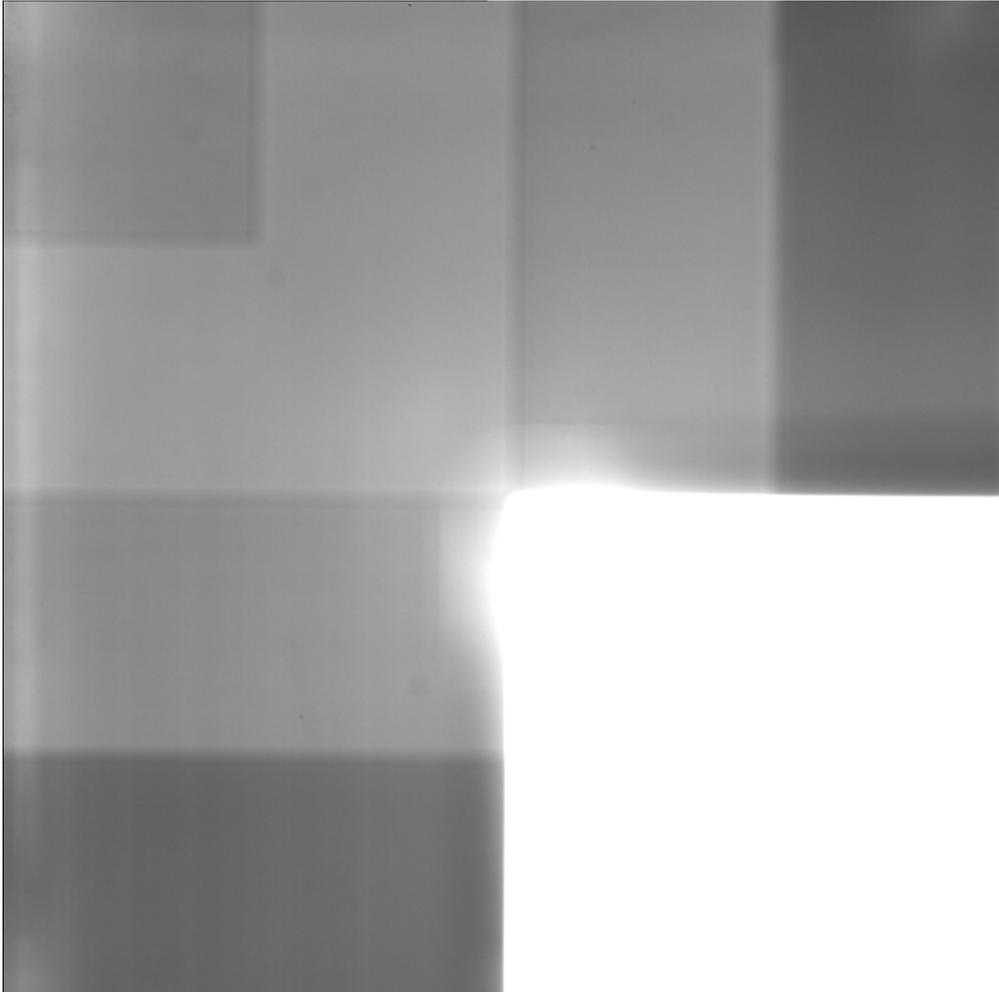


Figure 20: The flat field computed from laboratory images. The gray scale is adapted to the filter areas.

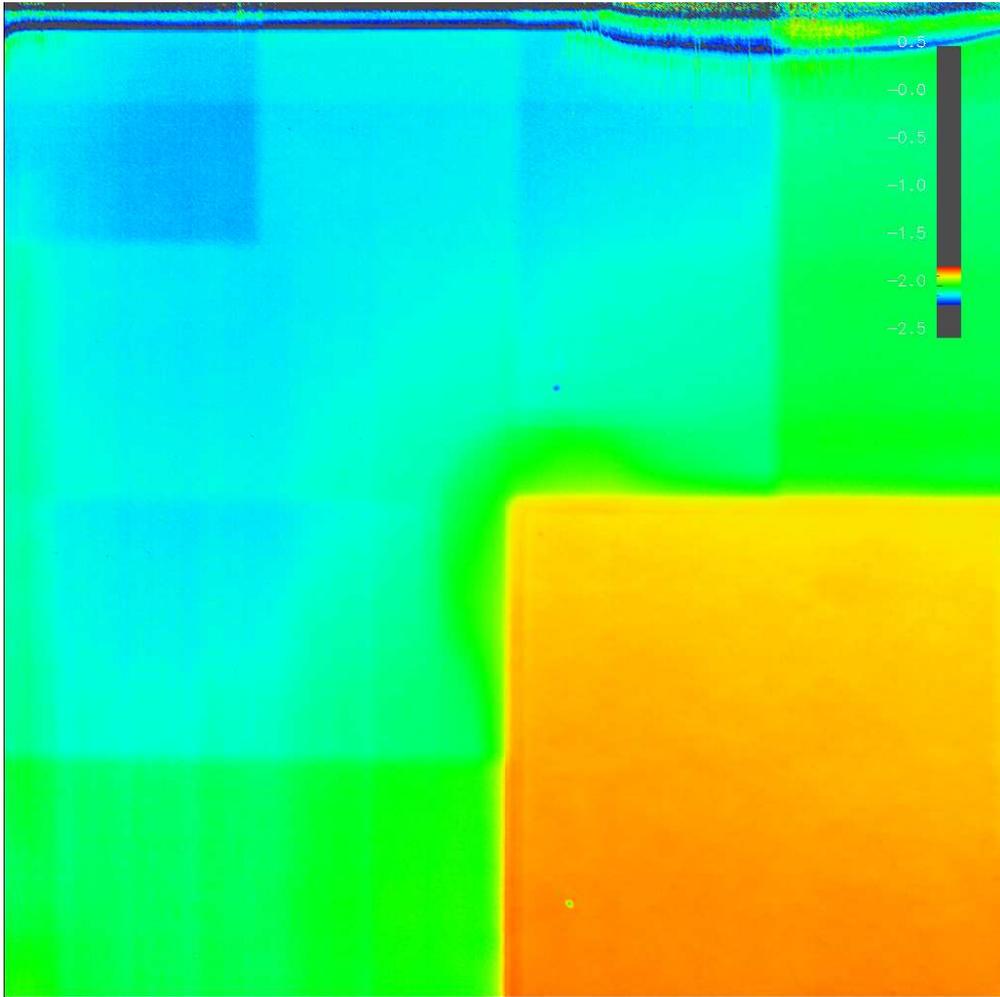


Figure 21: Logarithm of the ratio of space flat (Figs. 17 and 18) and lab flat (Figs. 19 and 20).

the dark bias, see Fig. 8. This is because when computing the lab flat, the dark correction was applied in a simplified way using only line averages.

The most striking difference between lab and space flat is that the brightness difference between filter and **NONE** filter area is much larger for the space flat. The reason is that the spectrum of the light from the lunar surface is different from the spectrum of the illumination used in the laboratory. One might think that this is not important unless we want to equalize the images from different filters, i. e., as long as we look only on individual filter images. However, because there is some space between the filter and the CCD sensor, the transition from one filter area to the other is smeared out over some ten pixels. The space flat corrects the strong brightness increase or decrease towards a filter edge much better than the lab flat.

The ratio between the space flat and the lab flat is show in Fig. 21. A scatter plot of the brightnesses of all pixels in the two flat fields is shown in

Fig. 22.

3.4 Applying the flat fielding

Flat fielding means just dividing the image by the flat field times exposure time, see Eq. (15). We apply flat fielding with the lab flat and with the space flat to the example image which has also been used to demonstrate the dark correction in section 2.3.2. The result of flat fielding with the lab flat is shown in Fig. 23, the result of flat fielding with the space flat is shown in Fig. 24. We notice that the space flat does much better equalize the different filters. The fact that we still note the filter boundaries is due to scattered light inside the camera, cf. section 5.4.

4 Performing the calibration

4.1 Cookbook for home grown calibration

4.1.1 Reading an image

The information needed to extract the binary image data from the PDS image file is provided in the ASCII file label. One peculiarity is the pointer to the image object. It says (the same in *all* files of the dataset) 36865 bytes. This means — in agreement with the PDS standard — the number of the first byte of image data in the file *with counting starting at one*. To get to the image data, the number of bytes you have to skip is one less, i. e., 36864.

The data type of a pixel value is `LSB_UNSIGNED_INTEGER` (2 Byte) for all raw images and `PC_REAL` (4 Byte) for all calibrated images. The numbers of rows and columns of the image are given by the keywords `LINES` and `LINE_SAMPLES`, respectively. The keywords `LINE_DISPLAY_DIRECTION` and `SAMPLE_DISPLAY_DIRECTION` are not given explicitly, thus the default values `DOWN` and `RIGHT`, respectively, apply.

4.1.2 Dark correction

We assume that the temperature dependence is perfectly as it should be theoretically, cf. section 2.1.3. This implies that we only need one master bias file and one master dark current file. A weighted sum of these two files is subtracted from the image which shall be corrected. The weights depend on the exposure time of the image and on the focal plane temperature at the time of image acquisition.

To perform a dark correction, you need to do the following:

- Implement the function for the temperature dependence of the dark current, see Eq. (4).
- Read in the master bias and master dark current images. They are also available in PDS format, so they can be read in the same way as the AMIE images. See section 2.2.2 for the names of these files.

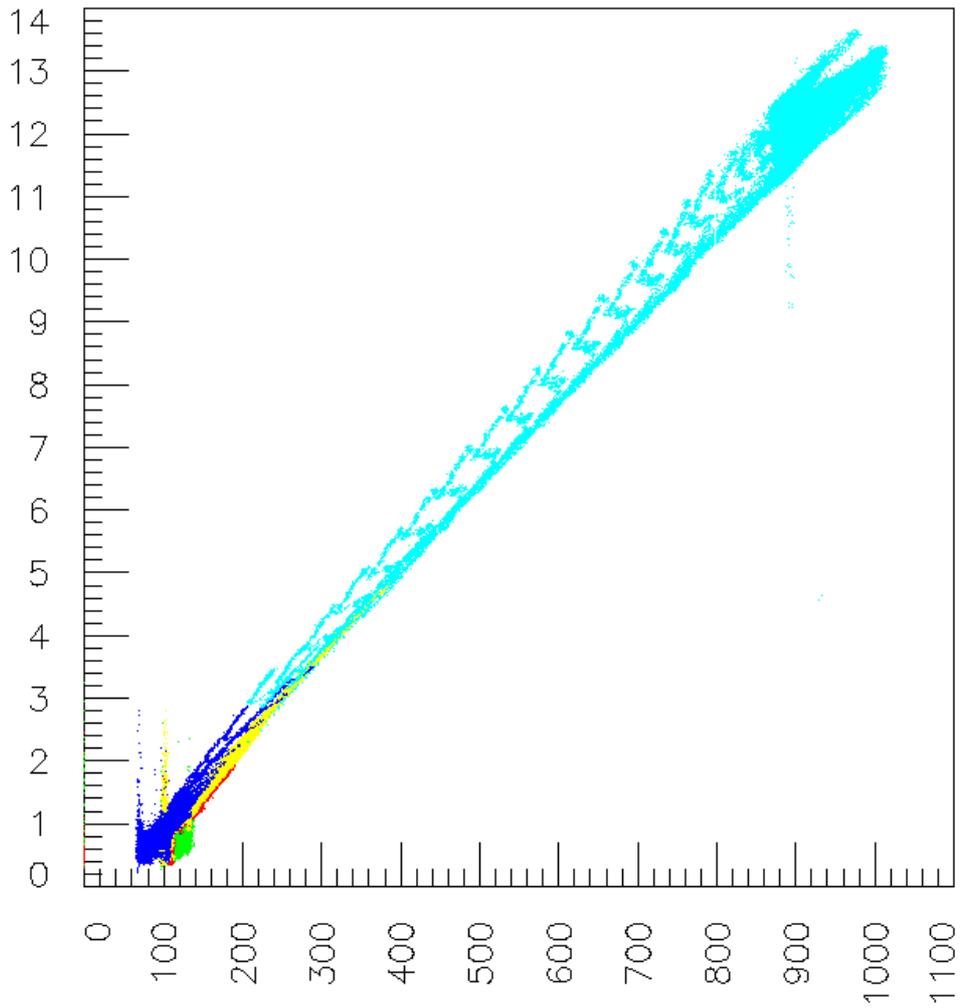


Figure 22: Scatter plot of the brightnesses of all pixels in space flat (vertical axis) and lab flat (horizontal axis). The color represents the filter type (cyan: NONE; blue: VIS; yellow: FeL; red: FeH; green: LASER). The strange regular pattern stems from the transition zone between filter and NONE filter areas.

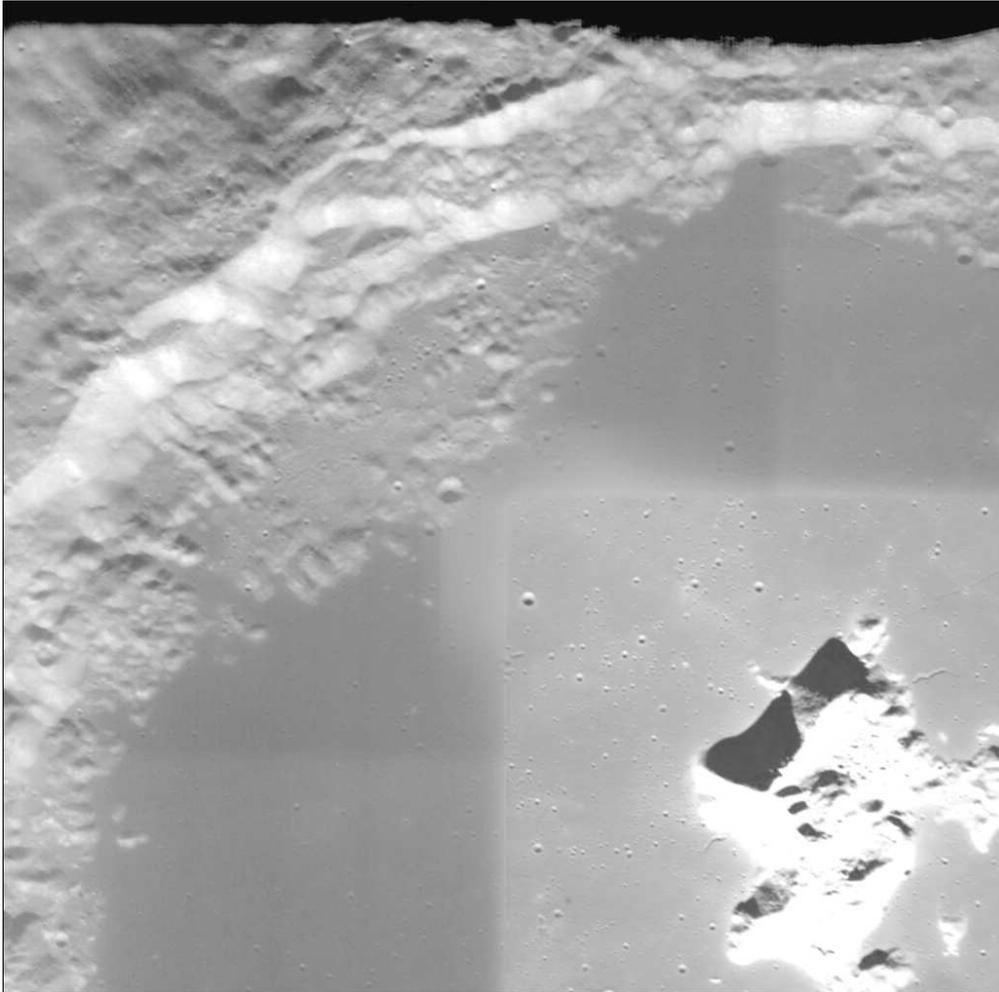


Figure 23: Image from orbit 1845, number 33, flat fielded with the lab flat. The dark corrected image before flat fielding is shown in Fig. 16.

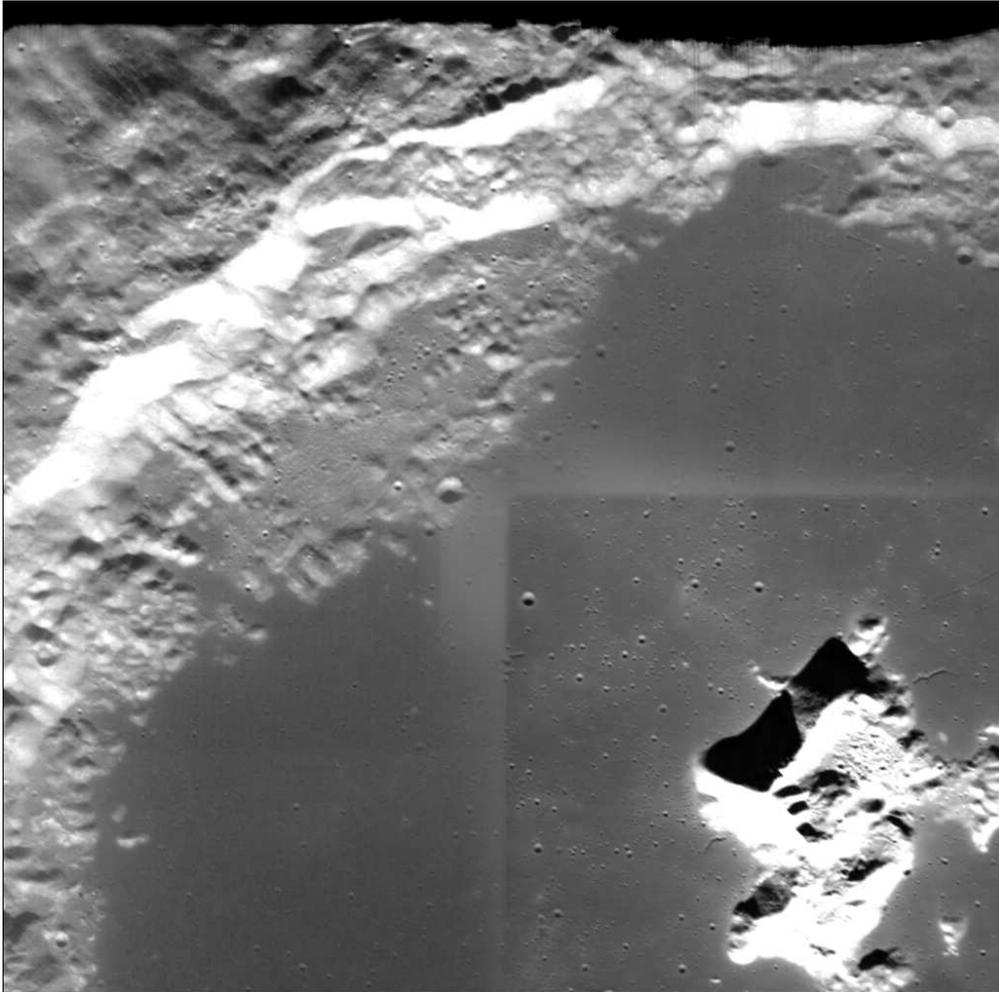


Figure 24: Image from orbit 1845, number 33, flat fielded with the space flat. The dark corrected image before flat fielding is shown in Fig. 16.

- Read in the image which shall be corrected, see section 4.1.1.
- Get the exposure time of the image. The exposure time is part of the image file name, e. g., for

`'AMI_LE3_R02887_00044_00014.IMG'`

the exposure time is 14 ms. The exposure time is also given in the PDS label:

`EXPOSURE_DURATION = 14 <ms>`

- Get the focal plane temperature at the time of image acquisition. The temperature is given in the PDS label:

`FOCAL_PLANE_TEMPERATURE = 290.36 <K>`

- Apply the dark correction, see Eq. (12), cf. section 2.1.

4.1.3 Optional filtering

If the image you are interested in exhibits a disturbing vertical pattern with 8 pixel spacing, you may try to filter this out, see section 5.3.

4.1.4 Flat fielding

Just read in also the flat field, see section 3.2, and divide the image by the flat field times exposure time, see section 3.1.

Calibration done.

4.2 Calibration using the provided IDL routines

4.2.1 Calibrating single images

IDL routines to read and calibrate AMIE images are provided in the 'IDL' subdirectory of the 'DOCUMENT' directory. There are three categories of routines:

1. Routines to read AMIE images,
2. calibration pipeline routines, and
3. readpds routines developed by the PDS Small Bodies Node.

Additionally, there is an example main program:

`'EXAMPLE_CALIBRATION.PRO'`

This program illustrates how to read and calibrate an AMIE image. It calls the same routines as the calibration pipeline, but just for one image.

The example program calls `'READ_AMIE.PRO'` to read an image. Please note that `'READ_AMIE.PRO'` returns a structure containing a browse image in addition to the actual image. For details see the comments in the source code. `'READ_AMIE.PRO'` can read calibrated images as well as raw images. The routine

‘READ_AMIEW.PRO’ just opens a widget to select the input file interactively and then calls ‘READ_AMIE.PRO’.

The routine ‘AMIE_COMBINE.PRO’ reads all filter images belonging to one full frame and creates a combined image. If raw images are combined, the filters may exhibit very different brightnesses. With the option /AUTOSCALE, all filters can be scaled to the same mean brightness and variance.

Both, ‘READ_AMIE.PRO’ and ‘AMIE_COMBINE.PRO’ accept the parameter ORIENTATION = "CCD_FRAME". This has to be used if the image is passed to calibration pipeline routines, but it returns a flipped image. To get the image with correct “handiness”, ORIENTATION = "REAL_WORLD" has to be used (which is the default).

4.2.2 Running the calibration pipeline

The main pipeline routine, ‘AMIE_PIPELINE.PRO’, is called with the name of a configuration file as argument. An example configuration file is also provided. The keywords which are set to a /dummy/ value are not required any more; they were used for calibration based on laboratory master darks and flats.

5 Open calibration issues

5.1 Missing calibration information

For a few images, not all information needed for the calibration is available. For 16 images, there are no corresponding POR files with the imaging commands, and therefore, the exposure time is not known. The commands for these images have been created on the fly in the control room and have not been stored. The field of five digits in the image file name which is reserved for the exposure time (cf. the EAICD in the ‘DOCUMENT’ directory) has been padded with ‘XXXXX’ for these images.

For 70 images, it was not possible to estimate the focal plane temperature at image acquisition time, because no close in time temperature measurement, neither from the AMIE unit 1 sensor nor from the temperature reference point sensor was available, cf. the document on temperature in the ‘DOCUMENT’ directory of the AMIE dataset.

These are the numbers of images for the different mission phases:

	raw	calib	no exp.	no temp.
Earth escape phase	4513	4499	14	–
Lunar phase	13036	13036	–	–
Extended mission	18909	18837	2	70
Total	36458	36372	16	70

5.2 The corrupted band adjacent to the CCD readout area

Many AMIE images exhibit a saturated band at the top, adjacent to the CCD readout area. The reason is that dark charge accumulates in the readout area when the instrument is switched on. This charge is swept away when an image is read out, however, if no image is taken for quite some time, a lot of charge can

accumulate, far beyond the nominal saturation value. If then an image is read out, incomplete charge transfer leaves some remnant charge which is then picked up by image pixels. The amount of such picked up remnant charge decreases exponentially with the distance from the readout area, therefore it is usually negligible. However, if charge has accumulated to a very high level, it becomes significant and yields the saturated band of pixels adjacent to the readout area.

When this problem had been noticed, an image cleaning was introduced before taking an image. This cleaning was just an extra command to take an image without downloading it. The readout of this cleaning image sweeps away the accumulated charge, and the subsequent actual image does not suffer from a saturated band.

Some of the dark sky images used to estimate the master dark frames exhibited a band of saturated pixels. It was not taken special care of when the master darks were estimated, thus the master darks exhibit a corrupted band adjacent to the readout area. Therefore the calibration using these master darks propagates this corrupted band to all calibrated images. Even raw images which did not exhibit a band of saturated pixels get corrupted in the band adjacent to the readout area by the calibration process.

It may be possible to avoid the corrupted band in the master dark frames by discarding saturated pixels in the estimation of the master dark frames. However, the number of valid images for pixels in this band may be too small for an accurate estimation. Therefore we recommend to reassess the set of available dark sky images before revising the procedure of estimating the master dark frames, cf. section 2.2.1.

5.3 Vertical stripes with 8 pixel spacing

Many images exhibit a spurious vertical pattern with 8 pixel spacing, particularly at low light level. An example is shown in Fig. 25. Although the amplitude is very small — about one data number — the pattern can be very disturbing. The origin of this 8 pixel pattern is not known. It was suspected to be related to interference with the triggering of the serial read out of the CCD. The pattern does not occur with cameras by Micro-Cameras & Space Exploration which use a newer version of the firmware.

The problem with the 8 pixel pattern is that its phase — i. e., the location on the CCD — is not constant but varies from image to image. Therefore it is not easy to correct. In fact, it is not visible at all in the dark current frames (Figs. 8 and 9), because it was averaged out when the dark frames were estimated from 154 images.

The 8 pixel pattern can be suppressed substantially by a seven point horizontal line median filter, applied after dark correction and before flat fielding. The filtered image is shown in Fig. 26. While the 8 pixel pattern is almost eliminated, there is also some drawback. The filtering does to some extent smoothen the image and by this degrades the brighter parts (i. e., brighter before flat fielding) where the 8 pixel pattern had never been visible.

To suppress the 8 pixel pattern in dark parts without degrading brighter parts (here again, and below, we refer to the brightness before flat fielding), we

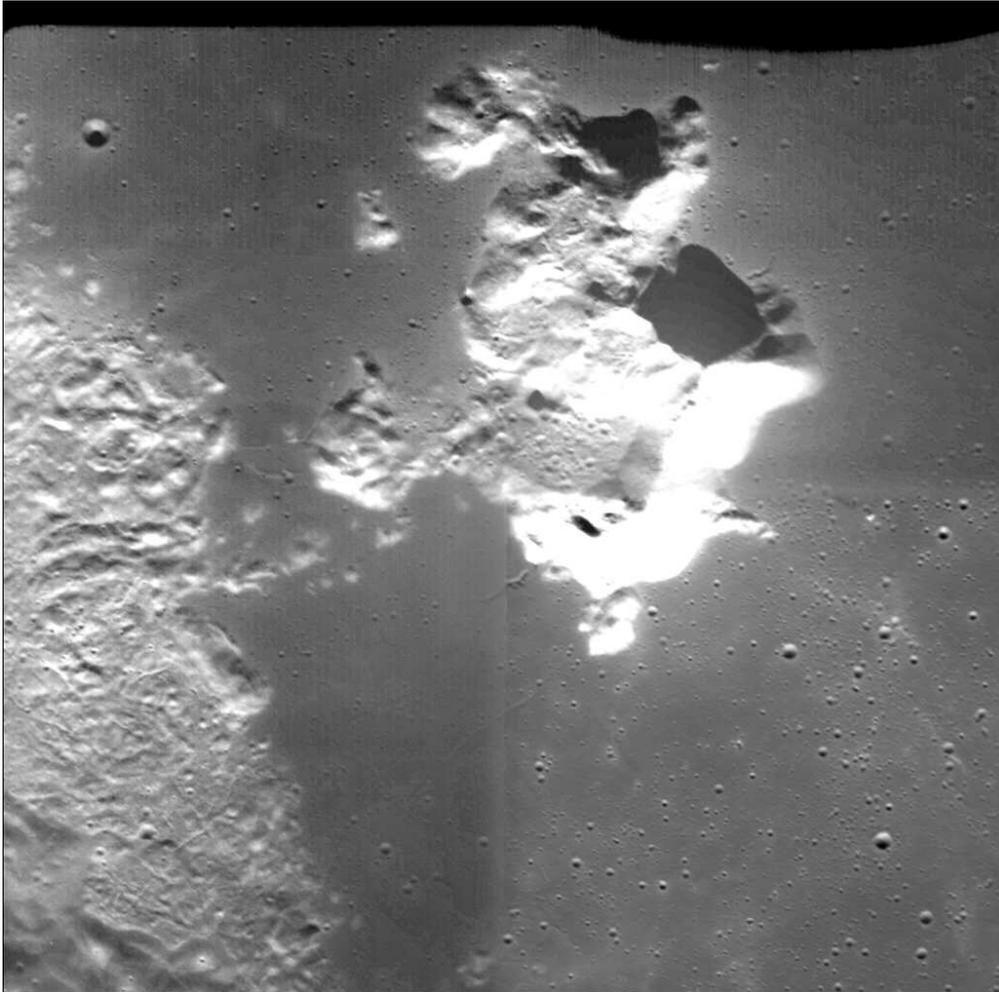


Figure 25: Calibrated image from orbit 2041, number 9. In the upper part, a spurious vertical pattern with 8 pixel spacing is visible.

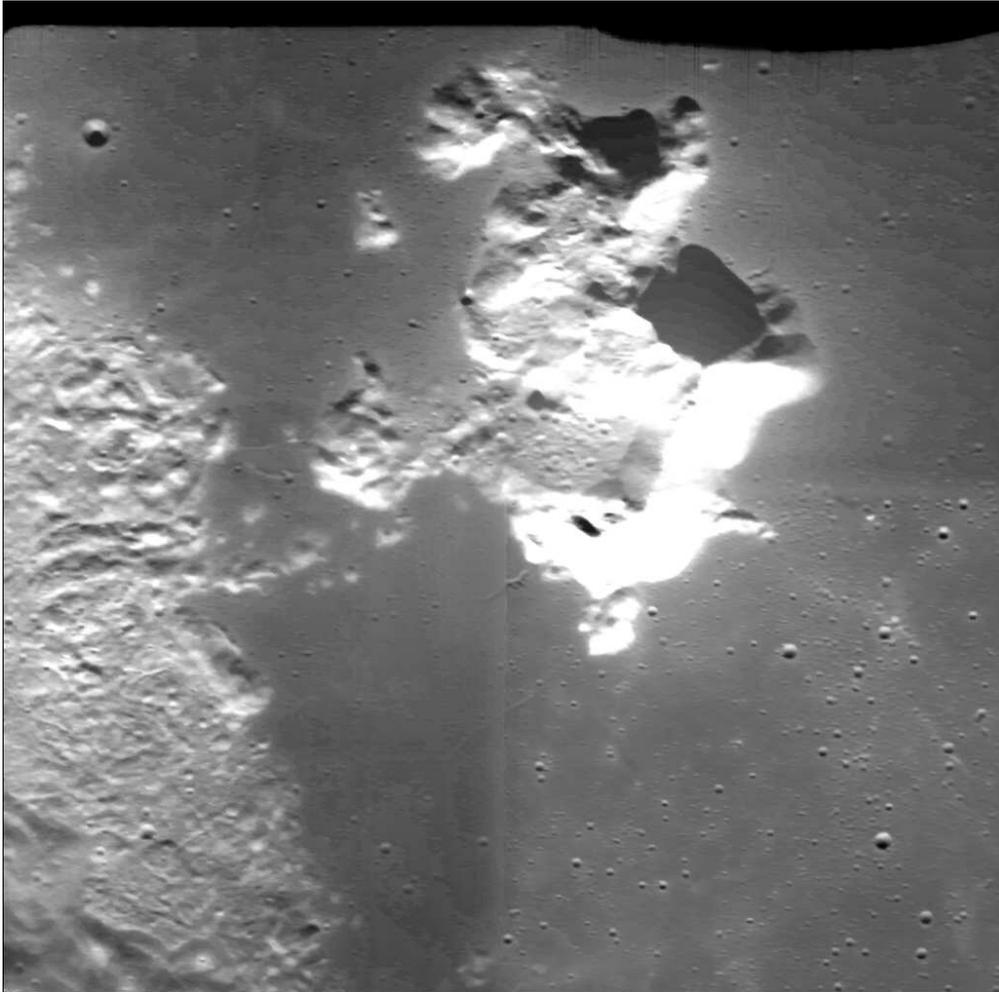


Figure 26: Filtered version of the image shown Fig. 25. A seven point horizontal line median filter has been applied homogeneously to the complete image.

can apply a weighted filtering. First, we filter the image as before with a seven point horizontal line median filter. But then, we take a weighted average of the filtered and the unfiltered image: Let D and D_f be the unfiltered and filtered data numbers, respectively. We set for each pixel

$$c = \exp\left(-\left(\frac{D_f}{64}\right)^2\right) \quad (16)$$

and take as resulting brightness

$$c \cdot D_f + (1 - c) \cdot D. \quad (17)$$

This implies that for small data numbers — i. e., small compared to 64 — we take the filtered image, while for large data numbers — i. e., large compared to 64 — we take the unfiltered image. The image resultant from this weighted filtering is shown in Fig. 27. While the 8 pixel pattern is suppressed similarly to the simple filtering, brighter parts (before flat fielding, e. g., the NONE filter area) are not degraded.

This weighted filtering procedure has to be applied after dark correction and before flat fielding. Thus one might start out from a raw image and apply dark correction, filtering, and flat fielding in this sequence. Alternatively, one might start out from a calibrated image, reverse the flat fielding by multiplying with the flat field times exposure time, then apply the filtering, and finally divide by the flat field times exposure time. The data number 64 which we used to separate dark pixels from bright pixels for the weighting may not be appropriate for all images.

5.4 Scattered light inside the camera

The indication of a substantial amount of scattered light inside the camera stems from a particular lunar limb image where the limb crosses the NONE filter area as well as a filter area. This raw limb image is shown in Fig. 28. It has been rotated by 180 degrees to put the sky up. Already in the raw image we note that the dark sky in the NONE filter area is considerably brighter than the dark sky in the filter area. This should not be the case unless some scattered light hits the CCD.

The dark corrected limb image is shown in Fig. 29. Now that the dark current contribution has been removed, the dark sky should really be dark. Histograms of the dark corrected image are shown in Fig. 30. The histogram of the NONE filter area exhibits a maximum at a data number of about 380. This represents the mean brightness of the lunar surface in the NONE filter area. The histogram of the NONE filter area exhibits a second peak at much lower data numbers, which represents the brightness of the dark sky.

To investigate this closer, the lower end of the histograms is shown expanded in Fig. 31. Now we can see that the dark sky has a data number of about 28 in the NONE filter area. There are three peaks for the filter areas. The middle one at a data number of about 27 represents the brightness of the lunar surface in the VIS filter area (the limb is passing through the VIS filter area). The peak

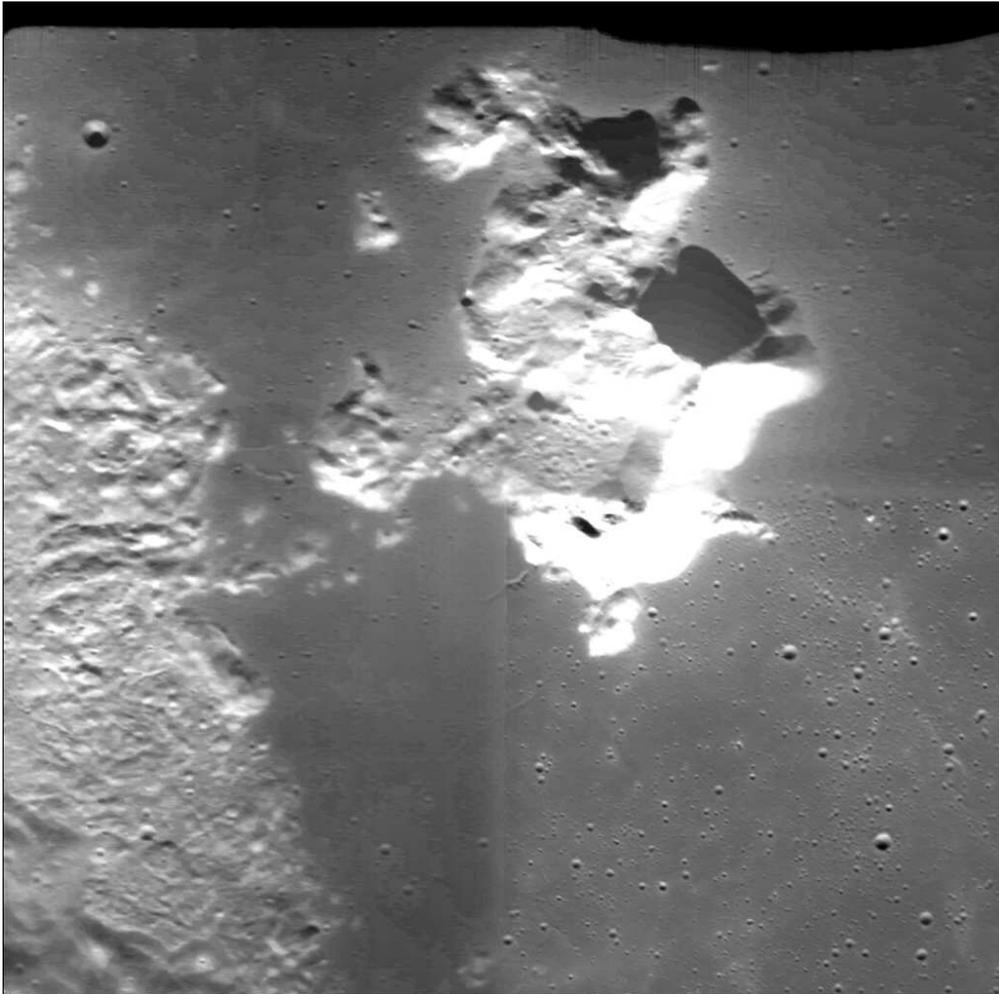


Figure 27: Weighted filtering applied to the image shown Fig. 25. The filter does only affect the darker parts (darker before flat fielding) of the image.

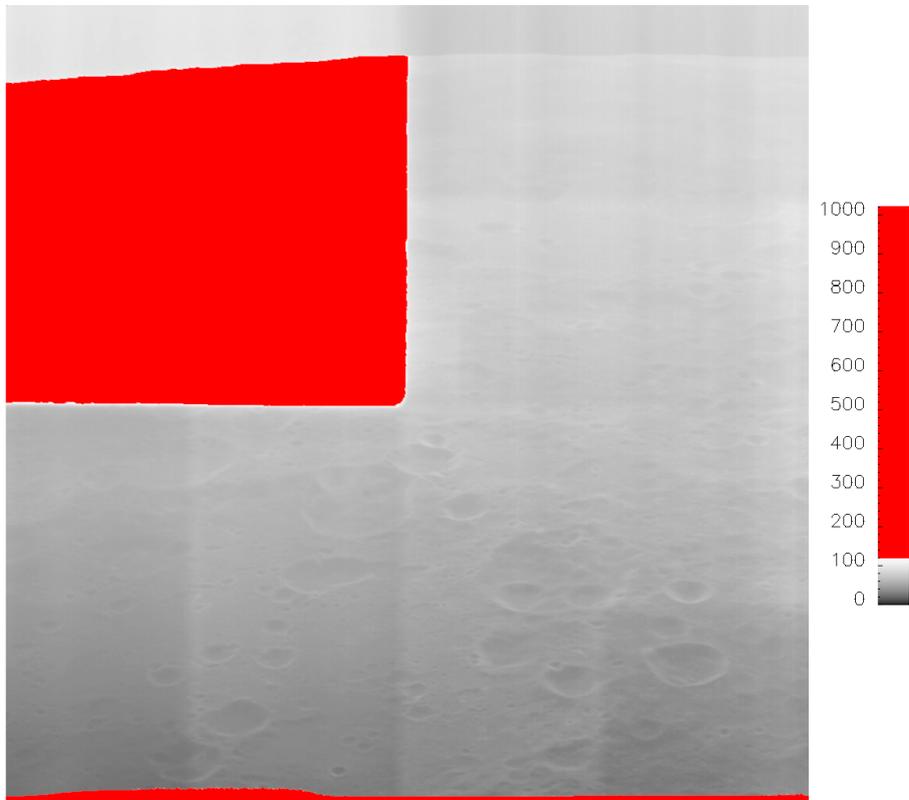


Figure 28: Raw limb image from orbit 2887, number 43.

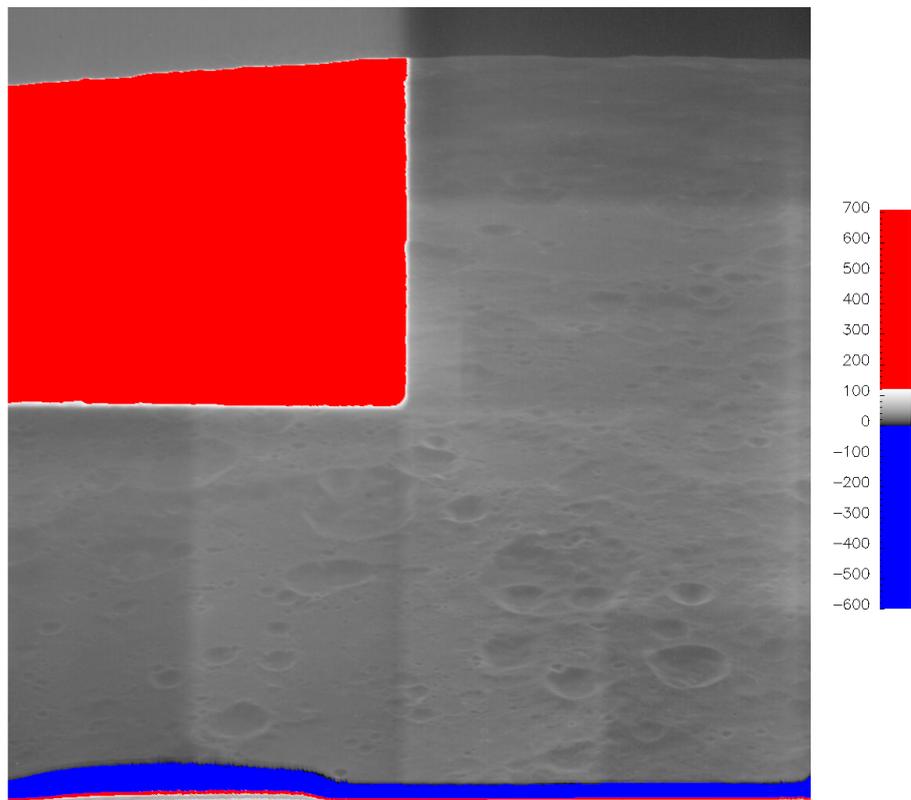


Figure 29: Dark corrected limb image in the same gray scale as the raw image shown in Fig. 28.

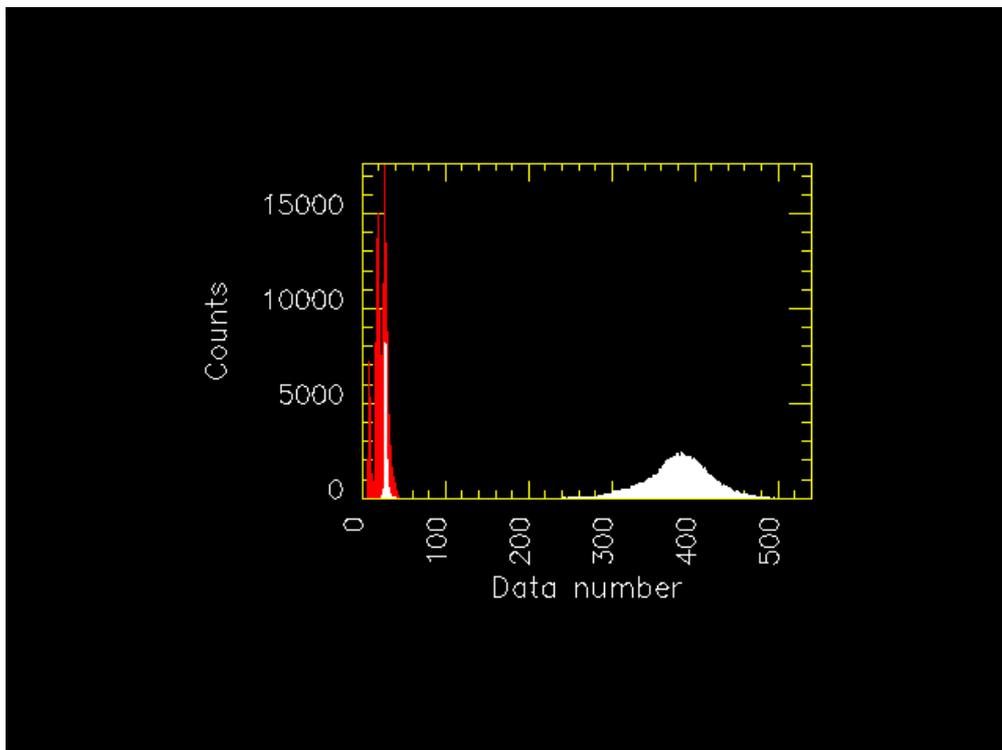


Figure 30: Histograms of the dark corrected image shown in Fig. 29. The white histogram represents the NONE filter area, the red histogram represents the filter areas.

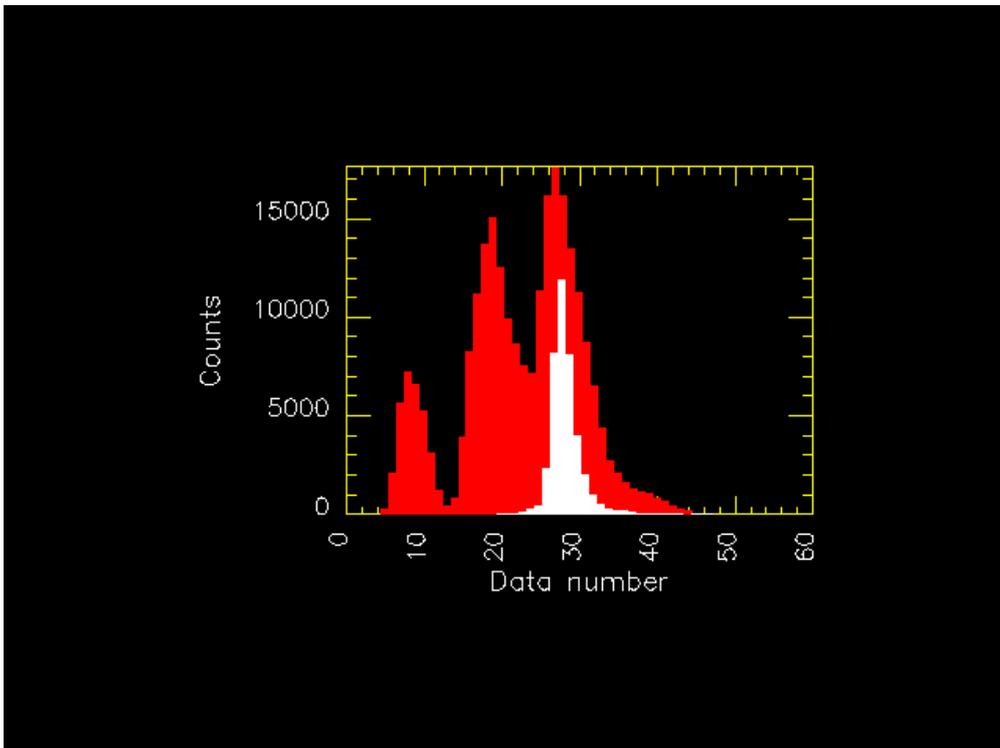


Figure 31: Expanded version of the histograms shown in Fig. 30. The white histogram represents the NONE filter area, the red histogram represents the filter areas.

at a data number of about 8 represents the brightness of the dark sky in the VIS filter area.

From these data numbers, we can draw some quantitative conclusions, which are summarized in this table:

	NONE	VIS	ratio
Lunar surface:	380	27	13.6
Dark sky:	28	8	3.5
Scattered light fraction:	7 %	30 %	
Direct light:	352	19	18.5
Box scattered light:	28	1.5	18.5
Filter scattered light:	0	6.5	

The lunar surface has a brightness of 380 in the NONE filter and 27 in the VIS filter, implying a ratio of 13.6 in the responseivity. The dark sky has a brightness of 28 in the NONE filter and 8 in the VIS filter, implying a much smaller ratio of 3.5. If we divide the dark sky brightness by the lunar surface brightness, we obtain the fraction of the scattered light, which is 7 % for the NONE filter and 30 % for the VIS filter. Particularly the latter number is quite high. The fact that the fraction is higher for the VIS filter than for the NONE filter means that scattered light gets below the filter from the side.

If we subtract the scattered light as given by the dark sky brightness from the lunar surface signal, we get the direct light which really comes from the lunar surface. It is 352 for the NONE filter and 19 for the VIS filter. The ratio of 18.5 is the true ratio of the responseivities. The value 13.6 was biased by the scattered light contribution.

We can separate the scattered light into two components. One that comes from the inside of the camera, from the “box”, and one that comes from the filter. The NONE filter part of the CCD should not see any (at least not any considerable) scattered light from the filter, therefore the dark sky brightness of 28 represents the box scattered light. With the responseivity ratio of 18.5, we get for the box scattered light after passing the filter a value of 1.5. This means that only 1.5 of the 8 data numbers of scattered light which we see below the filter come from the inside of the camera box, while 6.5 data numbers come from the filter itself.

The estimated fraction of scattered light — 7 % for the NONE filter and 30 % for the VIS filter — and particularly the spatial pattern of scattered light varies substantially from image to image depending on the brightness distribution in the scene. In Fig. 32, we show an image where there is excessive scattered light in the region of the filter areas close to the NONE filter area. This is a calibrated image. If this scattered light was present similarly in all images, it would have been taken up by the flat field and removed by the flat fielding. Therefore, this brightened region indicates above the average scattered light. And if this is above the average, there must also be cases below the average. An example of that is shown in Fig. 33. In this image, the region of the filter areas close to the NONE filter area is darkened, indicating below the average scattered light.

Because of this variability, there is not an easy way to correct for the scattered light. We would need to carry out a covariance analysis between the one

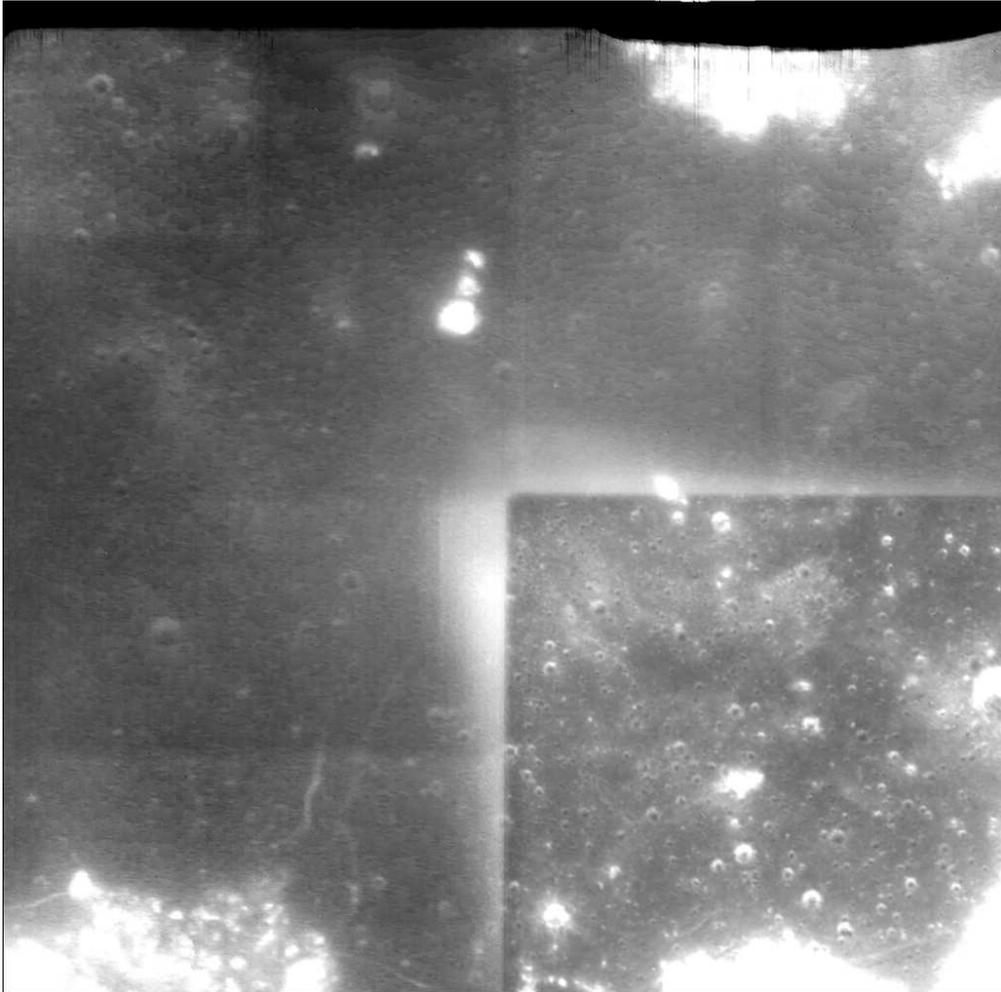


Figure 32: Image from orbit 2437, number 5, as example for above the average scattered light adjacent to the NONE filter area.

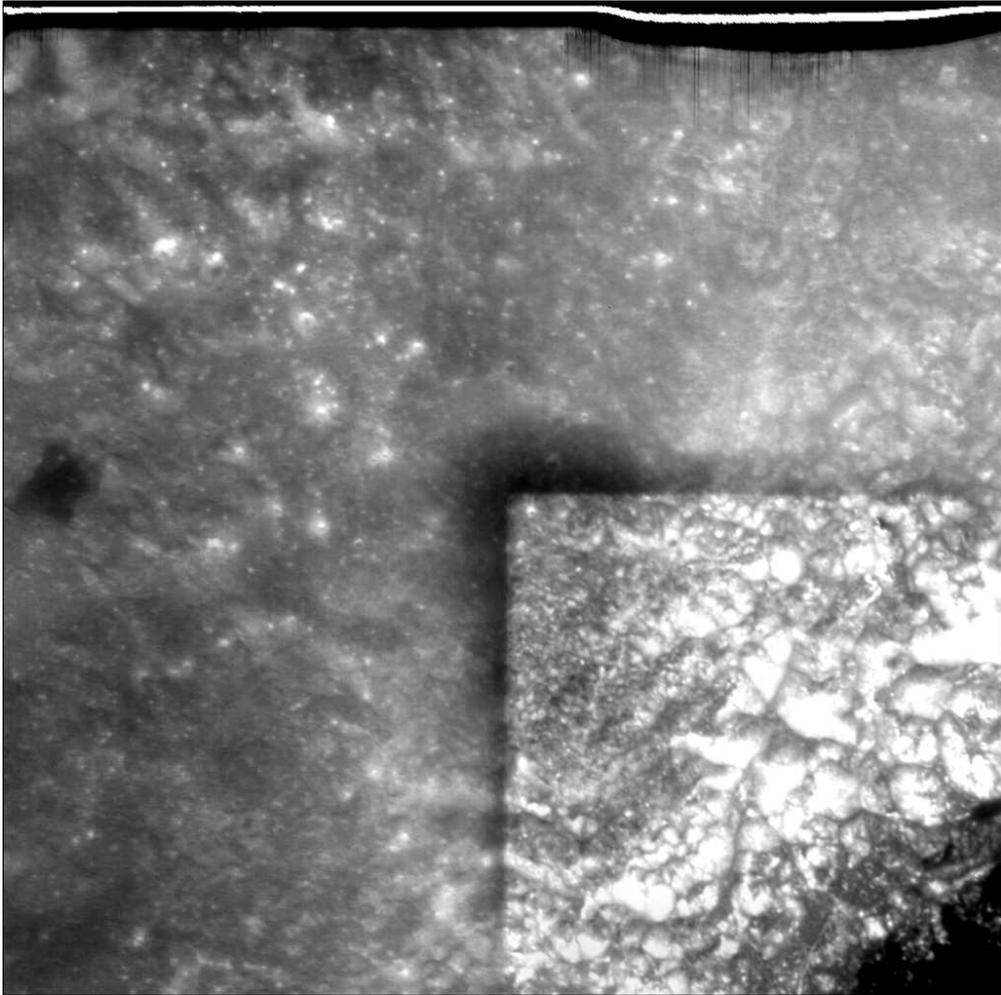


Figure 33: Image from orbit 589, number 25, as example for below the average scattered light adjacent to the NONE filter area.

million pixels. The size of the covariance matrix would be

$$1 \text{ million} \times 1 \text{ million} \times 4 \text{ Bytes} = 4 \text{ Terabytes.} \quad (18)$$

Such working memory may be beyond most contemporary computers. We may think of reducing the resolution of the images for the covariance analysis as the pattern of scattered light may mostly be smooth. However, considerable gradients may occur at the filter boundaries. A solution could be to use a variable resolution, which is higher near the filter boundaries. Even if a covariance analysis with reduced resolution could be handled, it requires considerable effort, and it is not clear if the scattered light can really be removed.

5.5 Absolute calibration

No absolute calibration has been performed so far. The unit of a pixel value of a calibrated image can not be given.

An absolute calibration may be achieved either based on laboratory measurements of on in flight data. Concerning laboratory measurements, the table of measurement results provided as the file ‘SPECTRAL_RESPONSE.PDF’ in the ‘DOCUMENT’ directory of the AMIE dataset may give a first hint to the spectral responseivity. In flight, bright stars have been observed. These observations may provide a basis for an absolute calibration.

5.6 Distortion

The distortion of AMIE images has not been investigated. It seems to be very small. Laboratory measurements and star field observations were performed to asses the distortion.