

**2001 Mars Odyssey
Thermal Emission Imaging System
(THEMIS)**

Data Processing User's Guide

Version 0.8
July 1, 2003

Prepared by:

Philip R. Christensen
THEMIS Principal Investigator
Arizona State University

Date

Approved:

Philip R. Christensen
THEMIS Principal Investigator

Date

Gaylen McSmith
Mars Odyssey Science Office Manager

Date

TBD
PDS / Mars Odyssey Interface Manager

Date

1 OVERVIEW

1.1 INSTRUMENT OVERVIEW

The Thermal Emission Imaging System (THEMIS) instrument consists of separate infrared and visible focal plane assemblies that share the telescope through a beamsplitter but have independent power and data interfaces to the spacecraft to provide system redundancy. The instrument was developed by Arizona State University (ASU) with Santa Barbara Remote Sensing (SBRS). The infrared focal plane and electronics were provided by the Santa Barbara Research Center (SBRS). The visible imager focal plane and electronics were provided by Malin Space Science Systems (MSSS). The telescope is a three-mirror anastigmat telescope with a 12-cm effective aperture and a speed of f/1.6. The IR and visible imagers have independent power and data interfaces to the spacecraft. The IR focal plane is an uncooled microbolometer with 320 cross-track pixels and 240 down-track pixels with an instantaneous field of view (IFOV) of 0.25 milli-radians (100 m from 400 km altitude). The focal plane is temperature stabilized to ~ 0.001 K using a temperature controller driving a thermo-electric heater/cooler. Ten stripe filters are placed over the microbolometer array to produce ten $\sim 1\text{-}\mu\text{m}$ wide spectral bands at nine separate wavelengths centered from 6.8 to 14.9 μm . Two filters (Bands 1 and 2) cover the same spectral region centered at 6.8 μm . The nine THEMIS IR wavelengths include eight surface-sensing wavelengths (Bands 1-9) and one atmospheric wavelength (Band 10). Time delay integration (TDI) of 16 consecutive rows is performed by firmware in the THEMIS IR imager. The THEMIS visible imager contains 1,024 cross-track pixels with an IFOV of 0.045 milli-radians (18 m from 400 km altitude) covering a 18.4 km swath bore-sighted with the IR imager. The visible imager has five stripe filters, each covering 192 rows on the detector. These 192-line framelets are collected every 1 second, producing 26 rows of downtrack overlap in successive framelets at an orbital velocity of 3.0 km/sec. An internal calibration flag, the only moving part in the instrument, provides thermal calibration and IR flat-fielding, and is used to protect the detectors from unintentional direct illumination from the Sun. The instrument weighs 11.2 kg, is 29 cm by 37 cm by 55 cm in size, and consumes an orbital-average power of 14 W.

THEMIS images are calibrated in flight using periodic views of the internal calibration flag. The calibration shutter flag is stored against a sidewall that maintains the flag at a known temperature. The flag is closed, imaged, and reopened at selectable intervals throughout each orbit. This process produces gores in the surface observations lasting approximately 50 seconds for each calibration. Calibration data are acquired every $\sim 3\text{-}10$ minutes as necessary. The optimum spacing of these observations is chosen to meet the calibration accuracy requirements, while minimizing the loss of surface observations.

The THEMIS IR imager internal analog-to-digital converter (ADC) produces a 12-bit signal (DN values of 0-4095). The digital output signal from the ADC is digitally adjusted by gain and offset terms and the result is converted to an 8-bit output value that is returned to Earth. The gain term determines how the dynamic range of the 12-bit

signal from the ADC is converted to the 8-bit output. At a gain of 1, the full 4096 DN input range is used and is converted to an 8-bit number by dividing by 16. At a gain of two the input dynamic range is 2048, etc, up to a gain of 16 that uses a dynamic range of 256. Thus, a gain of 1 uses the full input dynamic range, but with a factor of 16 loss in precision; a gain of 16 preserves the full precision from the ADC with a factor of 16 reduction in dynamic range. Gains higher than 16 are not appropriate as they result in a lower input dynamic range than is available in the 8-bit output. The position of the selected dynamic range relative to the full 4096 input values is determined by the offset term. The offset term is measured relative to the mid-point (2048) of the input signal. The input to output conversion is given by:

$$\text{output} = (\text{gain} / 16) * \text{input} - (256 * (\text{offset} + 8))$$

The analog gain of the THEMIS IR imager was adjusted during assembly to provide a 1-sigma variance of approximately 0.5 DN at the maximum useful digital gain of 16.

The data returned from the instrument in flight are converted to scene spectral radiance ($\text{W cm}^{-2} \text{str}^{-1} \mu\text{m}^{-1}$) by: (1) adjusting for the gain and offset that were applied in the instrument to optimize the dynamic range and signal resolution for each scene; (2) correcting for any signal drift or offset that occur between observations of the calibration flag; and (3) converting signal to radiance using the measured signal and temperature of the calibration flag, together with the instrument response function determined prior to launch.

1.2 PRE-LAUNCH CALIBRATION OVERVIEW

The calibration, test, and characterization of the THEMIS was performed at SBRS prior to instrument installation on the Mars 2001 Odyssey spacecraft at Lockheed Martin in Denver. The THEMIS calibration was performed by personnel from SBRS and ASU.

The primary objectives of these tests were to determine:

- 1) the geometric performance of the infrared and visible imagers, including the point-spread function, out-of-field response, camera model, and the focus performance over temperature;
- 2) the relative spectral response of the infrared and visible imagers; and
- 3) the radiometric performance of the infrared and visible imagers, including the instrument response function, the absolute radiometric calibration, the radiometric precision (noise equivalent spectral radiance), the performance metrics, and the analysis of systematic errors.

In addition to these calibrations, an extensive set of tests were performed under ambient and vacuum conditions to verify the instrument functional performance, including the

command and signal processors, the command and data links, and the calibration flag actuator.

2 INFRARED IMAGER CALIBRATION

2.1 INFRARED IMAGER CALIBRATION OVERVIEW

The THEMIS infrared imager is calibrated in flight using periodic observations of an internal calibration flag mounted in the THEMIS optical path immediately in front of the IR/visible beamsplitter. Observations of this flag are used to automatically flat field the IR detectors prior to data acquisition, and to radiometrically calibrate the THEMIS data at the end of each data acquisition.

Figure 2-1 shows the Mars data calibration flow for the IR imager. The basic process steps are outlined below; the details of the measurements and/or processes are discussed in subsequent sections.

- 1) Acquire raw Mars image and flag-closing image.
- 2) Adjust Mars and flag-closing images for applied gain and offset – convert to signal level (Gain=1; offset=0).
Integer values of gain and offset (DN at a gain of 1) are available. In flight only gains of 1, 2, 4, 8, or 16 will be used. The DN offset will be adjusted to prevent saturation at either the high or low end depending on scene temperature.
- 3) Remove image offset due to focal plane thermal equilibration using the measured calibration flag signal at end of image sequence.

The instrument firmware automatically sets the signal level when viewing the internal calibration flag to 128 DN for a gain of 1, offset of 0. During observations of external targets the THEMIS focal plane and internal temperatures drift slightly as these components attempt to equilibrate with the radiance of the observed scene (**Figure 2-2**). In flight this image offset component is removed by observing the calibration flag at the end of each observing sequence and finding the maximum or minimum calibration flag signal value as appropriate.

- 4) Compute the delta signal (scene minus calibration flag)
- 5) Correct for residual image signal drift.

The time constant of the thermal equilibration is approximately 30 seconds under typical observing conditions. Thus, the majority of the signal drift occurs within the first two minutes after the calibration flag is opened, and the effect of signal drift is minimized by waiting two minutes following a calibration before acquiring Mars images. Any minor signal drift that occurs during the image sequence can be removed if necessary.

- 6) Compute the delta spectral radiance using the instrument response function.
- 7) Compute the calibration flag spectral radiance using the measured calibration flag temperature, the calibration flag emissivity, and the spectral-radiance-versus-temperature function determined pre-launch.
- 8) Compute the initial scene spectral radiance using the calibration flag spectral radiance and the delta spectral radiance.
- 9) Remove systematic noise.

- a) Remove column- and row-correlated systematic noise.

Column-correlated noise occurs due to minor time-variable changes in detector response relative to the pre-launch response function. Row-correlated noise occurs due to minor fluctuations in the detector read-out bias voltage.

- b) Remove time-domain signal variations.

Minor (~ 0.001 °C) fluctuations in the focal-plane temperature over in time scales of ~ 10 - 60 seconds produce minor (~ 1 DN; gain – 16) signal variations in the time domain. These signal variations are constant across a row in each band and are correlated in time between bands. Because each band is viewing a different point on Mars at a given time, these correlations in time are offset in space.

- 10) Remove optical stray light. The calibrated spectral radiance from this step is archived in the THEMIS IR Reduced Data Record (RDR).
- 11) Compute brightness temperature in Band 9 and archive in the THEMIS IR Brightness Temperature Record (BTR).

2.2 INFRARED IMAGER CALIBRATION ALGORITHM

2.2.1 Acquire raw image.

THEMIS IR images are acquired using a single global gain and offset that applies to all bands over the entire image. A calibration reset is performed prior to Mars data acquisition. The calibration reset consists of closing the calibration flag, observing the flag, and setting the DN of each pixel to 128 at a gain of 1, offset of 0 (**Figure 2-2**). This reset produces a uniform flat field for all detectors. The THEMIS flight software requires that an image be acquired when a calibration reset is performed. This image is referred to as the Reset Image. Closing the calibration flag introduces a thermal perturbation to the focal plane assembly, and it is observed at SBRS and in flight that the instrument experiences a temperature reequilibration following this perturbation (**Figure 2-2**; see the THEMIS Calibration Report for details). The thermal reequilibration after the calibration

flag is opened produces a drift in the output signal that stabilizes with a time constant of approximately 30 seconds. Typically, Mars imaging does not begin within 2 minutes of the Reset Image to allow the signal drift to stabilize. The Reset Image occurs during the period of highest signal drift and is the least well calibrated. The Reset Image is typically limited to the minimum length possible (256 lines; 1 band), except for images acquired specifically to monitor and assess the degree of signal drift following a calibration reset.

2.2.2 Adjust Mars and flag-closing images for applied gain and offset

All THEMIS IR image data are converted to a floating point “signal” value, corresponding to a gain of 1, offset of 0 prior to radiance calibration.

2.2.3 Remove image offset due to focal plane thermal equilibration

A flag-closing image is acquired at the end of each THEMIS IR observing sequence. The davinci function “get_flag_dn” (author: P. Christensen) finds the DN value in this image that best corresponds to the flag DN in each band. All 320 columns are averaged to improve the signal-to-noise for this determination. The average of Band 3 lines 1 through 5 and 20 through 25 in flag-closing image are compared to determine if the scene was warmer than the flag (flag-closing DN decreasing with time) or colder (flag-closing DN increasing with time). If scene was warmer, the algorithm finds the minimum in the flag-closing image. If scene was colder, the algorithm uses the maximum in the flag-closing image. The function returns the calibration flag signal value (Gain = 1; offset = 0).

The maximum (or minimum) flag-closing DN cannot be observed in Bands 6-10 because of the time delay built into the THEMIS firmware to account for the fact that each successive band observes the same point on Mars at a later time. The firmware acts to align the images from each band in space, resulting in an offset in time. As a result, the flag has closed and the signal has begun to reequilibrate before data from Bands 6-10 are collected.

flag_option 1. This option returns the measured maximum (or minimum) DN values for Bands 1-5. The average DN of Bands 1-5 is used for the flag DN for Bands 6-10.

flag_option 2: Same as option 1 for Bands 1-5. The flag DN for Bands 6-10 is calculated using the DN offset between Band 5 and Bands 6-10 for the lines in Bands 6-10 that correspond to the identical time in Band 5. The DN offsets at this time are added to the flag DN for Band 5 to determine the flag DN in Bands 6-10. This option also filters the flag-closing image (flag_filter_option = 1; filter size = 3) in time before finding maximum (or minimum) value.

The “flag_option” and “flag_filter_option” used in each version of the calibration are given in Section 2.3.

2.2.4 Compute the delta signal (scene minus calibration flag)

The delta signal is computed by subtracting the calibration flag signal value from the Mars image signal for each band.

2.2.5 Correct for residual image signal drift

The davinci function “drift” (author: P. Christensen) removes residual signal image drift. This function includes a “no drift” option in which no correction is applied. The “no-drift” option is the default.

2.2.6 Compute the delta spectral radiance using the instrument response function

The pre-launch instrument response function is used to convert delta signal to delta spectral radiance. The SBRS pre-launch data were processed using the davinci function “make_response” (author: P. Christensen) and fit using a linear least-squares fit using davinci function “fit_response_full” (author: P. Christensen). The slope and offset terms for this linear fit are computed and stored for each of the 320 cross-track “pixels” in each band. These pixels correspond to the average of the 16 downtrack pixels that are merged using time-delay-integration (TDI) in the THEMIS firmware. When TDI is not used, the signals from a single row in each filter are used. A separate response function was not determined for these rows; if TDI is not used at Mars, the TDI-averaged response function will still be used and the resulting calculated radiance will be slightly (< 5%) in error.

The versions of “make_response” and the versions of the instrument response data files derived from “make_response” are described in Section 2.3. The versions of the davinci function “fit_response_full” and the versions of the instrument response function coefficient files derived from “fit_response_full” are described in Section 2.3. Section 2.3 also lists the instrument response function used in each version of the instrument calibration.

2.2.7 Compute the calibration flag spectral radiance.

The calibration flag spectral radiance is computed using the derived temperature of the flag and the radiance-to-temperature look-up table calculated from the relative spectral response of each THEMIS band. The temperature of the flag is determined from the average of the two values of the “flag_temp” telemetry point returned at the beginning of the flag-closing image. If this telemetry point is not available, then the values from the end of the flag-closing image are used. If these values are also not available, then the calibration_quality_flag (TBD) for all images that use this flag-closing image are set to “TBD (not available)”.

The radiance-to-temperature look-up table is computed using davinci function “make_temp_rad” (author: P. Christensen). The versions of the file produced by

“make_temp_rad” and the file used in each version of the calibration are given in Section 2.3.

2.2.8 Compute the initial spectral radiance using the calibration flag spectral radiance and the delta spectral radiance.

The calibrated spectral radiance for each pixel in the image is computed by adding the calibration flag spectral radiance to the delta spectral radiance.

2.2.9 Remove systematic noise

a) Remove column- and row-correlated systematic noise.

The davinci function “destripe” (author: P. Christensen) is used to remove column- and row-correlated noise. Column-correlated noise is due to minor variations in the instrument response in one or more of the detectors that are combined in each pixel by time-delay-integration. If these variations are constant throughout an image, then an entire column has a slightly high or low signal relative to the neighboring columns along the entire length of the image. This correlation along the image is used to remove this noise effect. This effect produces column-to-column variations of up to approximately 0.25 DN at a gain of 16.

Row-correlated noise is produced by fluctuations in the detector read-out bias voltage on time scales that are short compared with the read-out time for a row in the focal plane array, but long compared with the time between reading out a given row in the following frame (30 Hz). Thus, all pixels in a row have a similar bias, but the following row can have a slightly different bias. This variation produces average line-to-line variations of up to ~ 0.5 DN at a gain of 16.

Row-correlated noise is removed by first transposing the image in x and y, and then applying the identical algorithm described below. The resulting image and the computed difference vectors are transposed back within the “destripe” algorithm.

The first step in the algorithm takes the image average in the y direction for each band. The resulting average line vector is filtered using a unity weight filter of user-specified length (parameters “destripe_filter_x; destripe_filter_y”). The default filter size is 9 samples in both x and y. The average vector and filtered average vector are differenced to create a difference vector.

Option 1: Subtract the difference vector from each line in the image to create a new calibrated radiance with noise removed. The line and column difference vectors (diff_line; diff_column) are returned from the “destripe” function and stored in the image history structure.

Option 2: Same as Option 1 except the values of the difference vector less than a threshold value (parameter `thresh_size`) are set to 0. prior to subtracting the difference vector from image.

Option 3: Same as Option 1 except: Where the value of the difference vector is greater than a threshold value (`thresh_size`), the value of the average vector at this location is set to the average of nearest neighbors that are not themselves above the threshold values. This modified average vector is then filtered with a unity-weight filter of size “`destripe_filter_x` or `destripe_filter_y`”. The difference between the original average vector and the filtered modified average vector is used to create a new difference vector (`diff_line`; `diff_column`) that is subtracted from each line or column in the image. The effect of removing the “spikes” before filtering is to reduce the erroneous propagation of these spikes into neighboring values when the filtering is applied.

The “`destripe`” parameters (`destripe_option_x`; `destripe_option_y`; `destripe_filter_x`; `destripe_filter_y`) used in each version of the calibration are given in Section 2.3.

b) Remove focal-plane-correlated signal oscillations.

TBD

2.2.10 Remove optical stray light.

The THEMIS beamsplitter bracket has a surface that is out of the image ray bundle, but can produce a reflection of stray light that can reach the focal plane in IR Bands 3 through 8 and produce a “ghost” image. This ghost image is defocused slightly due to its longer path length through the optics, and moves through the focal plane in the opposite direction as the primary image. This opposite motion acts to blur the image in the downtrack direction when TDI is applied. The ghost image is removed in the standard processing in a three-step process using `davinci` function “`fix_ghost`” (author P. Christensen). The first step is to compute the ghost image. This is done by convolving the primary image with a two-dimensional, unity-weight “`defocus_filter`” that simulates the defocus blur of the ghost image, and a one-dimensional, variable-weight along-track “`tdi_smear_filter`” that simulates the blurring effect of TDI (`davinci` function “`convolve_ghost`”; author P. Christensen). This function has two parameters; “`defocus_filter`” that sets the size of the defocus filter, and “`tdi_smear_filter`” that sets the size and characteristics of the TDI smear filter. The second step is to offset the ghost image to simulate the downtrack optical offset and the resulting crosstrack motion due to planetary rotation (`davinci` function “`geom_ghost`”; authors N. Gorelick and M. Weiss-Malik). This function has a single parameter “`yoffset`” that determines the offset of the ghost image in each band in the time dimension. The third step is to subtract the computed ghost image from the radiance image (`davinci` function “`subtract_ghost2`”; author P. Christensen). This function uses removes the appropriate amount (“`percent`”) of the primary image that makes up the ghost image. This value is related to the reflectivity of the beamsplitter bracket surface in each band. The computed ghost image is further

offset in the cross-track direction by “xdelta” to account for the fact that the beamsplitter bracket is not exactly aligned with the focal plane. The values of “defocus_filter”, “tdi_smear_filter”, “yoffset”, “percent”, and “xdelta” used in the calibration are given in Section 2.3.

This step produces calibrated spectral radiance in each THEMIS IR band.

2.2.11 Compute Brightness Temperature.

The brightness temperature is computed from the calibrated spectral radiance for Band 9 assuming a surface emissivity of 1.0 and an atmospheric opacity of 0.0. This brightness temperature is determined using the radiance-to-temperature look-up table with davinci function “make_temp_rad” (author: P. Christensen). The versions of the file produced by “make_temp_rad” and the file used in each version of the calibration are given in Section 2.3.

2.3 IR CALIBRATION ALGORITHM AND PARAMETER HISTORY

ir_image_cal_qube

| version | date | Description |
|------------------------|---------|--|
| ir_image_cal_qube_v4.6 | 9/6/02 | inst. response = irf_fit_all_v3.0_tv6_1_2_3.0 temp_rad = temp_rad_v4 calib. script = cal_image_v1 flag_option = 2 flag_filter_option = 1 destripe_option_x = 3 destripe_option_y = 3 filt_size_x = 9 filt_size_y = 9 drift_option = 0 radiance_offset = radiance_offset_v1 defocus_option = 13 tdi_smear_option = 1 yoffset = 0, 0, 349, 299, 249, 202, 152, 103, 0, 0, xdel_option = 8 ydel = 0. percent_option = 5 |
| ir_image_cal_qube_v4.5 | 6/27/02 | inst. response = irf_fit_all_v3.0_tv6_1_2_3.0 temp_rad = temp_rad_v4 calib. script = cal_image_v1 flag_option = 2 flag_filter_option = 1 destripe_option_x = 3 destripe_option_y = 3 filt_size_x = 9 filt_size_y = 9 drift_option = 0 radiance_offset = radiance_offset_v1 |

| | | |
|------------------------|---------|--|
| ir_image_cal_qube_v4.4 | 5/25/02 | inst. response = irf_fit_all_v3.0_tv6_1_2_3.0 temp_rad = temp_rad_v4 calib. script = cal_image_v1 flag_option = 2 flag_filter_option = 1 destripe_option_x = 3 destripe_option_y = 3 filt_size_x = 9 filt_size_y = 9 drift_option = 0 deghost_option = 0 |
| ir_image_cal_qube_v4.3 | 5/21/02 | inst. response = irf_fit_all_v3.0_tv6_1_2_3.0 temp_rad = temp_rad_v4 calib. script = cal_image_v1 flag_option = 2 flag_filter_option = 1 destripe_option_x = 1 destripe_option_y = 1 filt_size_x = 9 filt_size_y = 9 drift_option = 0 deghost_option = 0 |
| ir_image_cal_qube_v4.2 | 5/20/02 | inst. response = irf_fit_all_v2.0_tv6_1_2_3.0 temp_rad = temp_rad_v4 calib. script = cal_image_v1 flag_option = 2 flag_filter_option = 1 destripe_option_x = 1 destripe_option_y = 1 filt_size_x = 9 filt_size_y = 9 drift_option = 0 deghost_option = 0 |
| ir_image_cal_qube_v4.1 | 5/19/02 | inst. response = irf_fit_all_v2.0_tv6_1_1_3.0 temp_rad = temp_rad_v4 calib. script = cal_image_v1 flag_option = 2 flag_filter_option = 1 destripe_option_x = 1 destripe_option_y = 1 filt_size_x = 9 filt_size_y = 9 drift_option = 0 |

| | | |
|----------------------|---------|---|
| ir_image_cal_qube_v3 | 5/18/02 | Same as unversioned ir_image_cal_qube. Gave identical results as unversioned version. inst. response = irf_fit_all_v2.0_tv6_1_1_3.0 temp_rad = temp_rad_v4 cal_image_v1 flag_option = 1 flag_filter_option = 0 destripe_option_x = 0 destripe_option_y = 0 drift_option = 0 |
| ir_image_cal_qube | 2/02 | Used from start of mapping to ~5/18/02 when version '_v3' was implemented |

bold = change from previous version

| | Script | Version | Date | Description |
|-------------|---------------|----------------|-------------|--|
| cal_image | cal_image_v1 | 5/17/02 | | Identical to unversioned version |
| | cal_image | 10/2/01 | | Used from start of mapping to _ when ‘_v1’ was implemented |
| get_flag_dn | | 5/15/02 | | Flag_option 2: Same as option 1 for Bands 1-5. The flag DN for Bands 6-10 is calculated using the DN offset between Band 5 and Bands 6-10 determined using the lines in Bands 6-10 that correspond to the identical time in Band 5. The DN offsets at this time are added to the flag DN for Band 5 to determine the flag DN in Bands 6-10. This option also filters the flag-closing image (flag_filter_option = 1; filter size = 3) in time before finding maximum (or minimum).. |
| | | 6/00 | | Flag_Option 1: Use two early points in flag-closing image to determine if scene was warmer or colder than flag. If scene was warmer, find minimum in flag-closing image. If scene was colder, find maximum in flag-closing image. The maximum (or minimum) DN cannot be observed in Bands 6-10 because of the time delay built into the THEMIS firmware to account for the fact that each successive band observes the same point on Mars at a later time. The firmware acts to align the images from each band in space, resulting in an offset in time. As a result, the flag has closed and the signal has begun to reequilibrate before data from Bands 6-10 are collected. Flag_option 1 returns the measured maximum (or minimum) DN values for Bands 1-5. The average DN of Bands 1-5 is used to approximate the flag DN for Bands 6-10. Flag_filter_option = 1 Filter flag-closing image by 3 samples in time prior to determining flag DN. |

| | | | |
|---------------|-----|----------|---|
| destripe | | 5/25/02 | Option 3: Same as Option 1 except: Set difference vector values greater than threshold value (thresh_size) to average of nearest neighbors that are not themselves above the threshold values. Filter this modified vector with boxcar filter of size "destripe_filter_x; destripe_filter_y". Take difference between modified vector and filtered modified vector to create a difference vector (diff_line; diff_column). Subtract this difference vector from each line or (column) in image. The effect of removing the "spikes" before filtering is to reduce the erroneous propagation of the spikes into neighboring values when the filtering is applied. Create and remove row vector (column noise) first, then create and remove column vector (row noise). Row and column vectors corresponding to the noise removed are stored in RDR header and can be reapplied (column vector first, then row vector). |
| | | 11/8/01 | Option 2: Same as Option 1 except: Set difference vector values less than threshold value (thresh_size) to 0., then difference from image. |
| | | 11/8/01 | Option 1: Take average in direction 1 (x or y) to create a vector in direction 2 (y or x). Filter this vector with boxcar filter of size "filter_size". Take difference between vector and filtered vector to create a difference vector (diff_used). Subtract this difference vector from each line or (column) in image. |
| make_response | 3.0 | 3/29/02 | used temp_rad_v4 (spectral radiance) |
| | 2.0 | 1/4/02 | used temp_rad_v2 (integrated radiance) |
| | 1.1 | 11/12/01 | used temp_rad_v1.1 (0.5 C intervals; integrated radiance) |
| | 1.0 | 7/00 | used temp_rad_1 (0.1 C intervals; integrated radiance) |

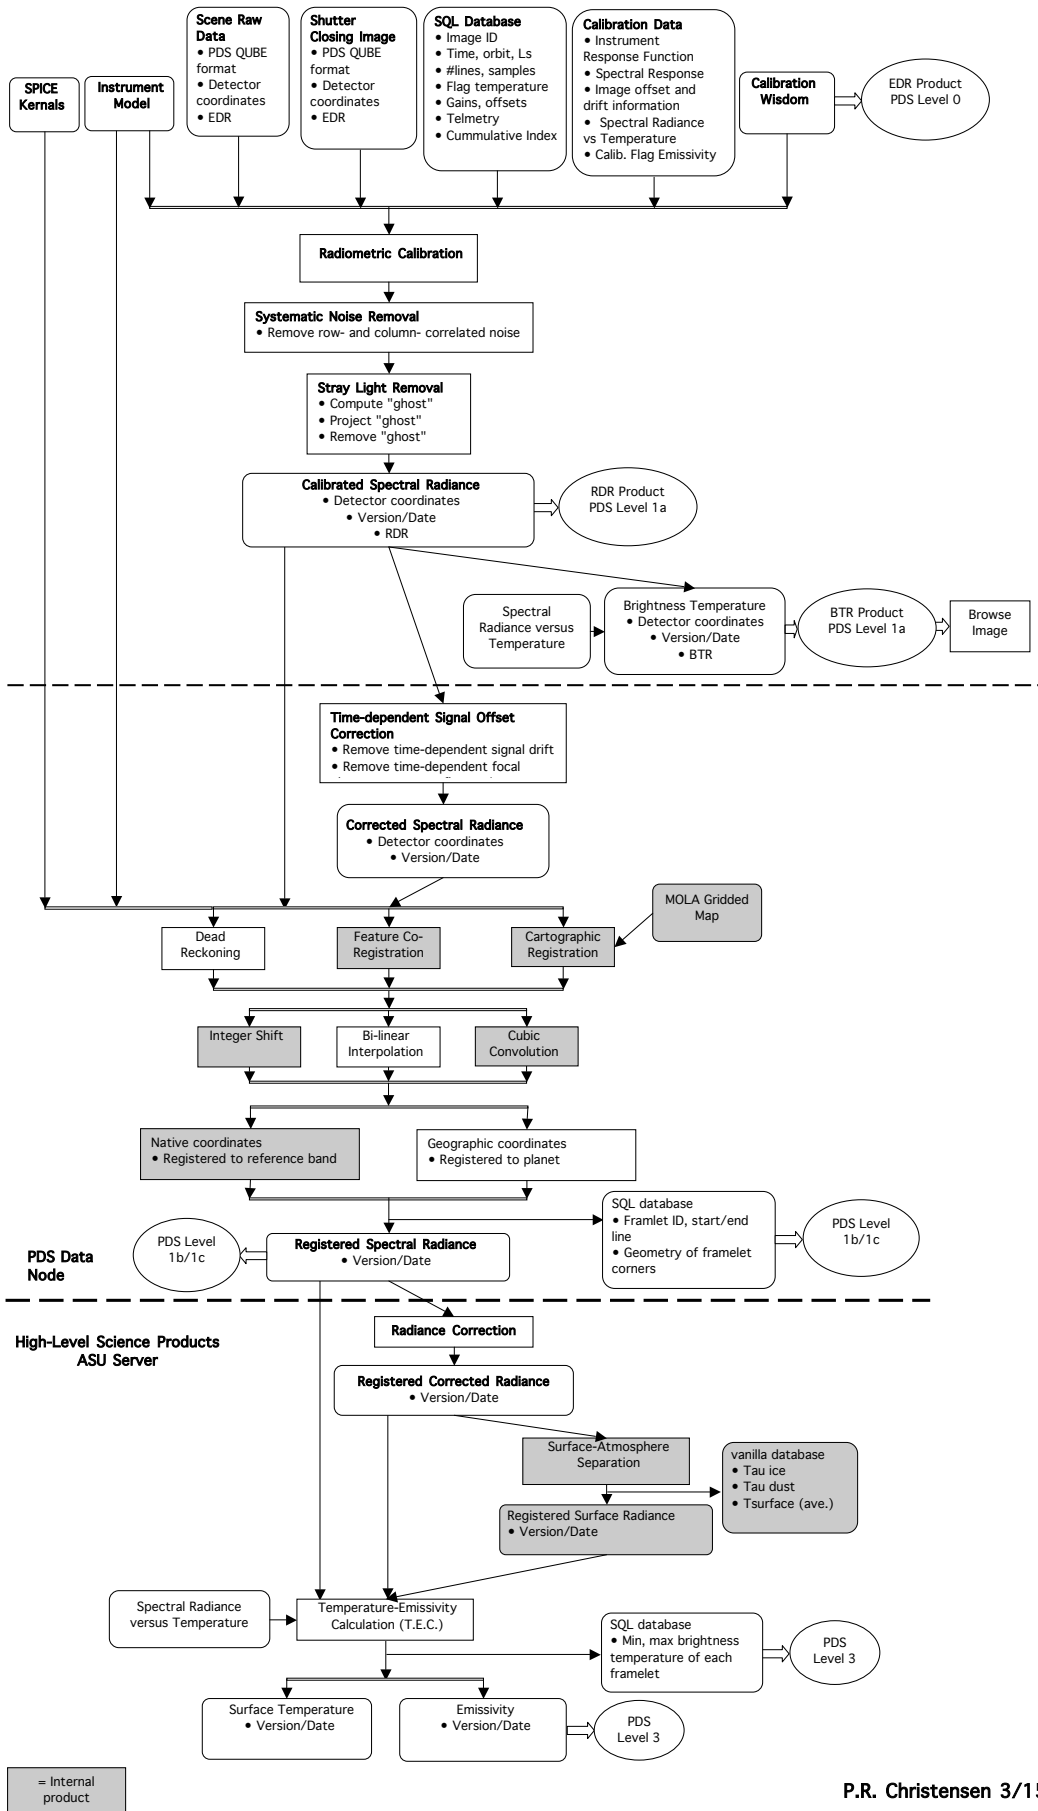
| | | | |
|-------------------|-----|---------|---|
| fit_response_full | 3.0 | 5/21/02 | Exclude 2 bad calibration test sets (Tinst=0 C, Target = 170 K; and Tinst = 35 C, Target = 307 K) |
| | 2.1 | 4/26/02 | Instrument temp groups 1-4 (-30, -15, 0, 15 C) |
| | 2.2 | 4/26/02 | Instrument temp groups 2-4 (-15, 0, 15 C) |
| | 2.3 | 4/26/02 | Instrument temp group 3 (0 C) |
| | 2.4 | 4/26/02 | Instrument temp groups 2-3 (-15, 0 C) |
| | 2.0 | 3/29/02 | Changed to use spectral radiance; all 5 instrument temperature groups (-30, -15, 0, 15, 30 C) |
| | 1.0 | 1/29/00 | Integrated radiance – early version |

| File | Version | Date | Description |
|-------------------------------------|----------------|-------------|---|
| temp_rad | | | |
| temp_rad_v4 | 3/29/02 | | Changed to use spectral radiance |
| temp_rad_v3 | 3/14/02 | | Changed temperature range to 100-380 C (from 120-360 C) |
| temp_rad_v2 | 1/4/02 | | Used better relative spectral response – filter* window* beamsplitter |
| temp_rad_v1.1 | 1/12/01 | | Temperature interval changed to 0.5 C |
| temp_rad_1 | 10/3/01 | | Temperature interval of 0.1 C |
| tv | | | |
| tv_6_1_2_v3.0 | 5/20/02 | | from make_response v3.0 cycle = 6; mirror_flag = 1; dn_flag = 2 (new get_flag_dn method; filter by 3); spectral radiance |
| tv_636_1_1_v3.0 | 5/16/02 | | from make_reponse v3.0; cycle = 636; used image 36 (gain =4; offset=0); mirror_flag = 1; dn_flag = 1; spectral radiance. (test case – did not make significant improvement) |
| tv_6_1_1_v3.0 | 3/29/02 | | from make_response v3.0 cycle = 6; mirror_flag = 1; dn_flag = 1; spectral radiance |
| tv_6_1_1_v2.0 | 2/14/02 | | from make_response v2.0 cycle = 6; mirror_flag = 1; dn_flag = 1 |
| tv_6_1_1_v1.1 | 11/12/01 | | from make_response v1.1 cycle = 6; mirror_flag = 1; dn_flag = 1 |
| tv_6_1_1_v1.0 | ?? | | from make_response v1.0; cycle = 6; mirror_flag = 1; dn_flag = 1 |
| irf_fit_all | | | |
| irf_fit_all_v3.0_tv6_1_2_3.0 | 5/21/02 | | fit of tv6_1_2_3.0 using fit_response_full v3.0 (new get_flag_dn; exclude 2 bad test sets in fit) |
| irf_fit_all_v2.0_tv6_1_2_3.0 | 5/20/02 | | fit of tv6_1_2_3.0 using fit_response_full v2.0 (new get_flag_dn; all 35 test sets) |
| irf_fit_all_v2.0_tv636_1_1_3.0 | 5/16/02 | | fit of tv636_1_1_3.0 using fit_response_full v2.0 |

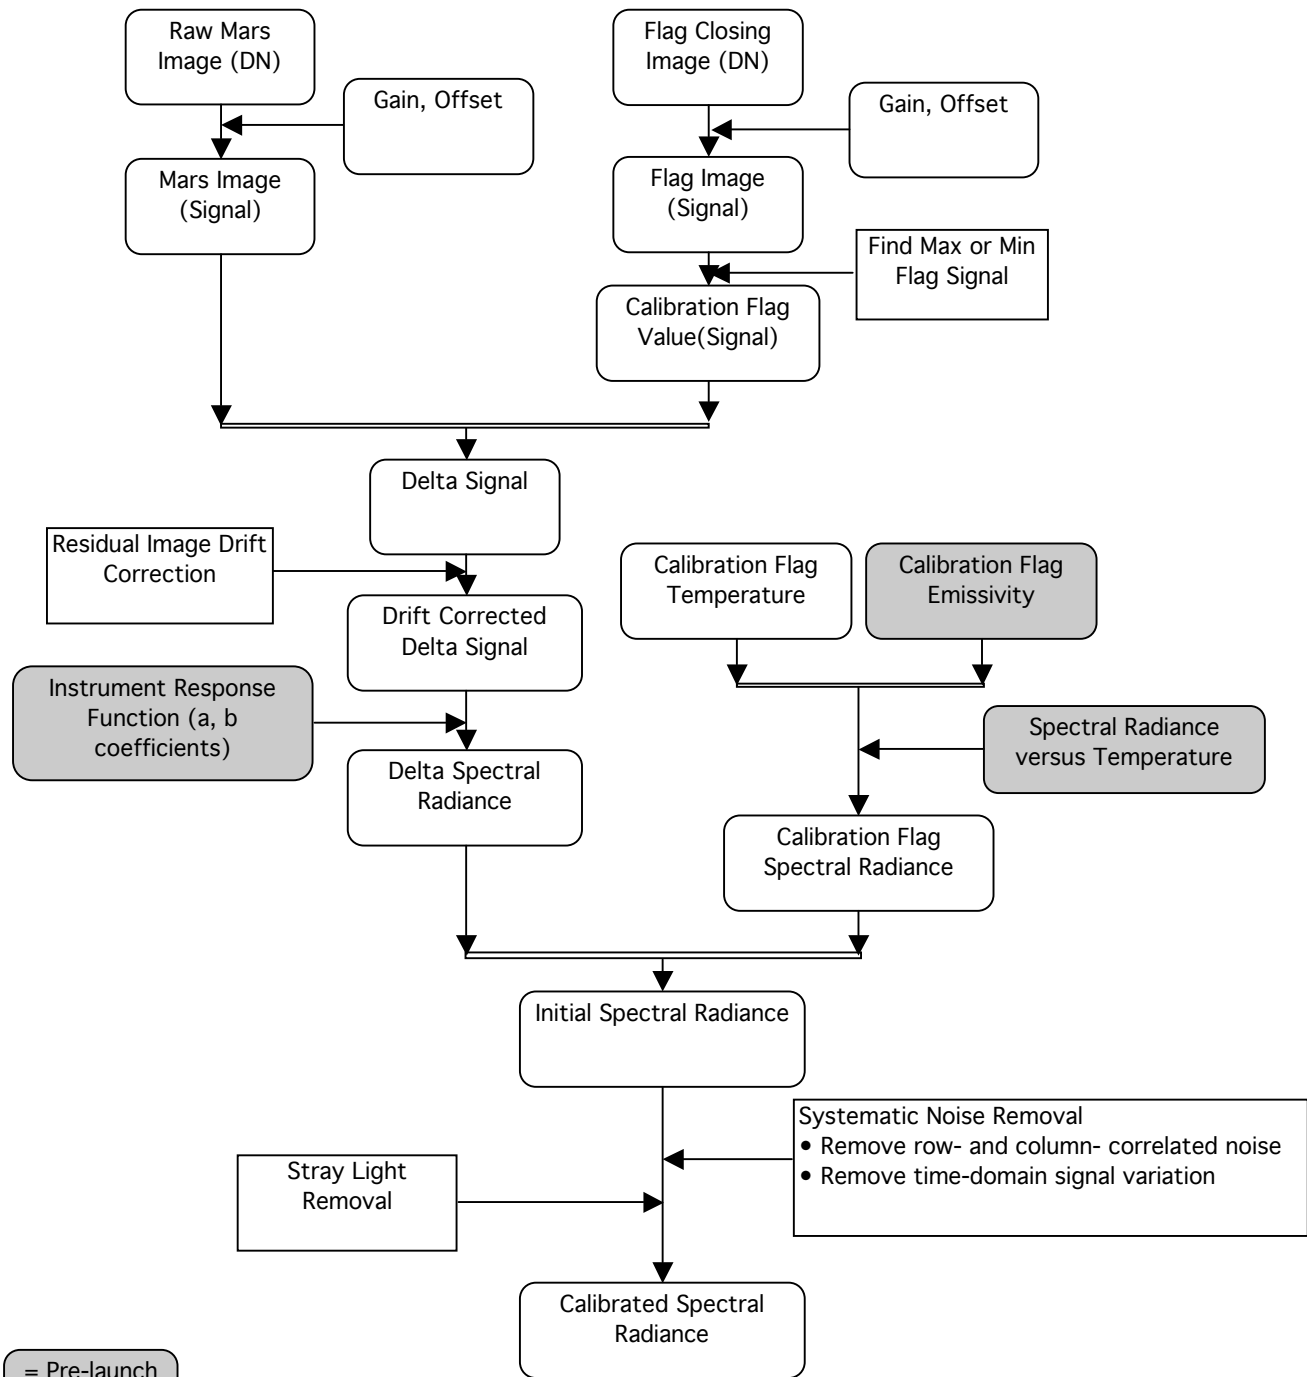
| | | |
|-------------------------------|----------------|--|
| irf_fit_all_v2.1_tv6_1_1_3.0 | 4/26/02 | fit of tv6_1_1_3.0 using fit_response_full v2.1 |
| irf_fit_all_v2.2_tv6_1_1_3.0 | 4/26/02 | fit of tv6_1_1_3.0 using fit_response_full v2.2 |
| irf_fit_all_v2.3_tv6_1_1_3.0 | 4/26/02 | fit of tv6_1_1_3.0 using fit_response_full v2.3 |
| irf_fit_all_v2.4_tv6_1_1_3.0 | 4/26/02 | fit of tv6_1_1_3.0 using fit_response_full v2.4 |
| irf_fit_all_v2.0_tv6_1_1_3.0 | 3/29/02 | fit of tv6_1_1_3.0 using fit_response_full v2.0 |
| irf_fit_all_v_1.0_tv6_1_1_2.0 | 2/18/02 | fit of tv6_1_1_2.0 using fit_response_full v1.0 |
| radiance_offset | | |
| radiance_offset_v1 | 6/27/02 | 3.09400e-7; 6.81523e-7; 9.26852e-6; 1.54012e-5; 1.42841e-5; 1.25434e-5; 1.19725e-5; 5.46783e-6; -1.98854e-6; -4.36478e-6 W cm-2 str-1 μm-1. |
| defocus_option | | |
| 13 | 9/6/02 | defocus_filter = ((29, 25, 19, 15, 9, 5) x 3) for bands 3 through 8 |
| tdi_smear_option | | |
| 1 | 9/6/02 | tdi_smear_filter = ((1,0,1,0,1,0,1,0,1,0, 1,0,1,0,1,0,1,0,1,0,1,0,1,0,1,0,1,0) x 3) filter |
| xdel_option | | |
| 8 | 9/6/02 | xdelta = (0, 0, 3, 3, 3, 3, 1, 1, 0, 0) pixel offset for bands 1 through 10 |
| percent_option | | |
| 5 | 9/6/02 | percent = (0., 0., 2.0, 4.5, 6.0, 5.5, 5.0, 5.0., 0., 0.) for bands 1 through 10 |

bold = currently used in standard calibration

THEMIS IR Data Processing Flow



THEMIS IR Mars Data Calibration Flow



= Pre-launch Data

3 VISIBLE IMAGER CALIBRATION

3.1 VISIBLE IMAGER CALIBRATION OVERVIEW

This section describes the calibration procedure that is applied to THEMIS-VIS raw Experiment Data Records (EDRs) in order to generate Reduced Data Records (RDRs). This section is provided by T. H. McConnochie and J. F. Bell, Cornell University (THEMIS-VIS Calibration Procedures, Version 2.0, 18 June, 2003).

The THEMIS-VIS calibration procedure ("pipeline") consists of, at present, seven steps: (1) 8-bit to 11-bit decoding; (2) identification of bad pixels; (3) bias subtraction; (4) "shutter" (frame transfer) smear subtraction; (5) correction of pixel sensitivity variations (*i.e.*, "flatfielding"); (6) stray light subtraction; and (7) radiance calibration. One known VIS calibration issue is not yet addressed by the calibration pipeline: namely, the current pipeline attempts only to identify and remove the spatially variable component of the stray light signal, and not the component that is uniform across the VIS detector. The outputs from the pipeline are VIS RDRs calibrated to radiance units ($\text{W cm}^{-2} \mu\text{m}^{-1} \text{sr}^{-1}$). It should be noted that atmospheric aerosols can be a substantial contribution to VIS radiances and can be highly variable spatially and temporally. VIS radiances, therefore, may not be simply interpretable in terms of surface reflectances.

3.2 VISIBLE DETECTOR DESCRIPTION

This discussion covers only those aspects of the VIS detector behavior that are directly relevant to the calibration. For a more detailed description of the THEMIS optical system and of the VIS detector, see Christensen *et al.* (2003) and Bell *et al.* (2003a). VIS uses a Kodak interline-transfer 1032x1024 CCD. For each column of photosites, there is a masked vertical register (v-register) adjacent to it. At the start of a VIS exposure, any charge accumulated on the detector is flushed. At the end of a nominal VIS exposure time, charge accumulated in the photosites is transferred immediately to the adjacent v-registers. The v-registers are masked to minimize accumulation of photo-electrons during readout. During the readout process, the charge is shifted "downstream," *i.e.*, down the v-registers and transferred to the CCDs horizontal register (h-register) one row at a time. In the EDR and RDR products, lines (*i.e.*, rows) closer to the h-register have higher line numbers. The h-register will be considered to be at the "bottom" of the detector array in this document, so that "down" and "downstream" are in the same direction. That is, if one row of pixels is said to be below or downstream of another row of pixels, then the line which is below or downstream is closer to the h-register. With this convention, the direction of spacecraft motion (south on the afternoon side of the orbit) is always down.

Physically, the detector has 1032 columns and 1024 rows. The 5 VIS filters are strips that are 1032 columns wide and approximately 200 rows tall and are bonded directly to the detector. In this document, as in the EDR and RDR labels on the PDS

files, the filters are labeled "filter 1", "filter 2", ... "filter 5" in order of their position on the detector array, with filter 1 at the bottom, closest to the register. In this naming scheme, "filter 1" is 860 nm, "filter 2" is 425 nm, "filter 3" is 654 nm, "filter 4" is 749 nm and "filter 5" is 540 nm. An alternate filter designation exists, using the labels "band 1", "band 2", ... "band 5" to refer to the filters in order of increasing wavelength [425, 540, 654, 749, and 860 nm (± 25 nm)]. The groundtrack of the Odyssey spacecraft runs approximately parallel to the detector columns. This allows multispectral images to be acquired using a timed sequence of VIS exposures. Any combination of one or more filters can be acquired by selectively reading out only those rows that correspond to the desired filters. Thus, a single VIS exposure will contain any or all of the five filters. The portion of an exposure that contains data for a single filter is referred to as a "framelet." The different filter framelets in a full-frame VIS exposure cover different regions of the surface, hence the need for a sequence of multiple exposures to build up a multispectral image using spacecraft downtrack motion to image the same parts of the surface through different filters (and thus at slightly different times). The exposure sequences are designed so that when a filter moves past the targeted region of the surface, it is no longer read out.

Rows and columns near the edges of the filters are always cropped during readout. The result is that each framelet has dimensions of 1024 by 192 detector pixels. The rows that are cropped are referred to henceforth as inter-framelet rows. When VIS EDR cubes are generated (in PDS QUB format), the VIS exposures in a sequence are broken up into their constituent framelets and re-assembled so that each framelet's position in the cube corresponds roughly to its spatial location on the surface of the planet. Thus, for a given filter, the corresponding plane of the EDR QUB is generated by concatenating the framelets from top to bottom in the order in which they were acquired. Assuming no spatial summing (spatial summing is discussed below) the first (top) framelet goes in rows 0 – 191, the second in rows 192 – 383, etc. Within a given EDR QUB plane, we assign the framelets a number, m , starting from the top, with the first framelet acquired in given filter being $m = 0$.

The bias and shutter smear removal procedures require that all of the framelets from a given exposure to be grouped together. Consider a framelet at position m_0 taken through filter number f_0 . The position m_1 of the framelet taken through filter number f_1 that was acquired in the same exposure as framelet (m_0, f_0) can be found by the following equation:

$$m_1 = m_0 + (f_0 - f_1). \quad (1)$$

If $m_1 < 0$ or $m_1 \geq n$, where n is the number of framelets in each filter in the EDR cube, then there was no framelet of filter f_1 acquired in the same exposure as framelet (m_0, f_0) . Note that this equation is written in terms of filter numbers, which are, as mentioned previously, not the same as band numbers. Conversely (and confusingly), the EDR and RDR QUB planes are stored in band order, not in filter order. The location of each filter within the EDR and RDR cubes are given by the BAND_BIN_FILTER keyword in the PDS labels.

VIS can operate in any of three spatial summing modes. The preceding discussion has assumed summing mode 1, *i.e.*, no spatial summing, for clarity. In summing modes 2 and 4, the VIS pixels are binned 2 x 2 and 4 x 4, respectively. This creates framelets with 512 columns by 96 rows, or 256 columns by 48 rows, respectively.

The amount of time required to read out a single VIS exposure depends both on the spatial summing mode and on the number of filters being read out. Each framelet of an exposure requires 141, 76, or 39 msec to read out, for summing modes 1, 2, and 4, respectively. A small amount of time (~ 4 msec) is also required to "dump" the charge from each filter that is not being read out. About 0.2 msec is required to dump each group of the filter-boundary rows, which are never read out. So, for example, it takes 703 msec to complete the readout process when summing mode 1 data from all five filters is being acquired. Since this readout time is long compared to the ~ 5 -10 msec typical exposure time, the opacity of the masks that protect the v-registers from photo-electrons, as well as the amount of time that each row of data spends in the v-registers during readout, are crucial factors for the VIS calibration process. Since VIS dumps the charge in filters not being read out, the amount of time a given row spends in the v-registers depends not only on its distance from the h-register, but also on which of the filters that are below it on the array are being read out.

Calibration steps 3 and 4 described below (bias and "shutter" smear subtraction) require knowledge of the timing for each data pixel as it is clocked down the v-registers to the h-register during readout. For pixels in a given filter, this timing depends on which of the lower-numbered (closer to h-register) filters are being read out in a particular exposure, and which are being dumped and not read out. Pixels in filter number f have $2^{(f-1)}$ possible timing "paths" during readout (*e.g.*, for filter 3, these are the read out filter combinations [3,2,1], [3,2], [3,1], and [3]). The total number of possible paths is therefore

$$\sum_{f=1}^5 2^{(f-1)} = 31 \quad (2).$$

These 31 possible paths will henceforth be referred to as "filter paths." Each path is given a filter path code, F . For a given filter with filter number f_0 , F is determined by:

$$F = \sum_{f=1}^{f_0} \mathbf{e}(f) \cdot 2^{(f-1)} \quad (3).$$

$\mathbf{e}(f) = 1$ IF filter f is included in the exposure

$\mathbf{e}(f) = 0$ IF filter f is not included in the exposure

A filter is included in an exposure if both of the following are true: (1) the filter is present in the EDR; and (2) the framelet position m of the filter (as calculated from equation (1)) satisfies $m_l < n$ and $m_l \geq 0$, where n is the number of framelets of each filter in the EDR cube.

3.3 VISIBLE IMAGER CALIBRATION ALGORITHM

3.3.1 8 bit to 11 bit decoding.

The VIS signal is originally sampled with 11 bits per pixel (0-2047 DN), but the data is reduced to 8 bits per pixel prior to downlink using a square-root encoding look-up table algorithm. This step simply applies an inverse look-up table to restore the VIS data to its original linear 11 bit format. [Section 3.4.1].

3.3.2 Identification of bad pixels.

Bad pixels, also known as null pixels, are assigned the value specified by the CORE_NULL keyword in the PDS RDR label, and are ignored in all subsequent VIS processing.

a) Threshold values

All pixels with a DN level of 2040 (the highest possible value) or 0 are flagged as null pixels.

b) Bad rows and columns

Some rows and columns near the edge of each framelet are always filled with unusable (noisy or saturated or zero) data. The pixels in these rows and columns [Section 3.4.2] are therefore flagged as null.

c) Saturated regions

In exposures where all or part of the detector is saturated, the DN value of some saturated pixels is "wrapped" to a value which is invariably small compared to the rest of the pixels in the array. These "wrapped" pixels are identified by considering each VIS framelet separately. Any pixel whose value is less than the framelet median by at least 1200 DN has probably been "wrapped" and is therefore flagged as null. The median is calculated using all pixels not flagged as null in steps 2.1 and 2.2.

d) Neighboring pixels

If too many of the pixels near a given pixel are flagged as null by steps 2.1 or 2.3, that pixel is also flagged as null. A 5 by 5 pixel region centered on each pixel is considered. This region is truncated if it would extend past the edge of the array. Pixels that lie in the "bad rows and columns" identified in step 2.2 are always treated as *valid* for this procedure. If more than 30% of the pixels in the 5x5 (truncated if necessary) region for a given pixel are null, then that pixel is also flagged as null.

3.3.3 Bias subtraction

a) Nighttime VIS images

The bias subtraction procedure uses VIS images acquired at night with nominal exposure times of zero. The bias is determined independently for each of the three available spatial summing modes using nighttime images with the same spatial summing mode. In practice there is no detectable difference between night images acquired with typical VIS exposure times (up to thirty milliseconds) and night images with zero exposure time, so non-zero exposure time night images are also included in the averaging process that is used to determine the bias. This also means that the dark current for typical VIS exposure times is *effectively* zero. However, because the VIS readout time is much longer than the VIS exposure times, some of the signal for zero millisecond nighttime exposures may be attributable to dark current accumulated during readout. The evidence for this is a slight ramp-up in the zero-millisecond signal towards the top of the detector. Fortunately, since the amount of dark current that accumulates during readout is not a function of exposure time, and since we have detected no temporal variability in the zero exposure night-time images, there is no need to distinguish between bias and readout dark current for the purpose of calibration. We will therefore henceforth lump both effects together and refer to the combination of the two simply as "bias."

b) Averaging to create bias frames

To properly account for the small ramp in the bias level, bias frames are created for each of the 31 filter paths and for each of the 3 spatial summing modes. Where possible, each bias frame is generated from a simple average of all available nighttime exposure framelets with the same summing mode and filter path.

c) Modeling to handle paths for which no empirical data is available

For many of the less common filter paths, nighttime exposure framelets have not yet been acquired. In these cases, we apply the following simple model of the detector readout process in order to approximate the expected bias signal:

Consider what happens to the charge from all upstream filters while a downstream filter is being read out. As our first simplifying assumption, we ignore the inter-framelet rows (there are 7 ± 2 inter-framelet rows at each filter boundary) and pretend that the top row of one framelet is always adjacent to the bottom row of the next. As each row of the downstream filter is shifted into the h-register, the upstream filter rows are shifted one row closer to the h-register, so that by the time the entire downstream filter has been read out, the rows of each upstream filter f have shifted so that they now lie under the filter $f - 1$. The charge that these rows accumulate in that time is characteristic of filter f . When the next downstream filter is read out, these same rows are shifted from lying under filter $f - 1$ to lying under filter $f - 2$, with the charge accumulated during the shift being characteristic of the starting filter for that shift. When the rows of filter f are finally in the

position where filter f rows are being transferred directly to the h-register, the charge that they accumulate is characteristic of filter 1. Thus, if all of the filters downstream from f are read out in an exposure, then the rows in filter f accumulate charge from filters $f, f-1, f-2, \dots, 1$.

If any of the filters downstream from f is not being read out, then the process of dumping that filter shifts f downstream very rapidly, so if the rows in filter f were under filter $f-x$, at the start of the dump, very little of the characteristic bias charge of filter $f-x$ is accumulated. Our second simplifying assumption is that the charge accumulated during such a dump is negligible.

The bias charge built up during readout for a certain filter following a given filter path with filter path code F can therefore be calculated simply by summing the characteristic bias charges, $\mathbf{E}(f)$, of each downstream filter. The bias frame $\mathbf{B}(F)$ is:

$$\mathbf{B}(F) = \sum_{f=1}^{f_0} \mathbf{e}(f) \cdot \mathbf{E}(f_0 - f + 1). \quad (4)$$

The values of f_0 and $\mathbf{e}(f)$ are properties of the filter path denoted by F and are defined by equation (3), *i.e.*, f_0 is the starting filter of the filter path and $\mathbf{e}(f)$ indicates whether the filter f is one of those being read out for that filter path. $\mathbf{E}(f)$ and $\mathbf{B}(F)$ are arrays with the same dimensions as the framelets for the spatial summing for which the $\mathbf{B}(F)$ are being derived. Note that $\mathbf{e}(f)$ is always equal to 1 when $f = f_0$, so the characteristic bias of the first filter $\mathbf{E}(1)$ is always present in the bias frame (this is one reason why it is always advantageous to obtain data in filter 1 (860 nm), despite the challenging nature of the scattered light correction for this filter as described below).

$\mathbf{E}(f)$ can be derived by using empirically known (via the averaging method described in section 3.2) $\mathbf{B}(F)$ arrays. We perform this derivation by using $\mathbf{B}(F)$ values for those filter paths for which all downstream filters are read out. These are $\mathbf{B}(31)$, $\mathbf{B}(15)$, $\mathbf{B}(7)$, $\mathbf{B}(3)$, and $\mathbf{B}(1)$. Writing out equation (4) explicitly for these five $\mathbf{B}(F)$ leads to the following trivial system of equations:

$$\begin{aligned} \mathbf{B}(F=31) &= \mathbf{E}(f=1) + \mathbf{E}(f=2) + \mathbf{E}(f=3) + \mathbf{E}(f=4) + \mathbf{E}(f=5) \\ \mathbf{B}(F=15) &= \mathbf{E}(f=1) + \mathbf{E}(f=2) + \mathbf{E}(f=3) + \mathbf{E}(f=4) \\ \mathbf{B}(F=7) &= \mathbf{E}(f=1) + \mathbf{E}(f=2) + \mathbf{E}(f=3) \\ \mathbf{B}(F=3) &= \mathbf{E}(f=1) + \mathbf{E}(f=2) \\ \mathbf{B}(F=1) &= \mathbf{E}(f=1) \end{aligned} \quad (5)$$

d) Bias subtraction algorithm

Bias frames for each summing mode are stored in FITS files each with 31 planes corresponding to the 31 filter paths. The bias is removed by determining the filter path for each framelet of each filter and then subtracting from it the stored bias frame for that

filter path and summing mode. The bias frames are stored in the FITS files in order of the F code, *i.e.*, the bias frame for a given filter path is found at plane number $F - 1$ for a zero-based cube plane counting system.

3.3.4 "Shutter" smear subtraction

a) Extrapolation of ground calibration data to zero exposure duration

The VIS ground calibration data set contains observations of an integrating sphere at 7 different lamp brightness settings with exposure times of 3, 6, 12, and 24 milliseconds. The "shutter" smear effect was not recognized during the ground calibration campaign, so no zero millisecond exposure times were used. However, we were able to derive a good approximation to a zero exposure image by fitting a line to DN level as a function of exposure time for every pixel at each lamp level. The bias level is found by applying the linear fitting procedure to data for which the lamps were turned off. The intercept of these linear fits is the zero exposure DN level, and the slope of these fits is the detector response in DN/msec.

The zero exposure DN levels show a ramp-up towards the top of the detector array that is monotonic and continuous across filter boundaries. The slope of this ramp, which we will call the "smear slope," is strongly but smoothly varying from column to column, producing an obvious pattern of broad vertical stripes. The zero exposure DN level does not go to zero at the bottom column of the array. Rather, it reaches, at the bottom row of filter 1, a nonzero value that is large compared with the slope and is roughly constant across all columns. We will use the term "smear intercept" to refer to this row of roughly constant nonzero values found at the bottom row of the bottom filter of zero exposure images.

The smear slope and smear intercept are both proportional to the lamp radiance. Since all of the lamps used during ground calibration tests had approximately the same color, it was not possible to determine the spectral behavior of the smear slope and smear intercept signal. The proportionality of smear to radiance suggests that the source of the zero exposure signal is light reaching both the v -registers (producing the smear slope) and the h -registers (producing the smear offset). The smear slope shows little variation from filter to filter, in stark contrast with the detector response, which is highly variable between filters due to spectral variations in filter transmissivity, lamp radiance, and detector quantum efficiency (QE). The small amount of filter-to-filter smear slope variation is not correlated with the detector response variations. Furthermore, the detector response shows strong intra-filter spatial variations (due to stray light, which is discussed in section 5), which are not highly correlated with the spatial variations in smear slope. These two pieces of evidence suggest that the zero exposure signal light reaching the h -registers and v -registers is leaking *around* the filters rather than passing through them, and so will have little spatial or spectral correlation with the target scene. We therefore model the shutter smear by assuming that the mean value of the smear is proportional to the mean broadband scene radiance. To remove the smear, we proceed by

deriving normalized shutter smear frames, and then establishing the scaling factors between mean radiance and mean shutter smear.

b) Zero duration daytime VIS images

In-flight VIS images acquired during daylight with a commanded exposure time of zero msec show exactly the same spatial pattern as the ground-calibration-derived zero exposure images. Clearly, a commanded exposure time of zero does not really produce an exposure time of exactly zero. However, in the majority of zero exposure flight data, there is no evidence of surface features visible through the shutter smear pattern. The spatial variation of the small amount of scene derived DN that is probably present in the zero exposure images will be washed out by averaging a sufficient number of images. Furthermore, the mean level of scene-derived contamination of the zero exposure images will be a negligible fraction of the mean signal of any given VIS image as long as the effective exposure time for a commanded exposure time of zero is a negligible fraction of the effective exposure time for the VIS image in question.

Shutter smear correction frames are generated using zero duration daytime VIS images which have already been bias corrected. The zero exposure daytime images are first averaged according to filter path and binning mode in the same manner as the nighttime VIS images are averaged to create bias frames. During the averaging process, the smear-frame generation algorithm also stores the mean DN levels for all of the framelets in each exposure. Using this mean DN level information, the smear frames for each filter path are normalized so that ratio of the means for any two filters paths in the smear frames is equal to the mean value of that filter path ratio for all of the individual VIS exposures in which both filter paths are present. The mean of the $F = 1$ filter path smear frame is set equal to unity.

Obviously, there are many filter paths that cannot occur in the same exposure as the $F = 1$ path. However, every filter path *does* occur in the same exposure as one of $F = 16$, $F = 8$, $F = 4$, $F = 2$, or $F = 1$ paths, which we will refer to as “direct” paths (these are the paths that apply to filters 5, 4, 3, 2, and 1, respectively, when there are no other filters downstream). The ratios of the direct filter paths are established with respect to the $F = 1$ path. Then the ratios of all other filter paths are established with respect to whichever direct path they occur with, and then rescaled using the ratio of that direct path to the $F = 1$ path, so that the ratios of all paths are known relative to the $F = 1$ path.

The direct-to- $F = 1$ ratios are established as follows: In VIS images that use all five filters, an exposure containing $F = 2$ always occurs immediately following an $F = 1$ exposure, an exposure containing $F = 4$ always occurs immediately following an $F = 2$ exposure, etc. So the $F = 2$ ratio is calculated relative to the $F = 1$ from the previous exposure in the same image, the $F = 4$ is calculated relative to the previous exposure $F = 2$, $F = 8$ relative to the previous $F = 4$, and $F = 16$ relative to the previous $F = 8$. These ratios are then rescaled in sequence so that they are expressed relative to $F = 1$.

The objective of the (admittedly complicated) ratio procedure described above is to establish the proper inter-filter and inter-filter-path normalization. This normalization must be as accurate as possible, because the only means that we have available for determining the scaling factor between radiance and shutter smear is to measure the changes in mean DN from framelet to framelet within an image as a result of changes in filter path, and then find the scaling factor that eliminates these changes in mean DN (this procedure is discussed further below). The normalization scheme assumes that the smear signals from the various filter paths are related to each other by a constant proportionality factor. We have examined the smear signal ratios that were used to generate the mean smear signal ratios, and have found that they are consistent with the constant proportionality hypothesis.

The normalizing ratios and normalized smear frames are generated independently for each summing mode. Obviously, the normalized smear frames are generated as described above only for the filter paths that are represented in our data set of zero exposure daylight images. Smear frames for filter paths that are not represented in the data set of zero exposure daylight images are generated by starting from the normalized smear frames and then applying the same method used for the bias frames, as described in section 3.3.

c) Deriving the radiance-to-smear scaling factors

As previously mentioned, we model the mean "shutter" smear as being proportional to mean radiance. In order to solve for the correct scaling factor, we postulate that, *on average*, once the shutter smear affect is removed, the framelet means will be independent of the filter path, provided that we consider framelets from the same filter. To take advantage of this postulate, we first determine, for each 5-filter image of a given summing mode, the factor \mathbf{y} that we need to multiply the normalized smear frames by so that when they are subtracted from the mean image framelets, the differences between framelet means are minimized in a least-squares sense. More explicitly, for each 5-filter image, we: (1) Identify the set of framelet pairs which satisfy: (a) both members of the pair are from the same filter; (b) the members of the pair come from adjacent exposures; and (c) the members of the pair have different filter paths. Let the elements of this set be indexed by i . (2) Calculate the difference in the mean (bias subtracted) DN for each pair, d_i , and the difference in the mean normalized shutter smear for the filters paths of the members of each pair, s_i . (3) Find the value of \mathbf{y} which minimizes

$$\sum_i (\mathbf{y}s_i - d_i)^2 \quad (6).$$

(4) For each member of each pair, store \mathbf{y} and M for later use. M is the shutter-smear-corrected mean DN for a framelet, *i.e.*, if S is the mean normalized shutter smear for a framelet's filter path and D is the mean uncorrected DN, then $M = D - \mathbf{y}S$. Clearly, the \mathbf{y} that we find for any given image will be biased by scene variations. However, these biases should be averaged away when we consider the entire set of \mathbf{y} generated by the full VIS data set.

We now seek to determine the scaling factor between \mathbf{y} and mean radiance. Note that at this stage of the calibration, we of course cannot solve for the true radiance. What we do have is shutter-smear corrected DN/msec, M / t , which we will call "photosite signal." The mean photosite signal is of course proportional to the mean radiance, but *that* proportionality constant will not be derived until step 7.

Based on our light leak model for the origin of the shutter smear, we expect \mathbf{y} to be proportional to the scene-averaged QE-weighted wavelength-integrated radiance reaching the VIS focal plane. Our goal, however, is to predict \mathbf{y} using what we have available, which is simply the narrow-band photosite signal from whatever filters are present in a given image. We therefore seek a proportionality constant, \mathbf{G}_f for each filter f that relates M_f / t to \mathbf{y} . Using j to index all of the M and \mathbf{y} values that have been solved for, \mathbf{G}_f is found by minimizing

$$\mathbf{c}_f^2 = \sum_j \left(\Gamma_f \frac{M_{jf}}{t_j} - \mathbf{y}_j \right)^2 \quad (7).$$

The results of the optimization procedure show that the lowest \mathbf{c}_f^2 occurs for $f = 3$. Since the $f = 3$ (650 nm) filter falls at the peak of the typical Mars radiance spectrum, it is not surprising that this filter is the best predictor of the shutter smear given our assumption that shutter smear is proportional to broad band radiance. For the $f = 3$ case, we keep the \mathbf{G}_f derived with equation (7). For the other four filters, we establish \mathbf{G}_f by determining the scaling (using least squares analogously to equation (7)) between M_3 and M_f , and then multiplying by \mathbf{G}_3 .

d) Subtraction of scaled shutter-smear frames

Smear frames for each summing mode are stored in FITS files each with 31 planes corresponding to the 31 filter paths. The smear frames are stored in order of the F code, *i.e.*, the bias frame for a given filter path is found at plane number $F - 1$ for a zero-based cube plane counting system. For the RDR data products, final shutter-smear subtraction is performed by first using the \mathbf{G}_f factors to calculate \mathbf{y} for each exposure, and then subtracting the \mathbf{y} -scaled smear frames for the appropriate filter path from each framelet of each exposure. To prevent abrupt changes in the multiplier \mathbf{y} , all \mathbf{y} values within a VIS image are calculated using the DN levels and \mathbf{G}_f factors for the same filter. \mathbf{y} is extrapolated for those exposures that do not contain the chosen f . The filter f used to calculate \mathbf{y} is chosen from the following priority list: 1, 3, 4, 5, 2, with filter 1 being the preferred filter and filter 2 being the filter of last resort. Once an f value is chosen, the multipliers for each exposure are calculated by:

$$\mathbf{y} = \frac{D_f \Gamma_f}{t + \Gamma_f S_F}, \quad (8)$$

where t is the exposure duration, S_F is the normalized mean shutter-smear for the filter path F followed by the filter f in the given exposure, and D_f is the uncorrected mean DN level.

3.3.5 Correction of pixel-to-pixel sensitivity variation.

When the THEMIS-VIS instrument is presented with a spatially uniform scene, the signal measured by the detector is non-uniform. The usual approach is to assume that the non-uniform detector response to a uniform scene is simply a result of sensitivity variations among the individual pixels. Unfortunately, the presence of a significant amount of stray light (discussed below) reaching the VIS detector means that this approach is invalid.

The crucial distinction between sensitivity variations and stray light is that sensitivity variations are a multiplicative effect, while stray light is additive, so that regions of the detector with greater sensitivity will have enhanced contrast, while regions with more stray light will have decreased contrast. This means that if we incorrectly attribute all of the non-uniform detector response to sensitivity variations, our calibration process will create spuriously decreased contrast in regions of high stray light. We have therefore adopted a method of measuring sensitivity variations that relies on the scene contrast directly, by using framelet to framelet signal *differences*, rather than averages as in the usual flatfielding approach.

The sensitivity correction is derived independently for each of the five filters, using shutter-smear corrected data as produced by step 4. Let \mathbf{L}_n and \mathbf{L}_n^* be the shutter-smear-corrected signal, respectively, from the early and later of members of framelet pair n . Let \mathbf{s}_n and \mathbf{s}_n^* be the arrays of standard errors for the member of the pair. These standard errors are based on Poisson counting noise statistics, *i.e.*, they are simply the square root of the number of electrons detected at each pixel, so:

$$\mathbf{s} = \sqrt{(\text{electrons per DN}) \cdot (\text{raw DN level})} / (\text{electrons per DN}). \quad (8)$$

The "raw DN level" refers to the DN level measured prior to bias subtraction, *i.e.*, after decoding to 11-bit values and null pixel identification. The gain is 25.4 electrons per DN for spatial summing mode 1, 25.4/2 for spatial summing mode 2, and 25.4/4 for spatial summing mode 4. A framelet pair consists of any two adjacent framelets from the same filter and same VIS data cube. The total number of such framelet pairs in the data set is labeled N . We assume that the stray light signal is slowly varying from framelet to framelet, so that the difference between framelets depends only on the changes in the scene and on counting noise. Since the magnitude of the differences caused by scene changes is proportional to sensitivity, we estimate the sensitivity \mathbf{S} by subtracting the counting noise estimate from the root-mean-square framelet-to-framelet differences:

$$\mathbf{S} = \frac{1}{N} \sqrt{\sum_{n=1}^N \frac{(\mathbf{L}_n^* - \mathbf{L}_n)^2 - (\mathbf{s}_n^*)^2 - (\mathbf{s}_n)^2}{\text{mean}(\mathbf{L}_n)^2}} \quad (9)$$

Unfortunately, the assumption of slowly varying stray light is not always valid, leading to some contamination of our sensitivity estimate by stray light effects. This contamination is most severe where the stray light is most severe, on the right and left edges of all filters, and across the entire filter 1 (band 5, 860 nm). However, excepting

the regions of severe contamination, we see no evidence that the sensitivity correction is a function of detector column. We therefore treat the sensitivity as being a function only of detector row, allowing us to derive it by averaging across the central, least contaminated pixels. We perform this column averaging using summing mode 2 values for \mathbf{S} , and we then normalize the resulting profile to unit mean. The sensitivity correction as function of row, $S(y)$, is therefore:

$$S(y) = \frac{\sum_{x=200}^{400} \mathbf{S}(x,y)}{\frac{1}{96} \sum_{y=0}^{95} \sum_{x=200}^{400} \mathbf{S}(x,y)} \quad (11)$$

Each column of each framelet is corrected by dividing it by $S(y)$. Since the $S(y)$ is originally generated using only summing mode 2 data and is thus a 96 element vector, it is resampled to 192 elements via linear interpolation for use with summing mode 1 data, and resampled to 48 elements via pixel averaging for use with summing mode 2 data.

3.3.6 Stray light subtraction.

The most prominent feature of un-calibrated VIS framelets is the broad bright features spanning most of the rows on the left and right edges of the framelets in all five filters. In 860 nm filter framelets this bright pattern spans all columns as well, occupying a large fraction of the framelet. The shape of these bright features is constant across all framelets in the data set, but their amplitude is variable. These bright features are seen to obscure, *i.e.*, reduce the apparent contrast of, scene elements. If the bright features were caused by increased sensitivity in that portion of the detector, then the apparent contrast of scene elements in that portion of the field of view would be enhanced rather than diminished. We therefore conclude that the bright features are caused by stray light that is being *added* to the scene brightness. Our calibration strategy is therefore to solve for the stray light contribution to a given framelet, and then subtract it.

Stray light correction frames are derived independently for each spatial summing mode using framelets that have been bias corrected, shutter-smear corrected, and corrected for pixel-to-pixel sensitivity variations. Let $\mathbf{M}_{x,y}^z$ be the array of framelet DN values for a given band z , M^z be the mean of $\mathbf{M}_{x,y}^z$, and $\mathbf{m}_{x,y}^z = \mathbf{M}_{x,y}^z - M^z$. Let $\mathbf{N}_{x,y}^z$ be the stray-light corrected DN values, and $\mathbf{D}_{x,y}^z$ be the stray light, with N^z and D^z being the means and $\mathbf{n}_{x,y}^z$ and $\mathbf{d}_{x,y}^z$ being the deviations from the mean. The framelet DN is modeled as a linear combination of the stray-light-corrected signal and the stray-light signal, *i.e.*,

$$\mathbf{M}_{x,y}^z = \mathbf{N}_{x,y}^z + \mathbf{D}_{x,y}^z . \quad (12)$$

We hypothesize that the stray light signal is caused by the same phenomenon as the shutter smear signal, namely, broadband light leaking under the edge of the filters. Just as this leaked broadband light reaches the v-registers, it also reaches the photosites, producing a characteristic pattern in every framelet. We have adopted the simplest model for determining $\mathbf{D}_{x,y}^z$ that is consistent with the broad band light leak assumption. This

model assumes that the stray signal for every pixel of every band, $\mathbf{D}_{x,y}^z$ is simply proportional to the mean signal in the 650 nm band, M^3 . The 650 nm band is chosen for consistency with the shutter smear subtraction method. Using this model, it is not possible to distinguish the mean stray-light-corrected signal N^z from the mean stray light signal D^z . We therefore in practice model only the spatially variable component of the stray light, $\mathbf{d}_{x,y}^z$. With proportionality constant $\gamma_{x,y}^z$, the model is

$$\mathbf{d}_{x,y}^z = \mathbf{g}_{x,y}^z M^3, \quad (13)$$

so by substituting into (12) and separating out the means and deviations from means, we have

$$M^z + \mathbf{m}_{x,y}^z = N^z + D^z + \mathbf{n}_{x,y}^z + \mathbf{g}_{x,y}^z M^3. \quad (14)$$

The means in (14) cancel, so we can write it as

$$\frac{\mathbf{m}_{x,y}^z}{M^3} = \frac{\mathbf{n}_{x,y}^z}{M^3} + \mathbf{g}_{x,y}^z \quad (15).$$

Since we have already removed the pixel-to-pixel sensitivity variations (the “flatfield”), and since there is no reason to believe that the set of scenes presented to VIS over the course of the mission is biased towards any particular pixel values, we expect the deviations from the mean of the stray-light corrected DN to average to zero over the entire data set, *i.e.*, indexing each framelet in the data set with j , with a total number of framelets J in each band,

$$\frac{1}{J} \sum_{j=0}^{J-1} \mathbf{n}_{x,y,j}^z = \mathbf{0} \quad (16).$$

By the same argument, we also expect

$$\frac{1}{J} \sum_{j=0}^{J-1} \frac{\mathbf{m}_{x,y,j}^z}{M_j^3} = \mathbf{0} \quad (17).$$

Since $\gamma_{x,y}^z$ is a constant for the data set, equations (15) and (17) immediately give us a solution for $\gamma_{x,y}^z$:

$$\mathbf{g}_{x,y}^z = \frac{1}{J} \sum_{j=0}^{J-1} \frac{\mathbf{m}_{x,y,j}^z}{M_j^3} \quad (18).$$

With this solution in hand we generate RDR framelets that have the spatially variable stray light component removed as follows:

$$D^z + \mathbf{N}_{x,y}^z = \mathbf{M}_{x,y}^z - \mathbf{g}_{x,y}^z M^3 \quad (19).$$

This stray light correction is performed on spatially coincident groups of framelets. This is in contrast to the shutter smear and bias corrections which are performed on groups of framelets that were acquired in the same exposure. A spatially coincident group of framelets is the set of framelets from different bands that have the same framelet number. In cases where band 3 is not present in a framelet set, the band 3 mean signal M^3 is estimated from mean signal in another band by using the data-set mean

ratio of that band's signal to the band 3 signal. The band used to estimate the band 3 mean is chosen in the following priority order: 3, 4, 5, 2, 1.

It is crucial to note that data stored in the RDRs is simply $D^z + \mathbf{N}_{x,y}^z$ scaled by the exposure time and a radiance conversion factor. In other words, the current version (v2.0) of calibrated RDRs still contains a spatially uniform stray light component D^z which, although constant within a given framelet, is variable across the data set. This spatially uniform stray light component will *not* bias the *mean* radiance in any given framelet, since the spatially uniform stray light component is also present in the data set used to generate the radiance conversion factors. However, it *will* decrease the contrast between low and high albedo regions, as well as shadowed and un-shadowed regions, and could lead to spurious color variations that are correlated with albedo and lighting variations.

Examination of the RDRs shows that the algorithm of (18) and (19) has not completely removed the spatially variable component of the stray light. Apparently, the stray light signal is not simply related to the framelet mean, and may depend on the brightness of regions outside of the field of view. However, we have not yet found a more complicated stray light model that produces significantly better results than the framelet-mean method we have described here. This will be the subject of additional analysis and will hopefully be addressed in future iterations of the VIS calibration pipeline.

3.3.7 Radiance calibration.

Absolute calibration of VIS images is required in order to generate true-color data products and to extract spectra from surface units to compare directly with laboratory mineral and mineral mixture spectra. Because the VIS pre-flight ground calibration data were not obtained using the same readout/clocking mode as the VIS flight data, additional analyses and corrections to the preflight data are required before attempting to use them to derive an absolute radiance calibration. In the meantime, we have adopted an *ad hoc* or "bootstrap" radiance calibration approach based on the use of well-calibrated Hubble Space Telescope (HST) Wide Field/Planetary Camera 2 clear-atmosphere observations from 255 to 1042 nm (e.g., Bell *et al.*, 1997; 2003b). Our approach is twofold: First, we generate a "VIS Average Mars" calibrated (with steps 1 through 6) DN spectrum by averaging all daylight (incidence angle $> 88^\circ$), 5-color, summing mode 2 framelets. Before including DN values in the average, we normalize them to a standard viewing geometry of 60° incidence, 0° emission by applying a Minnaert photometric function with $\kappa = 0.7$. We then generated an "HST VIS-like Average Mars" spectrum by averaging all of the non-limb and non-water ice cloud pixels from a photometrically-corrected 1999 global HST observational data set (Bell *et al.*, 2001, 2003b; Morris *et al.*, 2003). We use a Minnaert ($\kappa = 0.7$) photometric function to renormalize this HST average spectrum to the same standard (60° incidence, 0° emission) viewing geometry used for VIS. Convolution of the HST average radiance spectra over the VIS bandpasses and dividing by the VIS average DN spectrum yields a set of coefficients [in

(DN/msec)/(W/m²/nm/sr)] that is used to scale VIS images processed to Step 6 above to radiance. Optionally, the pipeline can also output calibrated images in I/F (where I is the observed flux and πF is the irradiance of the Sun at the top of the Martian atmosphere at the time of the VIS observation; Bell *et al.*, 2003b). VIS images calibrated using this bootstrap method agree to within ~5-10% of other independently-calibrated telescopic measurements of the absolute radiance of standard regions on Mars during low dust conditions (*e.g.*, Mustard and Bell, 1994). The validity of this approach is also being tested by new sets of simultaneous VIS and HST observations during 2003. The first set of simultaneous measurements was performed on March 8, 2003; 29 sets of 5-color VIS sequences were acquired within two hours of HST multispectral imaging from 220 nm to 892 nm using HST's new Advanced Camera for Surveys (ACS) instrument; 10 of these VIS image sets were acquired at exactly the same time as the HST measurements. Initial analysis of these simultaneous observations shows that our default average-based radiance coefficient approach appears to be accurate to within ~5% for VIS bands 1 through 4, and within ~15-20% for VIS band 5 (860 nm), which has the largest scattered light component. These results are preliminary and will be augmented by additional simultaneous HST/Odyssey observing campaigns. Notably, a planned campaign during the August 2003 opposition will take advantage of the fact that HST will be able to obtain a spatial resolution of ~4 km/pixel using ACS, meaning that typical 36 m/pixel color VIS image strips will correspond to about 5x15 HST pixels, a statistically more significant number for refining the calibration coefficients.

3.4 VISIBLE IMAGER CALIBRATION TABLES AND REFERENCES

3.4.1 Table 1: VIS 8-bit to 11-bit Inverse Look-up Table

| 8 bit | 11 bit | 8 bit | 11 bit | 8 bit | 11 bit | 8 bit | 11 bit | 8 bit | 11 bit | 8 bit | 11 bit | 8 bit | 11 bit |
|-------|--------|-------|--------|-------|--------|-------|--------|-------|--------|-------|--------|-------|--------|
| 0 | 0 | 32 | 45 | 64 | 150 | 96 | 315 | 128 | 542 | 160 | 829 | 192 | 1177 |
| 1 | 1 | 33 | 47 | 65 | 154 | 97 | 321 | 129 | 550 | 161 | 839 | 193 | 1189 |
| 2 | 2 | 34 | 50 | 66 | 158 | 98 | 328 | 130 | 558 | 162 | 849 | 194 | 1201 |
| 3 | 3 | 35 | 52 | 67 | 163 | 99 | 334 | 131 | 566 | 163 | 859 | 195 | 1212 |
| 4 | 3 | 36 | 55 | 68 | 167 | 100 | 340 | 132 | 574 | 164 | 869 | 196 | 1225 |
| 5 | 4 | 37 | 57 | 69 | 171 | 101 | 346 | 133 | 582 | 165 | 879 | 197 | 1237 |
| 6 | 5 | 38 | 60 | 70 | 176 | 102 | 353 | 134 | 591 | 166 | 889 | 198 | 1249 |
| 7 | 5 | 39 | 63 | 71 | 181 | 103 | 359 | 135 | 599 | 167 | 900 | 199 | 1261 |
| 8 | 6 | 40 | 65 | 72 | 185 | 104 | 366 | 136 | 608 | 168 | 910 | 200 | 1273 |
| 9 | 7 | 41 | 68 | 73 | 190 | 105 | 373 | 137 | 616 | 169 | 920 | 201 | 1286 |
| 10 | 8 | 42 | 71 | 74 | 195 | 106 | 379 | 138 | 625 | 170 | 931 | 202 | 1298 |
| 11 | 9 | 43 | 74 | 75 | 200 | 107 | 386 | 139 | 633 | 171 | 941 | 203 | 1310 |
| 12 | 10 | 44 | 77 | 76 | 205 | 108 | 393 | 140 | 642 | 172 | 952 | 204 | 1323 |
| 13 | 11 | 45 | 80 | 77 | 210 | 109 | 400 | 141 | 651 | 173 | 963 | 205 | 1336 |
| 14 | 13 | 46 | 83 | 78 | 215 | 110 | 407 | 142 | 660 | 174 | 973 | 206 | 1348 |
| 15 | 14 | 47 | 86 | 79 | 220 | 111 | 414 | 143 | 669 | 175 | 984 | 207 | 1361 |
| 16 | 15 | 48 | 90 | 80 | 225 | 112 | 421 | 144 | 678 | 176 | 995 | 208 | 1374 |
| 17 | 17 | 49 | 93 | 81 | 230 | 113 | 428 | 145 | 687 | 177 | 1006 | 209 | 1386 |
| 18 | 18 | 50 | 96 | 82 | 235 | 114 | 435 | 146 | 696 | 178 | 1017 | 210 | 1399 |
| 19 | 20 | 51 | 100 | 83 | 241 | 115 | 442 | 147 | 705 | 179 | 1028 | 211 | 1412 |
| 20 | 21 | 52 | 103 | 84 | 246 | 116 | 449 | 148 | 714 | 180 | 1039 | 212 | 1425 |
| 21 | 23 | 53 | 107 | 85 | 251 | 117 | 457 | 149 | 723 | 181 | 1050 | 213 | 1438 |
| 22 | 25 | 54 | 110 | 86 | 257 | 118 | 464 | 150 | 732 | 182 | 1061 | 214 | 1451 |
| 23 | 26 | 55 | 114 | 87 | 262 | 119 | 472 | 151 | 742 | 183 | 1073 | 215 | 1464 |
| 24 | 28 | 56 | 118 | 88 | 268 | 120 | 479 | 152 | 751 | 184 | 1084 | 216 | 1478 |
| 25 | 30 | 57 | 121 | 89 | 274 | 121 | 487 | 153 | 761 | 185 | 1095 | 217 | 1491 |
| 26 | 32 | 58 | 125 | 90 | 279 | 122 | 494 | 154 | 770 | 186 | 1107 | 218 | 1504 |
| 27 | 34 | 59 | 129 | 91 | 285 | 123 | 502 | 155 | 780 | 187 | 1118 | 219 | 1518 |
| 28 | 36 | 60 | 133 | 92 | 291 | 124 | 510 | 156 | 789 | 188 | 1130 | 220 | 1531 |
| 29 | 38 | 61 | 137 | 93 | 297 | 125 | 518 | 157 | 799 | 189 | 1142 | 221 | 1545 |
| 30 | 40 | 62 | 141 | 94 | 303 | 126 | 526 | 158 | 809 | 190 | 1153 | 222 | 1558 |
| 31 | 43 | 63 | 145 | 95 | 309 | 127 | 534 | 159 | 819 | 191 | 1165 | 223 | 1572 |
| | | 63 | 145 | 95 | 309 | 127 | 534 | 159 | 819 | 191 | 1165 | 223 | 1572 |

3.4.2 Table 2: Bad rows and columns

| Summing Mode | Bad Columns | Bad Rows |
|--------------|--------------------|----------|
| 1 | 0 - 9, 1000 - 1023 | 0 - 1 |
| 2 | 0 - 4, 500 - 511 | 0 |
| 4 | 0 - 1, 250 - 255 | 0 |

Note that all row and column numbers are zero based. An entry of zero refers to the first sample position (column) or line (row) of the array

3.4.3 Table 3: VIS Calibration Files

| Step | 3.4.3.1 Purpose | Summing Mode | 3.4.3.2 Format | 3.4.3.3 Filename |
|------|-----------------------------|--------------|----------------------|--------------------------|
| 3 | Bias frames | 1 | 1024 x 192 x 31 FITS | zeroframe1_bias.fits |
| 3 | Bias frames | 2 | 512 x 96 x 31 FITS | zeroframe2_bias.fits |
| 3 | Bias frames | 4 | 256 x 48 x 31 FITS | zeroframe4_bias.fits |
| 4 | Smear frames | 1 | 1024 x 192 x 31 FITS | zeroframe1_zero.fits |
| 4 | Smear frames | 2 | 512 x 96 x 31 FITS | zeroframe2_zero.fits |
| 4 | Smear frames | 4 | 256 x 48 x 31 FITS | zeroframe4_zero.fits |
| 4 | Smear scaling factors | 1 | 1 line text file | dezero1_coeffs.txt |
| 4 | Smear scaling factors | 2 | 1 line text file | dezero2_coeffs.txt |
| 4 | Smear scaling factors | 4 | 1 line text file | dezero4_coeffs.txt |
| 5 | Sensitivity correction | 1 | 1024 x 192 x 5 FITS | flat_framese2.prof1.fits |
| 5 | Sensitivity correction | 2 | 512 x 96 x 5 FITS | flat_framese2.prof2.fits |
| 5 | Sensitivity correction | 4 | 256 x 48 x 5 FITS | flat_framese2.prof4.fits |
| 6 | Stray light frames | 1 | 1024 x 192 x 5 FITS | destray2_frame1_v4.fits |
| 6 | Stray light frames | 2 | 512 x 96 x 5 FITS | destray2_frame2_v4.fits |
| 6 | Stray light frames | 4 | 256 x 48 x 5 FITS | destray2_frame4_v4.fits |
| 6 | Stray light scaling factors | 1 | 5 element FITS | destray2_frame1_r.fits |
| 6 | Stray light scaling factors | 2 | 5 element FITS | destray2_frame2_r.fits |
| 6 | Stray light scaling factors | 4 | 5 element FITS | destray2_frame4_r.fits |

Calibration files can be downloaded from the Themis Virtual Volume, calib directory.

3.4.4 References

- Bell, J.F. III, T.H. McConnochie, P.R. Christensen, M.C. Malin, M. Caplinger, and M. Ravine (2003a) Calibration and initial results from the THEMIS-VIS instrument on the Mars 2001 Odyssey Mission, *in preparation*.
- Bell III, J.F., M.J. Wolff, R.V. Morris, and W.H. Farrand (2003b) Global color and mineralogic variations on the martian surface from HST/WFPC2 multispectral imaging, in preparation for *Icarus*.
- Bell III, J.F., R. V. Morris, W. H. Farrand, and M. J. Wolff (2001a). A Re-Assessment of Global Color Units on Mars from Hubble Space Telescope Visible to Near-IR Imaging and Spectroscopy. In *Lunar and Planet. Sci. XXXII*, Abstract #1484, LPI, Houston (CDROM).
- Bell III, J.F., M.J. Wolff, P.B. James, R.T. Clancy, S.W. Lee, and L.J. Martin (1997) Mars surface mineralogy from Hubble Space Telescope imaging during 1994-1995: Observations, calibration, and initial results, *J. Geophys. Res.*, 102, 9109-9123.
- Christensen, P.R., B.M. Jakosky, H.H. Kieffer, M.C. Malin, H.Y. McSween, Jr., K. Neelson, G.L. Mehall, S.H. Silverman, and M. Caplinger (2003) The Thermal Emission Imaging System (THEMIS) for the Mars 2001 Odyssey Mission, *J. Geophys. Res.*, in press.
- Morris, R.V., J.F. Bell III, M.J. Wolff, W.H. Farrand, and P.R. Christensen (2003) A Hubble Space Telescope view of Martian surface compositions from the Mars Global Surveyor Thermal Emission Spectrometer, in preparation for *Icarus*.
- Mustard J.F. and J.F. Bell III (1994) New composite reflectance spectra of Mars from 0.4 to 3.14 μm , *Geophys. Res. Lett.*, 21, 353-356.