



Contents lists available at SciVerse ScienceDirect

Icarus

journal homepage: www.elsevier.com/locate/icarus

EPOXI instrument calibration

Kenneth P. Klaasen^{a,*}, Michael A'Hearn^b, Sebastian Besse^b, Dennis Bodewits^b, Brian Carcich^c, Tony Farnham^b, Lori Feaga^b, Olivier Groussin^d, Donald Hampton^e, Marty Huisjen^f, Michael S. Kelley^b, Stephanie McLaughlin^b, Frederic Merlin^b, Silvia Protopapa^b, Jessica Sunshine^b, Peter Thomas^c, Dennis Wellnitz^b

^aJet Propulsion Laboratory, California Institute of Technology, Mail Stop 306-392, 4800 Oak Grove Dr., Pasadena, CA 91109, USA

^bDepartment of Astronomy, University of Maryland, College Park, MD 20742-2421, USA

^cCornell University, Space Sciences Building, Ithaca, NY 14853, USA

^dAix Marseille Université, CNRS, LAM (Laboratoire d'Astrophysique de Marseille) UMR 7326, 13388 Marseille, France

^eGeophysical Institute, University of Alaska Fairbanks, 903 Koyukuk Drive, Fairbanks, AK 99775-7320, USA

^fBall Aerospace and Technology Corporation, 1600 Commerce, Boulder, CO 80301, USA

ARTICLE INFO

Article history:

Received 3 December 2012

Revised 14 March 2013

Accepted 25 March 2013

Available online 25 April 2013

Keywords:

Instrumentation

Data reduction techniques

Infrared observations

ABSTRACT

NASA's EPOXI mission used the Deep Impact (DI) Flyby spacecraft to deliver a payload of three scientific instruments, two visible cameras and an IR spectrometer, to a close flyby of Comet 103P/Hartley 2 in November 2010. Interpretation of the scientific measurements made using these instruments depends on accurate calibration of the instruments' performance. Updates to the instrument calibrations achieved during the Deep Impact primary mission and results of continued monitoring of their performance during EPOXI are reported here. The instruments' performance has remained remarkably stable over the nearly 7 years of flight. Significant improvements in the understanding and calibration of the IR spectrometer response non-linearity, time-varying background level, flat field, wavelength map, and absolute spectral response have been achieved. Techniques for reducing some semi-coherent horizontal noise stripes in the visible cameras' readouts were developed, and some adjustments have been made to their absolute radiometric conversion constants. The data processing pipeline has been updated to incorporate the improvements in the instrument calibrations.

© 2013 Elsevier Inc. All rights reserved.

1. Overview

NASA's Deep Impact Flyby spacecraft carried three science instruments during its EPOXI mission to Comet Hartley 2 – two visible (VIS) cameras (the high-resolution HRIVIS and the medium-resolution MRIVIS) and an IR imaging spectrometer (Hampton et al., 2005). These instruments were also used previously on the Deep Impact mission to Comet Tempel 1. The high-resolution instrument (HRI) consists of an $f/35$ telescope with 10.5-m focal length with a spectrally split optical path feeding the filtered CCD camera and the IR spectrometer slit. The HRIVIS camera has an angular pixel size (IFOV) of 2 μ rad. The medium-resolution instrument (MRI) consists of an $f/17.5$ telescope with a 2.1-m focal length feeding a filtered CCD camera. Its IFOV is 10 μ rad. Both cameras use a Fairchild split-frame-transfer CCD with 1024×1024 active pixels. The CCD is divided into four independent 512×512 -pixel quadrants, each with a separate readout chain. The IR spectrometer is a two-prism (CaF₂ and ZnSe) imaging spectrometer

imaged on two contiguous quadrants of a Rockwell HAWAII-1R 1024×1024 -pixel HgCdTe mid-wave IR array. The array is divided into four 512×512 quadrants for readout, only two of which are read out to make a single spectrometer data frame. Its IFOV is 5 μ rad/pixel unbinned; most data are returned with 2×2 -pixel binning. For reference, Fig. 1 shows the format of a single IR spectrometer data frame illustrating its spectral and spatial coverage. An anti-saturation filter is placed over the central third of the spectrometer slit; this filter attenuates light longward of 2.7 μ m to avoid saturating a comet nucleus image with thermal emission while integrating long enough to detect the dim emission lines in the coma. The CCDs and IR FPA are read out and digitized to 14 bits by a set of dedicated instrument electronics, one set per instrument. The HRI and MRI are mounted on an instrument platform on the spacecraft, along with the spacecraft star trackers and inertial reference unit.

The calibration of the instruments for scientific purposes during the Deep Impact mission was described by Klaasen et al. (2008). Here we discuss the continuing calibration of these instruments during the EPOXI extended mission including a number of improvements and corrections to the previous calibrations.

* Corresponding author. Fax: +1 818 393 3290.

E-mail address: kenneth.p.klaasen@jpl.nasa.gov (K.P. Klaasen).

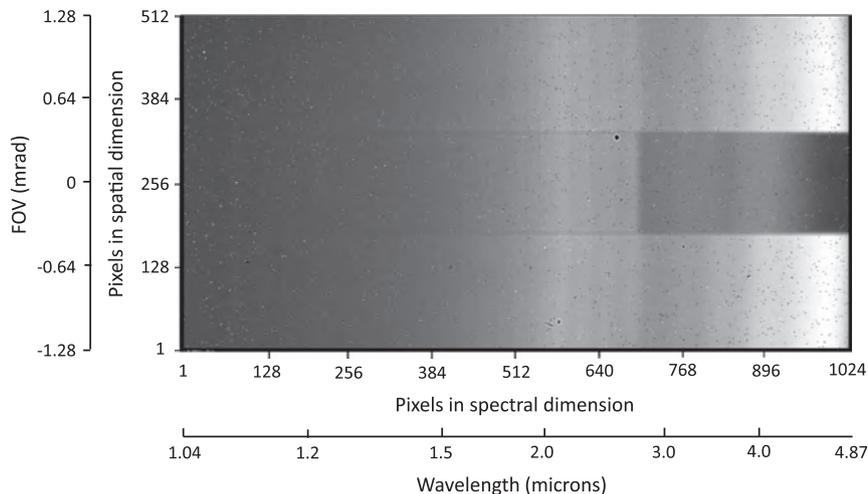


Fig. 1. The format of an IR spectrometer full data frame showing spectral and spatial dimensions; pixels are unbinned physical pixels. This raw, unprocessed frame shows the spectrum of the Moon filling the spectrometer slit. The effect of the anti-saturation filter is obvious in the central 1/3 rows. Also noticeable are scattered bad pixels and the effect of the spectral smile causing spectral lines to be slightly curved rather than strictly vertical. The spectral map shifts ~ 0.7 pixels/ $^{\circ}\text{C}$ toward higher column numbers with decreasing instrument temperature. A calibrated version of this frame is shown in Fig. 62.

First we provide a brief description of the EPOXI mission and summarize the calibration data collected during its lifetime. Instrument calibration results are then discussed with emphasis on the differences from the previously reported calibrations. The IR spectrometer calibration is addressed in the areas of geometric measurement accuracy, spatial resolution achieved, and radiometric measurement accuracy. Calibration of the VIS cameras is next addressed covering these same aspects of their performance. Finally, the data calibration processing pipeline updates for EPOXI are summarized, and differences between the Planetary Data System (PDS) delivered Version 1 (McLaughlin et al., 2011a,b,c) and Version 2 (delivery planned in 2013) data calibrations are described.

2. Mission description

The EPOXI mission is a NASA Discovery-class mission that was initiated in June 2007. The mission used the already flying Deep Impact spacecraft that had successfully completed its primary mission to Comet Tempel 1 in 2005 but remained healthy with adequate consumables onboard to continue an extended mission. EPOXI was funded in response to two scientific proposals for the extended mission – the Extrasolar Planetary Observation and Characterization (EPOCh) proposal, with Dr. Drake Deming as Principal Investigator, and the Deep Impact eXtended Investigation (DIXI), with Dr. Michael A'Hearn as Principal Investigator. These two proposals were merged into a single mission called EPOXI.

The EPOCH observations of stars known to have planetary companions plus observations of Mars and our own Earth/Moon system from a deep-space vantage point were conducted in 2008 during the cruise to the primary target for DIXI, Comet 103P/Hartley 2. Five flybys of the Earth occurred during the flight to Hartley 2, three of which were used to give gravity assists to efficiently shape the spacecraft trajectory. The DIXI observations of Hartley 2 occurred around its close flyby in November 2010 (A'Hearn et al., 2011). Instrument calibrations were conducted periodically throughout the mission, with the final calibration data set being acquired in November 2011.

Instrument calibration results specific to the EPOCH objectives have been published previously (Ballard et al., 2010; Barry et al., 2010; Christiansen et al., 2010, 2011; Livengood et al., 2011). This paper concentrates on the calibrations supporting the DIXI investigation, although calibration results from both investigations are shared between the two science teams.

3. Calibration data collection

The EPOXI instruments were exercised in a wide variety of ways during the cruise phase as well as at the Hartley 2 encounter. Besides the EPOCH observations, various instrument calibrations were conducted to monitor instrument health and performance and to better characterize certain aspects of their performance. Table 1 summarizes the data taken that are useful for instrument calibration. Calibrations included dark frames, star images, and observations of the Moon. Calibration data were acquired and returned uncompressed. The EPOCH exoplanet, Earth, Mars, and microlensing observations are not included in this table.

4. IR spectrometer calibrations

4.1. Geometric calibration

A check of the focal length of the IR spectrometer was performed by comparing positions of stars in a spatially reconstructed scan of the star cluster 47 Tuc with those from a simultaneously shuttered MRIVIS image of the cluster taken during the September 2010 calibration. The signal from columns 500–700 ($1.79\text{--}2.75\ \mu\text{m}$) was averaged in the IR data to achieve adequate SNR to locate the stars. Fig. 2 shows the reconstructed IR scan and the MRI image with the locations of a few selected stars indicated in both images. The horizontal dimension of the IR scan is derived from scanning the spectrometer slit; 50 frames are abutted here, and the scan rate is assumed to be exactly one slit width per integration time and exactly in the cross-slit direction (as commanded). Because the IR spatial geometry is only fixed in the vertical (along-slit) direction, the focal length can only be measured in flight in that direction. Comparison of the relative separations of pairs of stars located near the top and bottom of the reconstructed IR scan with those from the MRIVIS frame showed that the IR focal length is within $<0.2\%$ of the expected value compared to the MRIVIS focal length.

The relative alignment of the boresight of the IR spectrometer with respect to those of the two VIS cameras is discussed in Section 5.1.

4.2. Spatial resolution

The spatial resolution of the IR spectrometer was evaluated in the along-slit direction using cross-slit scans of the star Beta Hyi.

Table 1
Summary of instrument calibration data acquired during the EPOXI mission.

Date of acquisition	DOY(s) of downlink	Activity	Data types			Target(s)	Purpose	Comments
			HRIVIS	MRIVIS	IR			
10/4/2007	283	Instrument checkout	Images	Images	Frames	Dark sky	Test instruments	
11/4–9/2007	312–313	EPOCH photometry test	Images	Images		HD 80606 & HD 80607	Assess HRIVIS photometric accuracy	
12/4/2007	338	Instrument checkout retest			Frames	Dark sky	Test IR spectrometer functionality	
12/17/2007	351–352	HRI scattered light cal	Darks, short and long exposures			Moon	Measure scattered light in selected HRIVIS filters	
12/29/2007	363–364	Lunar calibration	Radiometry	Radiometry, scattered light	Along and cross-slit scans	Moon	1R flat field, anti-saturation filter characterization, scattered light, radiometry	Long-wavelength saturation of flat-field data
1/9/2008	9, 10	Cruise calibration	Stars, clusters, darks, STIMs	Stars, clusters, darks, STI Ms	Stars, clusters, darks	Achernar, iCar, 16CygA, Vega, Canopus, bet Hyi, HD60753, 47Tuc, NGC3114, NGC7027	Radiometry, focal length, geometric distortion, psf, VIS crosstalk, IR alignment	Good radiometry, geometry, linearity data
1/16–17/2008	19	IR darks			Darks	None	Determine IR background level	
6/23–25/2008	175–178	Cruise calibration	Stars, clusters, darks, STI Ms	Stars, clusters, darks, STI Ms	Stars, clusters, darks	Achernar, iCar, 16CygA, Vega, Canopus, bet Hyi, HD60753, 47Tuc, NGC3114, NGC7027	Radiometry, focal length, geometric distortion, psf, VIS crosstalk, IR alignment	Good radiometry, geometry, linearity data
9/18/2008	262	IR reciprocity			Darks	None	Background rate vs. exposure and detector flush history	
9/26/2008	270	IR flush test			Darks	None	Background rate vs. exposure and detector flush history	
10/2/2008	276	DI encounter rerun			Darks	None	Obtain darks applicable to the DI encounter frames	
10/8/2008		IR gap test			Darks	None	Background rate vs. mode, gap time, and flushes	
12/17/2008	352–353	HRIVIS PSF test	Star images, darks, STI Ms			Achernar, Canopus	Improved HRIVIS psf	
1/26/2009	26	IR subframe gain test		Context frames	Cross-slit scans at different modes, rates, gap times	Moon	Determine if subframe response differences are a gain or an offset effect	
6/1–2/2009	153, 156	IR flat field		Context frames	Along-slit scans	Moon	IR flat field, radiometry	Poor short-wave SNR
6/9/2009	161	IR ASF scattering		Context frames	Cross-slit scans	Moon	Radiometry, anti-sat filter scattering	
6/18/2009	169	IR linearity			Darks	Dark sky	Linearity	
10/13–24/2009	286–297	Bet Hyi IR radiometry		Context frames	Star scans	Bet Hyi	Absolute spectral radiometric response	
11/20/2009	325	IR subframe darks			Darks	None	Background signal rate in subframes vs. integration time and gap time	
12/5, 12 and 18/2009	339–352	IR flat field and radiometry		Context frames	Cross-and along-slit scans	Moon	IR flat field, radiometry	Good flat-field data
2/16/2010	47–49	Cruise calibration	Stars, clusters, darks, STI Ms	Stars, clusters, darks, STI Ms	Stars, clusters, darks	Achernar, iCar, 16CygA, Vega, Canopus, bet Hyi, HD60753, 47Tuc, NGC3114, NGC7027	Radiometry, focal length, geometric distortion, psf, VIS crosstalk, IR alignment	Good radiometry, geometry, linearity data
4/20/2010	110	IR no-save flush test			Darks	None	Test flush frames to eliminate first-frame transient in background level	
5/3–17/2010	123–138	Bet Hyi IR radiometry		Context frames	Star scans	Bet Hyi	Absolute spectral radiometric response	
7/12/2010	193	DO SI DO test		Stars		Random space	Test S/C ability to support approach observing plan	

Table 1 (continued)

Date of acquisition	DOY(s) of downlink	Activity	Data types		Target(s)	Purpose	Comments
			HRIVIS	MRIVIS			
9/3/2010	246	HRIVIS PSF test	Star images, darks, STIMs		Canopus	Improved HRIVIS psf	
9/28–29/2010	271–272	Pre-encounter calibration	Stars, clusters, darks, STIMs	Stars, clusters, darks, STIMs	Achernar, iCar, 16Cyga, Vega, Canopus, bet Hyi, HD60753, 47Tuc, NGC3114, NGC7027	Radiometry, focal length, geometric distortion, psf, VIS crosstalk, IR alignment	Good radiometry, geometry, linearity data
11/27–28/2010	331, 332, 334	Post-encounter calibration	Stars, clusters, darks, STIMs	Stars, clusters, darks, STIMs	Achernar, iCar, 16Cyga, Vega, Canopus, bet Hyi, HD60753, 47Tuc, NGC3114, NGC7027	Radiometry, focal length, geometric distortion, psf, VIS crosstalk, IR alignment	Good radiometry, geometry, linearity data
2/5–6/2011	36–63	IR extended linearity		Darks	None	Gather enough data to allow pixel-by-pixel linearity determination	
10/12/2011	287	STIM LED off filter	STIMs	STIMs	None	Test signal from STIM LED light reflected off filters	Part of CN anomaly test
10/12–13/2011	287	Warm pre-amp test	Darks	Darks	None	Check effect of elevated pre-amp temperatures on background signal level	Part of CN anomaly test
10/31–11/2/2011	307–309	Sunlight leak test		Darks	None	Measure any sunlight leaking to MRI CCD at Hartley approach attitude wrt the Sun direction	Part of CN anomaly test
11/12–23/2011	317–360	Vega IR scans		Star scans	Vega	Absolute spectral radiometric response	

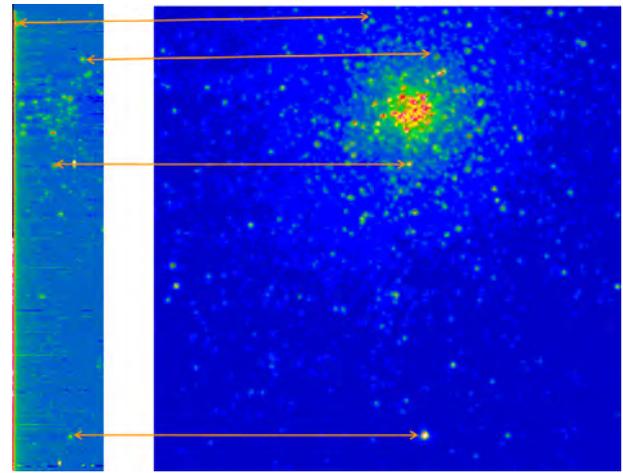


Fig. 2. Comparison of star locations in a spatially reconstructed IR spectrometer scan and a simultaneously shuttered MRIVIS frame verifies that the IR focal length is within <0.2% of the expected value relative to that of the MRIVIS. Note that the brightness distribution of stars in the cluster is dominated by a few bright red giants in the IR (~2 μm), while in the visible (~0.6 μm) the stellar brightness distribution is more uniform.

A function consisting of a Gaussian plus a quadratic polynomial was fit to the dark-subtracted star signal along each column. The full width at half maximum (FWHM) of the best-fit Gaussian was determined. This technique is the same as was used during the prime mission (Klaasen et al., 2008).

A significant degradation in spatial resolution was observed during EPOXI compared to that during the pre-encounter prime mission. Fig. 3 shows the Gaussian FWHM values vs. column number for four calibrations spanning the primary mission and EPOXI time periods. Below about column 600 (~2 μm), the FWHM has grown progressively wider with time. The DI post-encounter calibration shows the start of the degradation. By the time of the EPOXI post-encounter calibration, the shortwave FWHM had increased from about 1.7 pixels to as large as 2.7 pixels. The spatial resolution longward of ~3 μm (column 750) remains diffraction limited.

The IR spectrometer has collimating and focusing mirrors that are very sensitive to precise alignment. A shift may have been introduced around the DI encounter with Comet Tempel 1, perhaps associated with the release of the Impactor spacecraft, which may have introduced significant aberrations in the re-focused images of the slit within the spectrometer.

In addition, a hint of optical distortion has been observed in the data from 2010. This effect is most noticeable in detector Quad A (i.e., at <1.8 μm) and shows up as a wider point-spread function (PSF) on the lower-line side of the centroid. This asymmetry shows up more obviously when the PSF centroid is located in the center half of a row pixel than when it is located nearer to a pixel boundary. An example showing this effect is presented in Fig. 4.

4.3. Radiometric calibration

4.3.1. Response linearity

4.3.1.1. *Quadrant-based non-linearity correction.* The IR spectrometer response non-linearity for the prime mission was measured from recurring inflight calibration data collections during cruise to Tempel 1. Average mode-dependent linearity correction coefficients were determined for each quadrant (Klaasen et al., 2008) using only 1–2 frames of data per exposure time. The IR spectrometer response non-linearity characteristic was monitored throughout EPOXI using data from each standard calibration sequence. Detector illumination was by thermal emission from the surround-

ing instrument structure and was, therefore, nearly spatially uniform across all pixels. The data reduction technique was the same as that used during the primary mission (Klaasen et al., 2008). Fig. 5 shows the resulting third-order polynomial best-fits to the full-frame, quad-average DN/ms response rate vs. mean DN level normalized to 1.0 at a mean DN level of 5000 DN. Changes have been minimal with the various EPOXI calibration results differing by <1%.

However, ongoing science analysis has revealed that using quadrant-averaged non-linearity corrections leaves a non-uniform background in the IR spectrometer frames. In particular, near the quadrant boundary the background is about 1% higher than elsewhere, and low-level horizontal striping is visible when ratioing a short-integration frame to a long-integration frame.

Corrections in the non-linearity for the gradient across the quadrants and anomalies at the location of the quadrant boundary were not made prior to 2011 due to the limited data collected in 2005 before the Tempel 1 encounter and during cruise to Hartley 2. The quadrant approach to the non-linearity analysis was a limiting factor in the accuracy of the dark subtraction for the low-signal coma scans, at levels of 2% and, therefore, the overall fidelity of the absolute calibration and uncertainties. Fig. 6 shows an example of the residual signal gradient across a representative Hartley 2 dark image including the quadrant boundary, horizontal striping, and non-zero dark regions of the frame remaining after the quadrant linearity correction and background subtraction steps as described in Klaasen et al. (2008) were applied. The quadrant

non-linearity correction and master dark (see Section 4.3.3) used in this example were derived from dark frames collected in 2010 and were delivered to the PDS in 2011 with Version 1 of the Hartley 2 data. As can be seen by the instrumental artifacts that remain in the frame (gradient across the chip, quadrant boundary, striping, and non-zero dark regions), this treatment is not sufficiently accurate for the low-signal data acquired of Hartley 2's coma several days from closest approach. A more sophisticated linearization and dark subtraction process should remove most of these artifacts, and signal levels in the frame should evenly bound 0 DN.

4.3.1.2. Pixel-by-pixel non-linearity correction. For EPOXI, a pixel-dependent (vs. quadrant averaged) non-linearity correction was implemented to improve the calibration. A large linearity image set was taken during cruise to enable this approach. Polynomial fits were made to the relative response for each detector element separately for each binning mode. These polynomials, one for each pixel and each binning mode, were implemented as a pixel-dependent non-linearity correction. For the two sub-frame modes, an additional correction had to be applied to the linearity image set because pixels not in the sub-frame are not reset, and their saturation influences the in-frame response. This type of non-linearity treatment is used in the pipeline for processing PDS Version 2 of the Hartley 2 data. Detailed discussion of the process for deriving this correction and of the resulting calibration improvements follow.

Since the standard cruise calibration data sets contain too few frames to provide adequate signal-to-noise ratios (SNRs) for

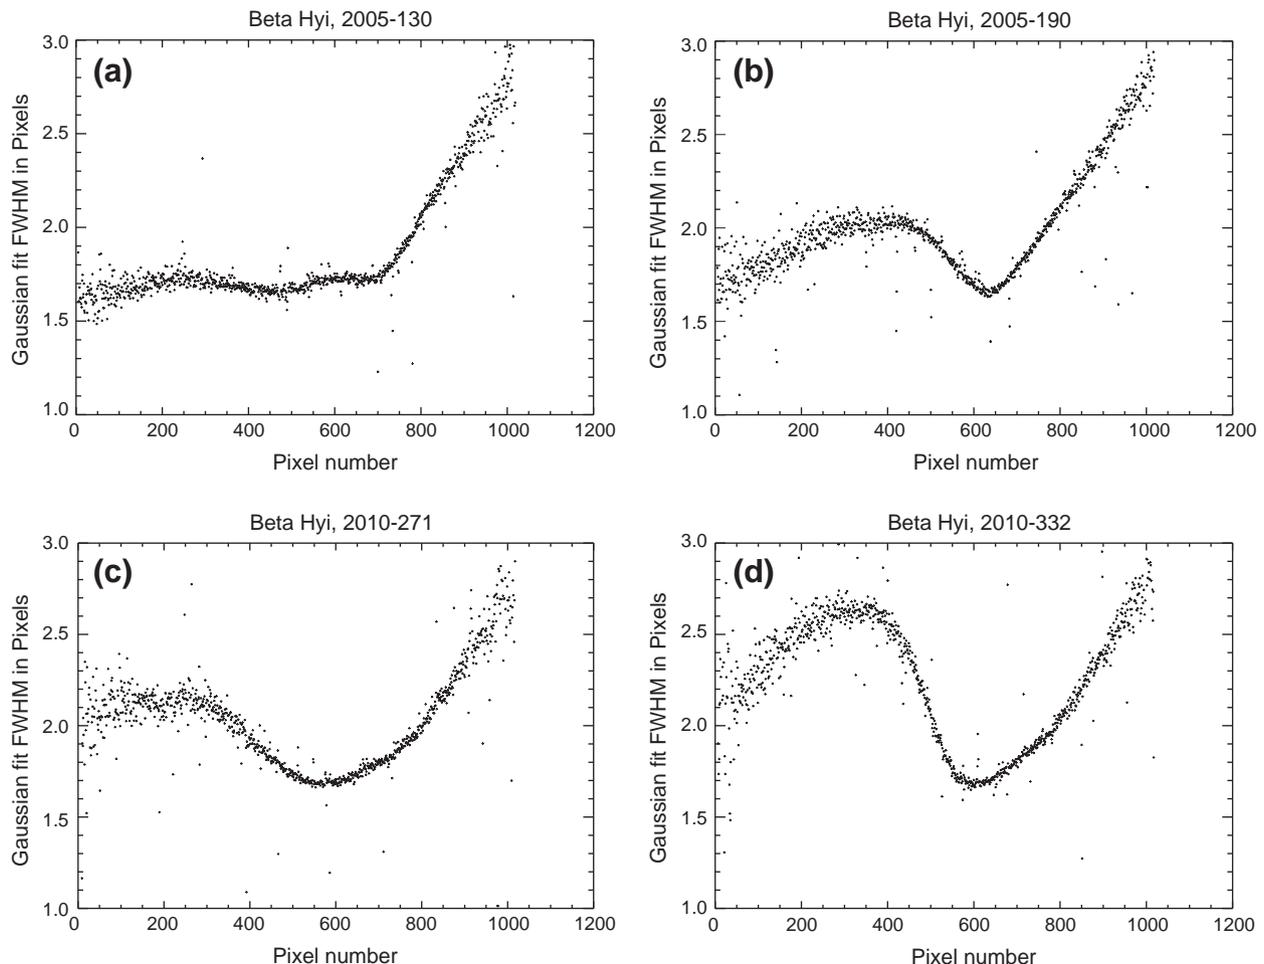


Fig. 3. FWHM of Gaussian fits to star signals along columns vs. column number from: (a) 5/10/05, (b) 7/9/05, (c) 9/28/10, and (d) 11/28/10. Spatial resolution below column 600 ($\sim 2 \mu\text{m}$) is seen to have degraded with time.

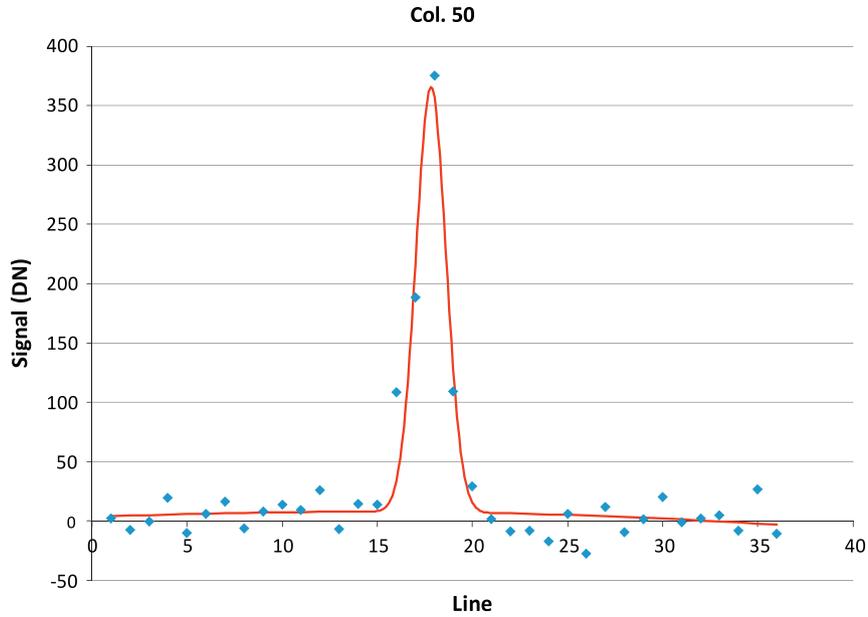


Fig. 4. Best Gaussian fit to star signal along column 50 showing asymmetry in the along-slit point-spread function.

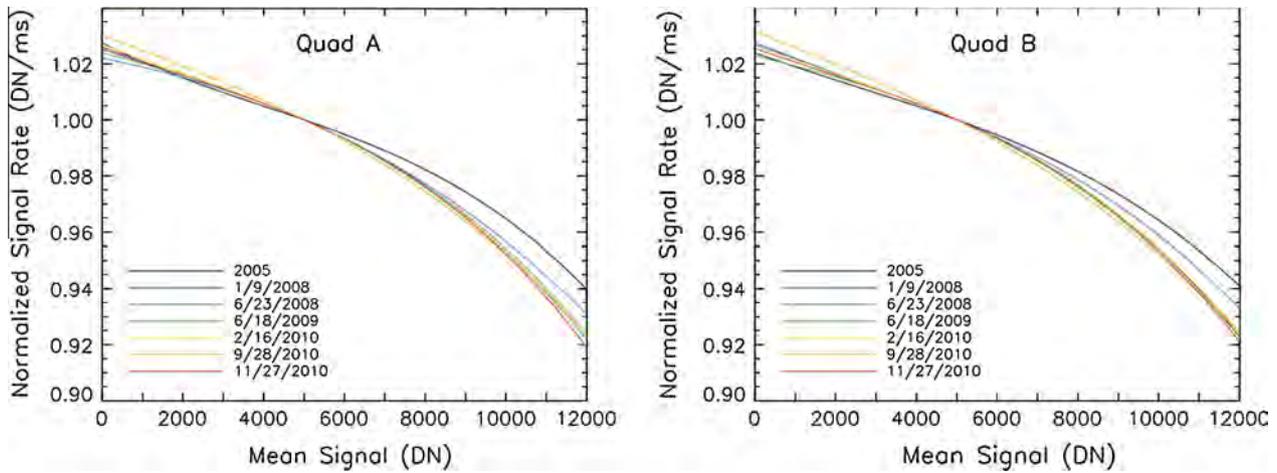


Fig. 5. Time history of 3rd-order polynomial fits to the quad-averaged IR spectrometer non-linearity response function normalized at a mean signal level of 5000 DN; changes have been minimal during EPOXI.

pixel-by-pixel linearity solutions, a special expanded IR linearity data set was acquired in February 2011. The 2011 calibration sequence included 50 frames of data at each of 12 different integration times for the unbinned imaging mode (Mode 4) and 14 integration times in each of the four binned imaging modes (Modes 1–3 and 5). The integration times varied from 721 ms to 19,000 ms and were not the same across all five modes in order to collect data at as many different signal levels as possible. With 50 data points for each integration time and mode combination, more dark frames were acquired for this non-linearity calibration test than for all other non-linearity sequences combined. Therefore, there was sufficient signal-to-noise to derive reliable pixel-by-pixel response equations, robustly define bad pixels, and create new mode-dependent master dark frames (see Sections 4.3.3 and 4.3.6).

In order to properly calculate the response equations for good pixels, certain data points were omitted from the fits. To remove the chances of including a saturated pixel in the longest exposures, integration times >16,200 ms were not included in the fits. In addition, the first 10 frames (of 50) for each integration time were not

included because of the known dark-level instability and decay, where the IR dark-frame level observed in a series of successive frames decreases significantly between the first and second frame and more gradually for each frame thereafter, as discussed in Klaasen et al. (2008) and Section 4.3.3.5.

The response of a representative subset of good pixels in all modes is shown in Fig. 7. The response is close to linear with integration time, but measurable departures from linearity are observed. Although the data were collected at different integration times for each operational mode, a vertical slice through the plot at a single integration time shows that the sub-frame modes consistently demonstrate higher response than either full-frame mode, whereas the three full-frame modes, binned, unbinned and alternating binned, show about the same response.

The two sub-frame modes (BINSF1 and BINSF2) use only 50% or 25% of the total pixels available, and the out-of-frame pixels are not reset in these imaging modes. In the sub-frame modes, the out-of-frame pixels have saturated, and the built-up charge from these pixels induces an increased response rate in the sub-frame

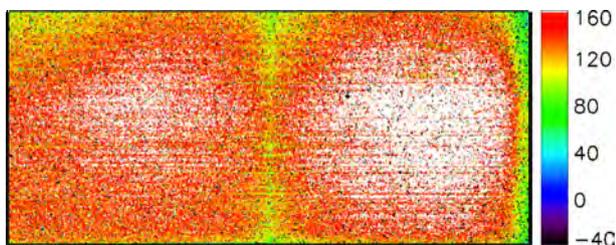


Fig. 6. Representative dark data frame from the Hartley 2 observation sequence on October 17, 2010 (E-34d, DOY 290, hi0340609368_4000001_017.fit) with corrections using quadrant-based non-linearity correction coefficients and temperature-dependent functional scaling of the master dark (see Section 4.3.3) but failing to remove a number of instrumental artifacts. The color bar stretch is linear from -40 to 160 DN.

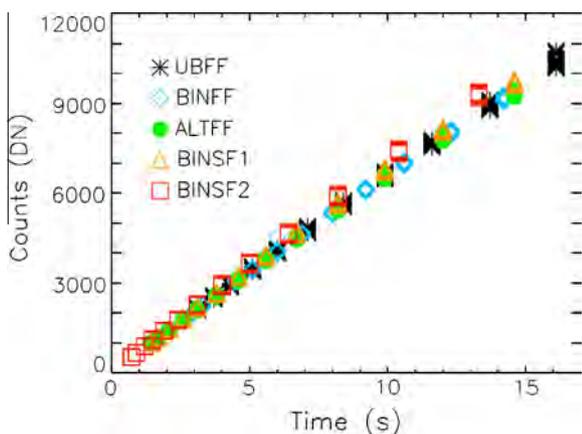


Fig. 7. The raw data number response of a representative subset of good pixels (same region on the detector in all modes) vs. detector integration time. Imaging modes are identified by color and symbol.

pixels (Section 4.3.3.6), as illustrated in Fig. 8A. In order to remove this effect, linear fits to the measured rates for the good pixels from frames with exposures between 4 and 10 sec were derived for each mode, and signal-level-rate ratios of the full-frame to sub-frame modes were calculated. The derived ratios are 0.97 for BINSF1 and 0.93 for BINSF2 for this data set. The difference between these ratios is consistent with BINSF2 having 50% more out-of-frame unreset pixels than BINSF1. Once these factors are applied to the data, by multiplying the sub-frame signal by the corresponding ratio, the linearity curves agree well for all the modes (Fig. 8B).

The data were next normalized to 1 DN/ms at the exposure time per mode closest to 6000 ms. This exposure time was chosen for normalization because it has moderate signal levels (~ 4000 DN) and is neither the first (i.e. shortest) nor the last (i.e. longest) exposure time for any mode of the linearity dark sequence. The good data points from each pixel were then fit with fourth-order polynomials. Fig. 9 shows the subset of BINFF and BINSF2 data from Figs. 7 and 8 fit with fourth-order polynomials to show the general non-linearity trends. Comparison between Figs. 7 and 10 shows that once the linearization correction is applied to the data, all modes at all exposure times exhibit consistent response behavior.

Using this data set, pixel-by-pixel response linearity functions were derived for all good pixels in the array (the bad pixels are flagged as such, see Section 4.3.6). Using pixel-by-pixel linearity corrections eliminated the spatial non-uniformities in the resulting background image as shown in Fig. 11. The improved calibration achieved by using pixel-by-pixel linearity corrections was propagated throughout the derivation of the calibration files to be discussed in upcoming sections, such as the master dark file, the

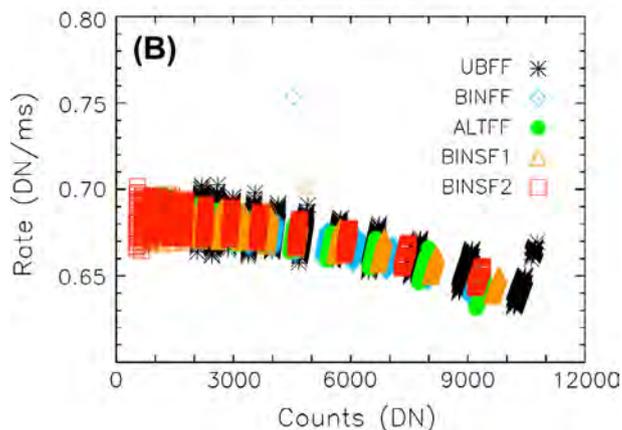
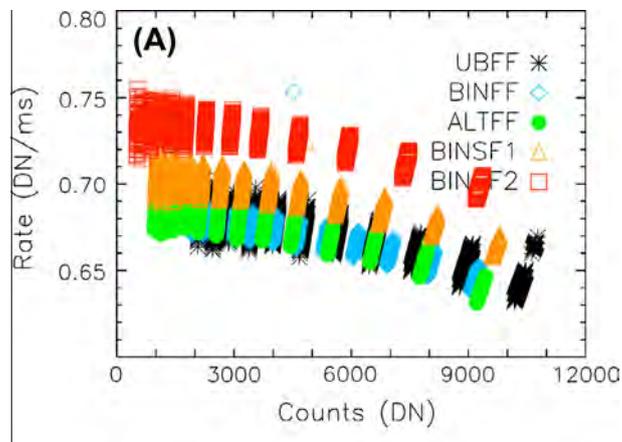


Fig. 8. The IR spectrometer response rate is calculated using the subset of data shown in Fig. 7. Plot A shows the sub-frame imaging modes with a higher response rate. Plot B was created by applying the correction ratios of 0.97 and 0.93 for BINSF1 and BINSF2, respectively. As in Fig. 7, imaging modes are identified by color and symbol. The non-linearity in the IR response becomes more apparent in these curves.

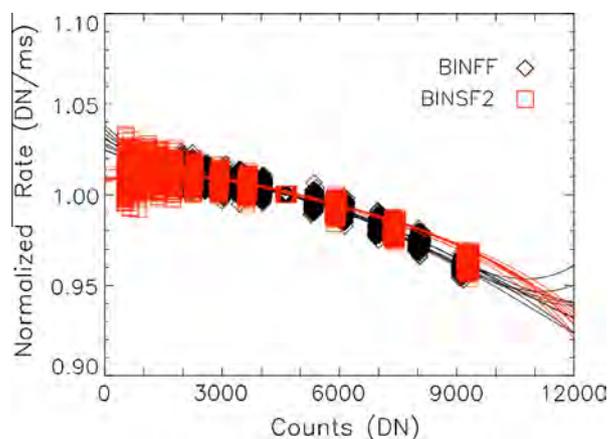


Fig. 9. The IR spectrometer non-linearity equations are derived after all the data are normalized to 1 DN/ms at the exposure time per mode closest to 6000 ms and the sub-frame scaling ratios are applied. The binned full-frame (BINFF) and the smallest binned sub-frame (BINSF2) mode data plotted in Figs. 7 and 8 are shown here as representative fits.

absolute radiometric calibration curve, bad pixel maps, the flat-field correction file, and the anti-saturation filter spectral transmission curve.

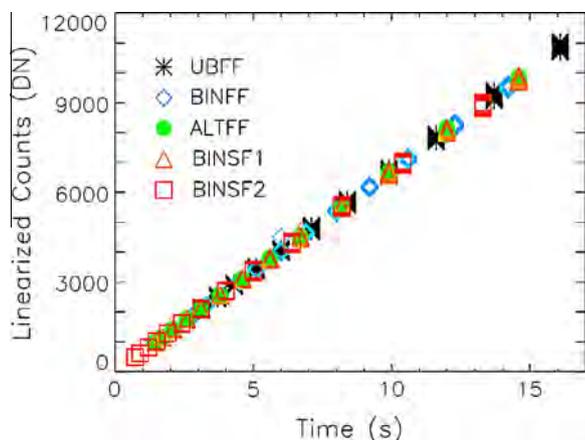


Fig. 10. The response of a representative subset of good pixels (same region on the detector in all modes) after applying the linearity correction vs. raw signal level. Imaging modes are identified by color and symbol.

During DI and for EPOXI PDS Version 1, the pipeline included the gain adjustment as part of the linearity correction when using the subframe modes (Klaasen et al., 2008). As discussed in Section 4.3.3.6 below, the difference in background signal generation rate seen in the subframe modes is now understood not to be a gain effect. Therefore, for EPOXI the subframe gain adjustment has been eliminated from the linearity correction step in the processing pipeline, and the extra background signal seen in the subframe modes is corrected for as part of the dark frame subtraction step (Section 4.3.3.6).

4.3.2. Gain/full well

The IR spectrometer system gain (i.e., the conversion factor from signal electrons to output DN [data number]) and detector full-well capacity (i.e., the maximum signal for which the detector response remains nearly linear with integration time) were determined using the photon transfer technique (Janesick et al., 1987). As for the linearity measurements, detector illumination was by thermal emission from the surrounding instrument structure.

Measurements were made with data from each standard calibration activity (1/08, 6/08, 2/10, 9/10, and 11/10). Processing was as described by Klaasen et al. (2008). The first frame taken in each data set was excluded due to the transient first-frame effect discussed in Section 4.3.3.5. Data were linearized prior to further processing using quad-average corrections. Fig. 12 shows an example of a typical photon transfer curve for the spectrometer. In the unbinned mode, the gain remained between 15.5 and 16 e^-/DN . No significant differences were observed between the two detector quadrants. In the binned modes, the gain remained between 58

and 60 e^-/DN . The gains are consistent with those measured during the prime mission. No systematic changes were observed in the gain over the course of the EPOXI mission.

The full-well levels also remained the same as measured during the prime mission – 8000 DN for saturation of a few pixels and 11,000 DN for the start of saturation of the majority of pixels.

4.3.3. Zero-exposure background

4.3.3.1. Background signal sources. The signal in the IR spectrometer dark frames is a combination of the detector dark current, the background thermal IR radiation emitted by components on the spectral imaging module (SIM) bench, and a dark-sky background. As the HRI-IR does not have a shutter, darks are acquired by looking at a patch of dark sky that is assumed to be flat across the field of view. The zero-exposure background signal is dominated by detection of thermal emission photons from the instrument SIM bench structure surrounding the detector. The detector continuously integrates signal between successive detector resets. Resets are followed by detector reads after a delay that defines the integration time for that frame.

The spatial distribution of background signal due to thermal emission from the structure has been characterized by a master dark frame representing the linearized DN/ms signal rate at each pixel across the array for a given reference instrument temperature (as represented by the prism temperature). Fig. 13 shows the dependence of the response rate of the detector on the temperature of the optical bench.

4.3.3.2. Detector readout timing. Because the detector continuously integrates signal between successive detector resets, the background signal level is critically dependent on the detailed internal timing between resets and reads. During EPOXI we acquired additional calibration data that allowed refinement in our knowledge of the instrument timing patterns. Updated timing diagrams are presented in Fig. 14. The primary differences from those presented in Klaasen et al. (2008) involve the minimum delay between ending the execution of one command and initiating execution of the subsequent command and the timing and order of reset/reads in the alternating modes.

4.3.3.3. Background spatial stability/uniformity. The stability of the spatial pattern in the IR background was examined using dark frames taken throughout the EPOXI mission to see if a single master dark frame would be applicable at all times in each IR operating mode. Linearized (using quadrant-wide average coefficients) dark frames (excluding any that had obvious externally generated signal) were divided by the master dark derived during the DI prime mission for the same mode and contrast enhanced over the range of 0.99–1.01 times the median value. Frames containing first-frame effects or unsaturated non-reset regions for the subframes modes

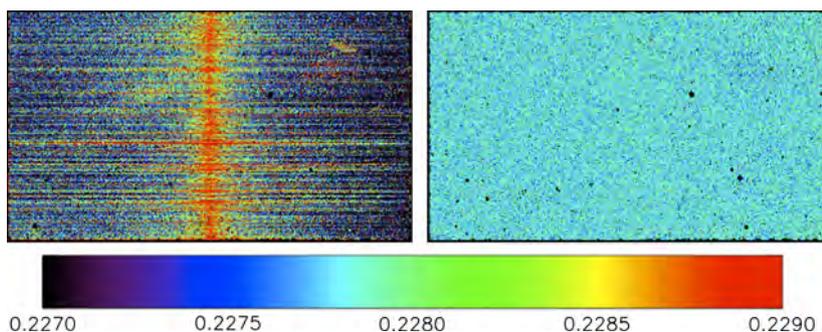


Fig. 11. Ratio of a 3.12-s dark frame divided by a 13.7-s dark frame after linearity correction using a quadrant-average function (left) and using pixel-by-pixel functions (right). The pixel-by-pixel linearity correction results in a much more uniform background at all signal levels.

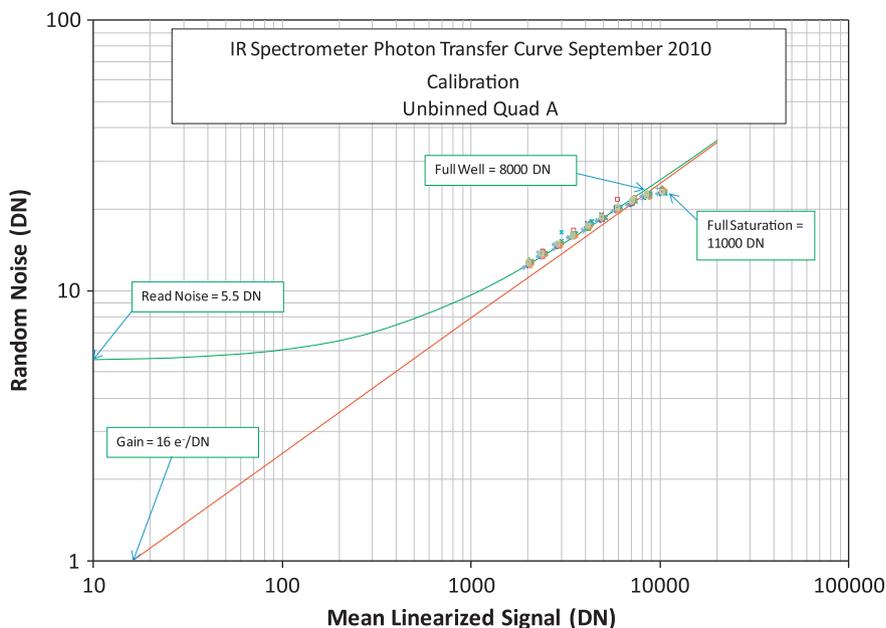


Fig. 12. Representative IR spectrometer photon transfer curve.

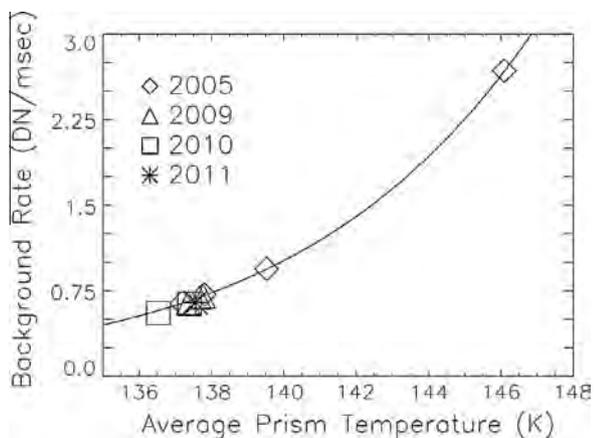


Fig. 13. IR spectrometer background rate as a function of the prism temperature for the inflight calibration data surrounding the Hartley 2 encounter. Standard calibrations were acquired on September 28, 2010 and November 27, 2010, while the extended linearity calibration was performed on February 5–6, 2011. For comparison, the Hartley 2 closest approach data from November 4, 2010, were acquired at an optical bench temperature of 137.06 K. Also shown are background rates observed during the prime mission in 2005 and a lunar calibration in 2009. Symbols indicate the measured values of the total signal on different dates, while the line indicates the computed best-fit modeled contribution from the instrument thermal emission previously defined for the detector (Klaasen et al., 2008).

were excluded. A total of 74 images acquired between December 29, 2007 and September 28, 2010 were analyzed.

This analysis showed that the IR dark background pattern during EPOXI is measurably different than it was at the end of the DI mission, as seen in Fig. 15. All modes show enhanced signal levels at the longwave end of the detector (~5%) and a slight enhancement in the high-line, low-sample corner of full-frame modes (~2%). The relative signal level is slightly reduced in the center of each quadrant (~1%). No systematic changes in the background pattern within a quadrant were seen during the EPOXI mission. However, some change in the balance between the two quadrants was observed in all the modes, with the most pronounced change being seen starting with data taken on 6/18/09. These changes are at the <1% level and are most likely due to small errors in the de-

rived linearization factors rather than real changes in the detector or its illumination. As discussed in Section 4.3.1, linearization was performed on a quadrant-average basis for this analysis, and the linearization coefficients were updated with each cruise calibration, including that of 6/18/09.

Although the signal-to-noise ratio (SNR) levels are significantly different for the two exposure levels in each mode shown in Fig. 15, there is a suggestion that the dark pattern does vary somewhat with signal level. To investigate this possibility further, ratios were taken between the long and short exposures in each mode for the frames in Fig. 15. The results are shown in Fig. 16, contrast-enhanced from 0.99 to 1.01 times the median value. A systematic dependence of the linearized dark pattern is seen with exposure time and/or signal level. The columns near the quadrant boundary have a relatively lower signal level than other columns as the signal level and/or exposure time increases. The effect is quite small (of order 1%) and is consistent with the results discussed previously in Section 4.3.1. This effect is apparently due to these columns having a different linearization function than the others in their quadrant.

The initial analysis of dark background pattern stability in the ALTF mode showed a distinctive vertical gradient in the ratios to the BINFF master dark (which was thought to be applicable to ALTF as well, since they have the same frame format). Additional research into the IR Spectrometer internal frame timing revealed that in the alternating readout modes (ALTF and DIAG) the reset frames and the read frames do not take the same amount of time to complete (as they do in the interleaved modes). As shown in Fig. 14, the resulting total integration time decreases from $1.456s + \text{INTEG_DELAY}$ for the first line to $1.410s + \text{INTEG_DELAY}$ for the last line in ALTF. The need for a separate ALTF master dark thus became apparent, and the need for using a line-dependent integration time in the algorithm for radiometric calibration of ALTF frames in the data processing pipeline was identified and subsequently implemented.

Fig. 17 shows the dark background pattern for ALTF after dividing by a BINFF master dark that had been scaled by line-dependent factors representing the actual integration time for each line divided by the average integration time. Use of this modified master dark results in a dark background pattern that matches

closely that of the BINFF mode (see Fig. 15) verifying that the line-dependent ALTFB integration time formula is correct. Fig. 18 presents the ratio between a high- and a low-signal ALTFB frame pair; the variation in the dark background pattern with signal level is less than is seen for the other modes (see Fig. 16).

4.3.3.4. Master dark frame update. For EPOXI, the DI prime mission master dark frame was initially replaced by an updated frame with improved SNR derived from a large set of 2187 unbinned full-frame (UBFF) dark-sky frames acquired while scanning across the star β Hyi in October 2009. The data were first linearized using quad-averaged corrections, and then a resistant mean for each pixel was computed over the set of frames to produce the master dark. This master dark, appropriately binned and cropped for all other modes, was used in producing Version 1 of the PDS IR data archive (McLaughlin et al., 2011a,b,c).

With the abundance of dark frames collected in each operational mode during the 2011 pixel-by-pixel non-linearity calibration, we were able to create a master dark for each mode independently to better characterize and remove the mode-dependent dark-level structure within a given frame. The new

mode-dependent master darks were created using the last 40 frames (of 50) of each exposure time <16,200 ms from the February 2011 linearity calibration data (Table 2). As in the linearity analysis, selecting only these frames avoids most saturation and first-frame signal-level-decay effects. The pixel-by-pixel non-linear response correction was applied to each frame prior to creating the master dark.

In order to incorporate the measurements from different exposure times into a single master dark, the frames were normalized to a standard and well-behaved location on the detector. Experiments showed that the relative background level across the chip is very stable, at least near the nominal operating conditions. Although there are more- and less-sensitive regions on the array, relative contrast across the array is constant with exposure time (up to the saturation point). Therefore, a sub-frame region can be used to represent the background level changes in the analyses. For the full-frame, unbinned mode (Mode 4), the region defined by rows 191–319 (inside the anti-saturation filter) and columns 192–320 (wavelengths from 1.2 to 1.37 μm) was chosen to provide the representative measure of the background, and it was assumed that region would reflect proportional changes across the entire

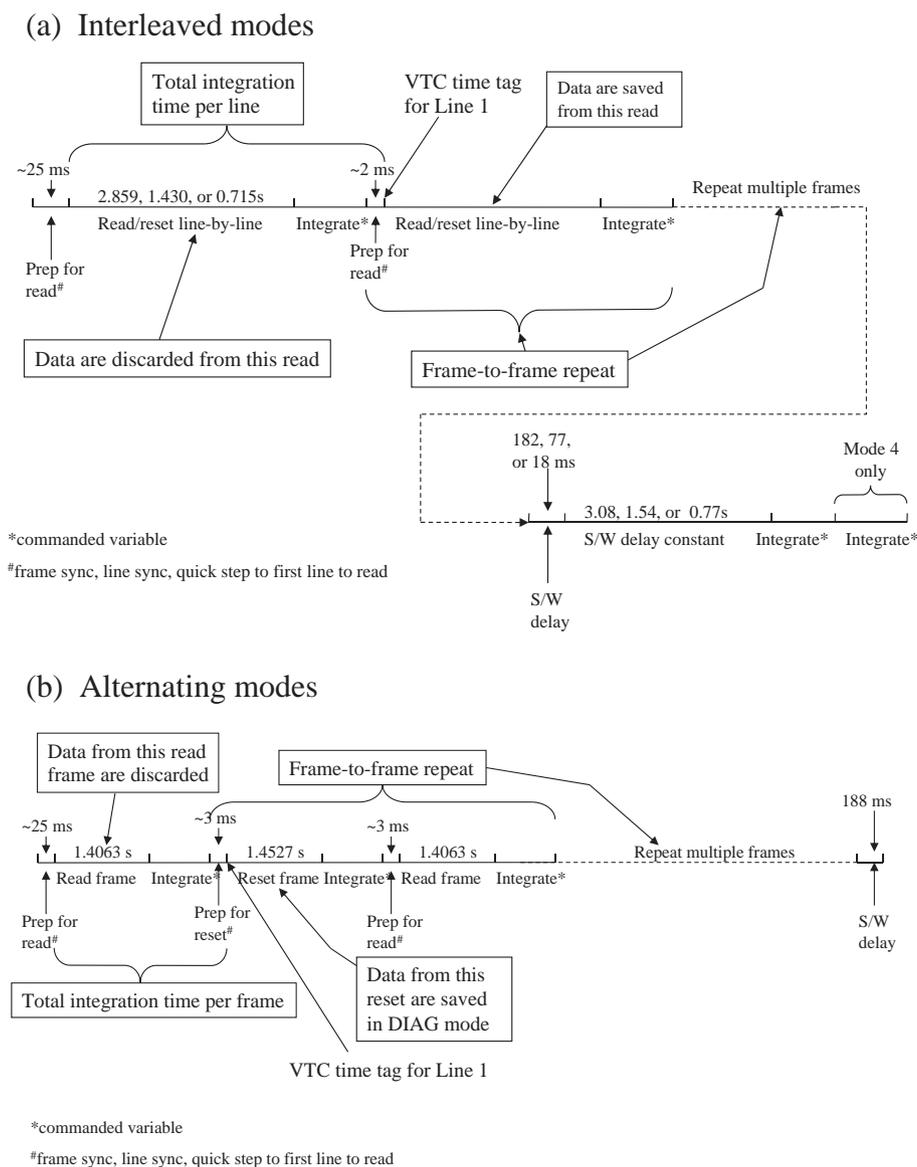


Fig. 14. IR spectrometer focal plane timing for both readout modes, (a) interleaved (Modes 1–4) and (b) alternating (Modes 5 and 6).

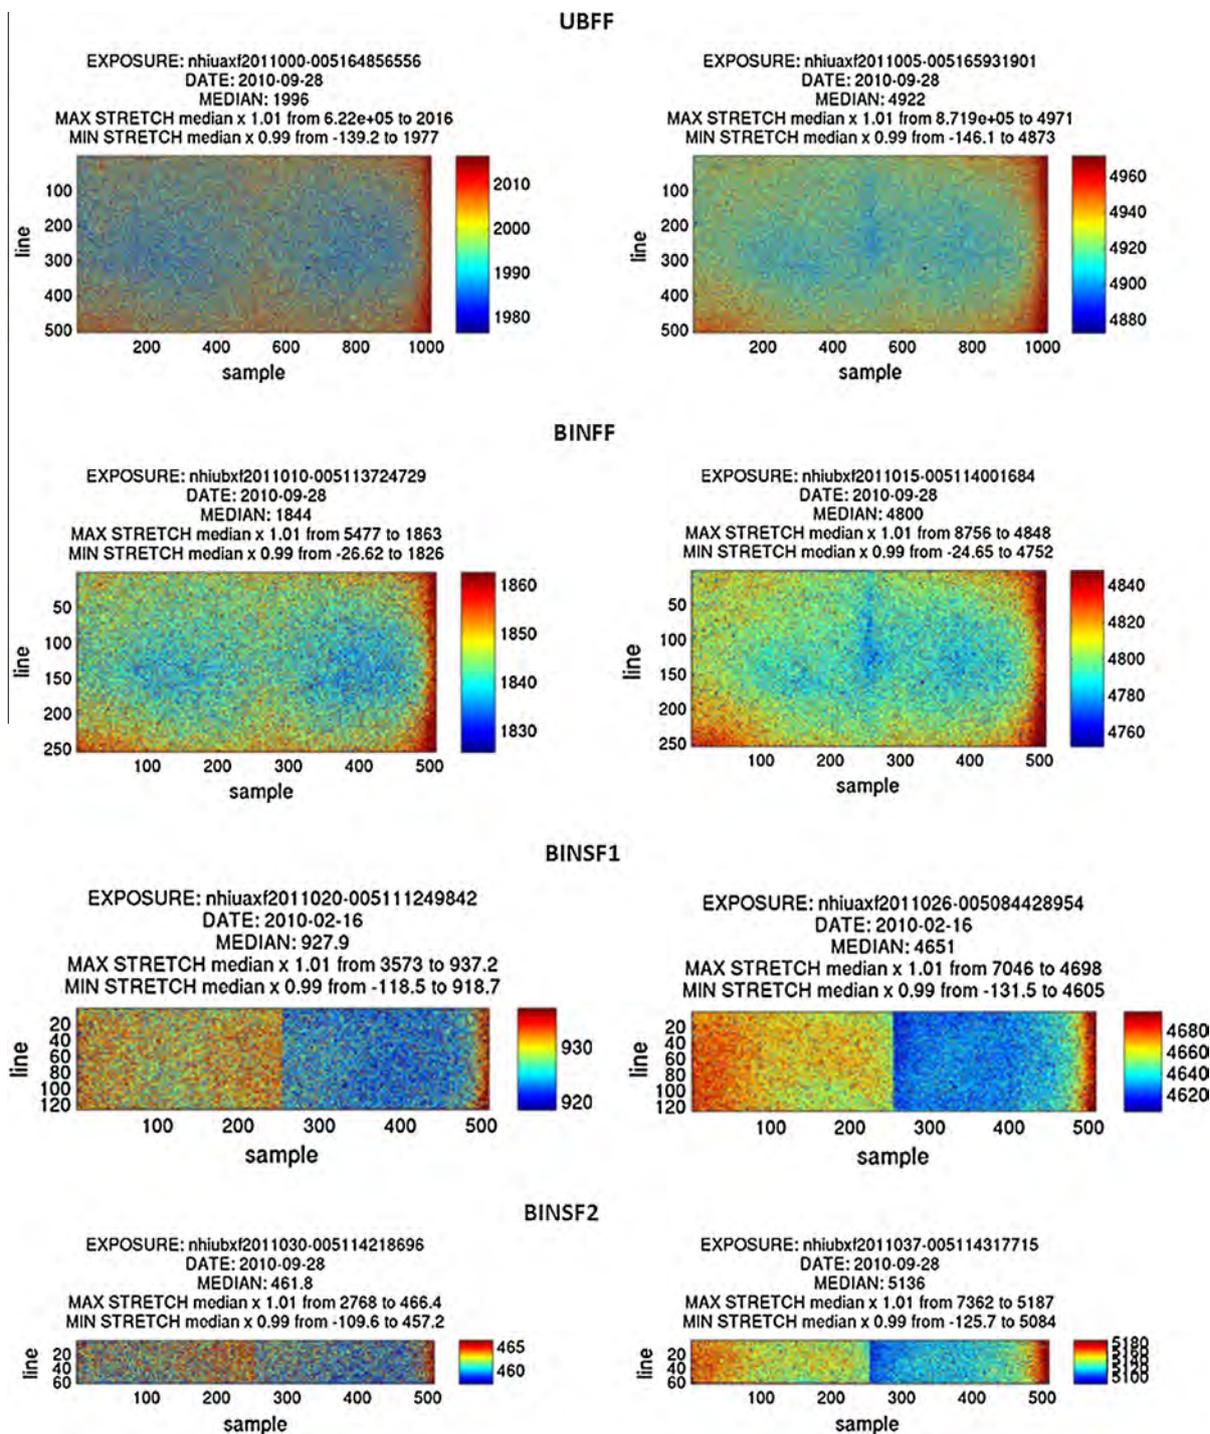


Fig. 15. Examples of ratios of EPOXI IR spectrometer dark frames to the DI master dark in each mode at two different integration times. Contrast stretch is from 0.99 to 1.01 times the median value.

field of view. (Note: a row contains one full spectrum and a column contains spatially adjacent pixels within the slit at a single wavelength.) This region was selected because it has a good signal with little gradient, and the same pixels can be used in all modes (though they are binned in the other modes). The region also avoids the boundaries of the anti-saturation filter rows.

A resistant mean algorithm was applied to the selected region to remove any pixels that differ by more than 2.5σ from the mean, and the average of the remaining pixels was used to provide a single value that would represent the background level for the frame. The frame was then divided by this average background level to

derive a normalized frame. A single master dark was created for each mode by calculating the resistant mean of each pixel in all of the normalized frames for a given mode (Fig. 19). Note that for the ALTF mode (Mode 5), the signal rate rather than signal was used to create the master dark frame in order to first remove the timing gradient, and hence signal gradient, across lines due to the readout algorithm of that mode. These new master darks will be used in creating Version 2 of the PDS IR data archive.

4.3.3.5. Residual signal trapping. As discussed by Klaasen et al. (2008), the IR background level also varies as a function of the

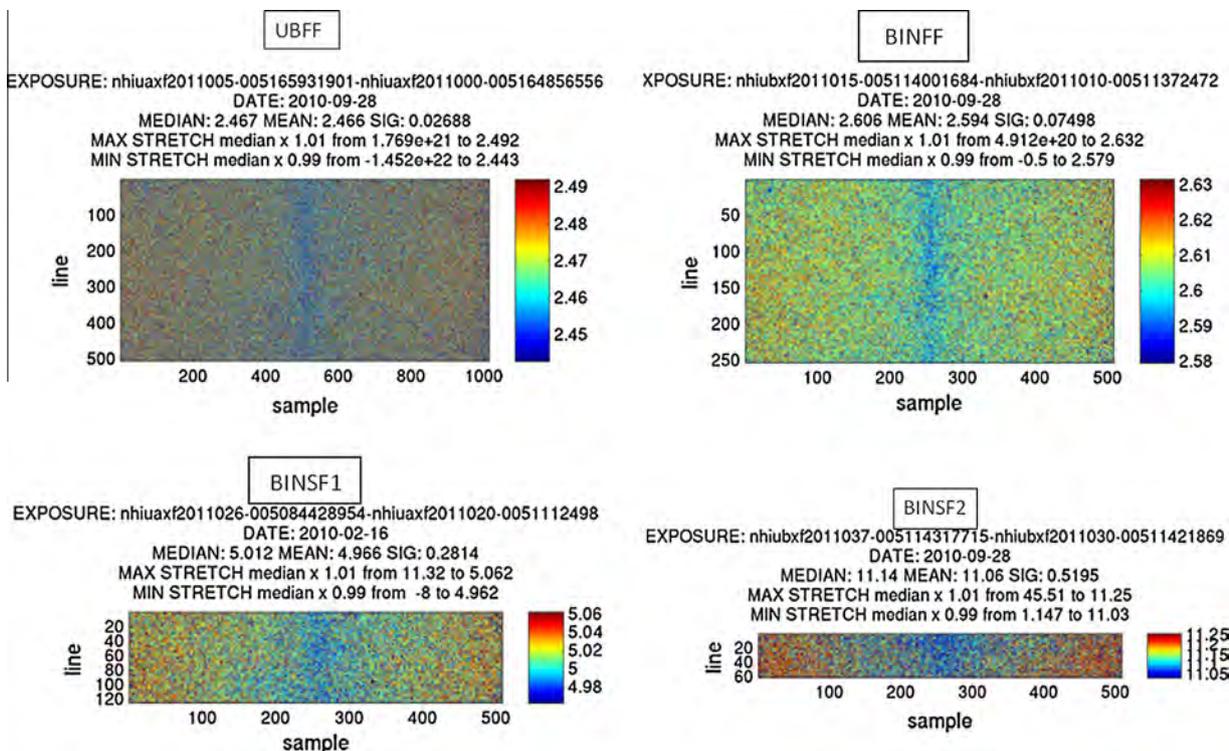


Fig. 16. Ratio of EPOXI IR spectrometer high-/low-signal dark frames from Fig. 15, contrast stretched from 0.99 to 1.01 times the median value.

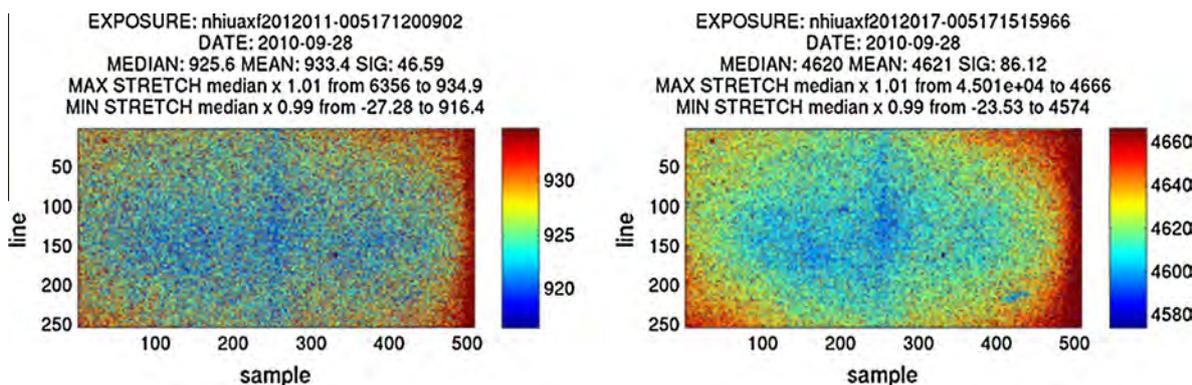


Fig. 17. Examples of ratios of EPOXI IR spectrometer ALTF dark frames to a master dark modified to reflect the line-dependent integration time in this mode; two different integration times are shown. Contrast stretch is from 0.99 to 1.01 times the median value.

history of resets and reads preceding a given frame. A reset frame was seen to not fully clear the previously integrated signal; therefore, the first frame in a sequence of successive IR frames with identical integration times usually has a measurably different background level than those of subsequent frames. We have now empirically derived a model of these transient effects by assuming that they are due to signal electron trap states within each pixel well of the IR detector (Smith et al., 2008a,b). The transient effect is predicted by the amount of charge that is transferred between the trap states and the photocharge well during an exposure frame. We assume five different trap states exist each with different exponential time constants for capturing and releasing charge. The steady-state charge held in each trap is assumed to be proportional to the steady-state photocharge in the well.

The transient background model was derived by fitting binned full-frame data from the Hartley 2 approach and departure sequences and from the extended linearity calibration data from Feb-

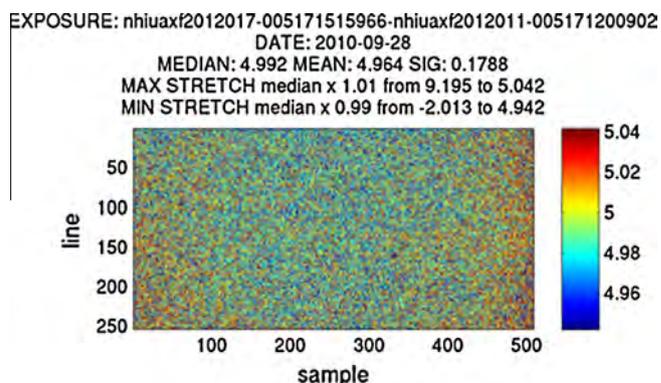


Fig. 18. Ratio of EPOXI IR spectrometer ALTF high-/low-signal dark frames from Fig. 17, contrast stretched from 0.99 to 1.01 times the median value.

ruary 2011 using a standard reference area near the center of Quad A. The data were all fit using the same trap proportionality constants and time constants. The resulting model parameter fits are given in Table 3. Fig. 20 shows some examples of the goodness of the model fit. The fits are generally good to a few DN, usually <1 DN after the first frame of a sequence. For this model to work properly, a complete history of detector resets and reads is required as an input to the pipeline. Corrections for the transient background have not been applied in the standard EPOXI calibration pipeline; the model is available as a separate user option (Section 6.3).

4.3.3.6. Subframe glow. An additional effect on the background signal is seen in the IR spectrometer subframe modes. As was observed in the linearity data reported by Klaasen et al. (2008) and in Section 4.3.1, the background generation rates in the subframe modes were higher than that in the full-frame modes by a factor of about 1.04 for the BINSF1 (128-line) mode and 1.08 for the BINSF2 (64-line) mode. The two subframe modes use only 50% or 25% of the total detector lines available, and the out-of-frame top and bottom border pixels are not reset in these modes. Therefore, by about 20s after the last full-frame reset, these border pixels have become saturated with charge. For DI, this unreset charge in the border areas was suspected to be the cause of the increased background signal rates in the subframe modes, but no explanation for its mechanism had yet been derived. We now offer a probable explanation and an improved model for correcting for this excess charge.

We believe that the excess background charge produced in the subframe modes when the border regions have remained unreset for ≥ 20 s is due to photon emission glow from the saturated border pixels, which have become forward biased due to their saturation. Photons are internally reflected in the CdZnTe substrate of the detector, and the illumination spreads from each glowing pixel in

a pattern approximately proportional to $\cos(r)$, where r = the distance from the glowing pixel. Several pieces of evidence support this conclusion. First, the background signal generation rate increase in the subframe areas is directly correlated with the amount of built-up charge in the unreset regions. As Fig. 21 shows, the excess background charge does not exist when the unreset signal charge is less than $\sim 11,000$ DN, gradually increases until the unreset charge reaches $\sim 20,000$ DN, then remains constant beyond that point. We interpret this result as being due to pixels saturating at various charge levels between 11,000 DN and 20,000 DN, with some pixels starting to go forward biased for charges as low as 11,000 DN and all pixels going forward biased for charge levels above 20,000 DN. This saturation level is consistent with those observed in the linearity data and the photon transfer data. Second, the background signal rates in the subframe areas are highest nearest the boundaries with the unreset regions, as shown in Fig. 22. The vertical profile of the glow signal is fit well by a quadratic function of line number in each column. The ratio of the subframe background at the center line of the detector to that of a similar BINFF frame is the same as the linearity signal ratios seen in each subframe mode. And third, the glow signal accumulates at a constant rate over the commanded integration time.

A correction for the excess background charge in the subframes due to unreset pixel glow can be implemented in the pipeline as part of the dark frame subtraction step. However, this subtraction will only be correct when the unreset pixels are all fully saturated. Correction for the glow effect is not properly modeled when the unreset pixel charge is $< 11,000$ DN or when partial saturation of the unreset pixels occurs due to mean charge levels between 11,000 and 20,000 DN or to non-uniform scenes that extend into the unreset region causing spatially uneven saturation. Note that the excess fully-saturated background glow signal in the subframe modes is proportional to the integration time and the inherent glow rate. The excess charge is not a function of either the

Table 2
February 2011 frames used for IR spectrometer master dark creation.

Mode	# of non-saturated exposure times	# of well-behaved frames	Total frames used in master dark
BINFF (1)	12	40	480
BINSF1 (2)	13	40	520
BINSF2 (3)	13	40	520
UBFF (4)	11	40	440
ALTF (5)	13	40	520

Table 3
Background signal transient model parameter values. A = proportionality constants, TauE = time constants for emptying traps, TauF = time constants for filling traps.

	Long	Med. lng	Med	Med. sh.	Short
A	0.0250	0.003	0.002	0.003	0.0051
TauE (s)	208.3	71.3	25.4	4.6	0.50
TauF (s)	208.3	71.3	25.4	4.6	0.50

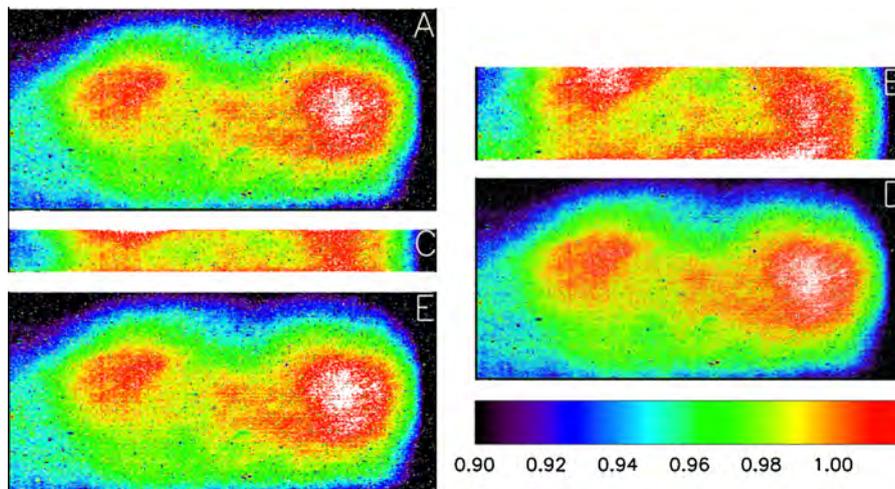


Fig. 19. Normalized mode-dependent master dark frames. (A) BINFF, (B) BINSF1, (C) BINSF2, (D) UBFF, and (E) ALTF. An identical stretch from 0.9 to 1.02 is applied to all of the master darks displayed.

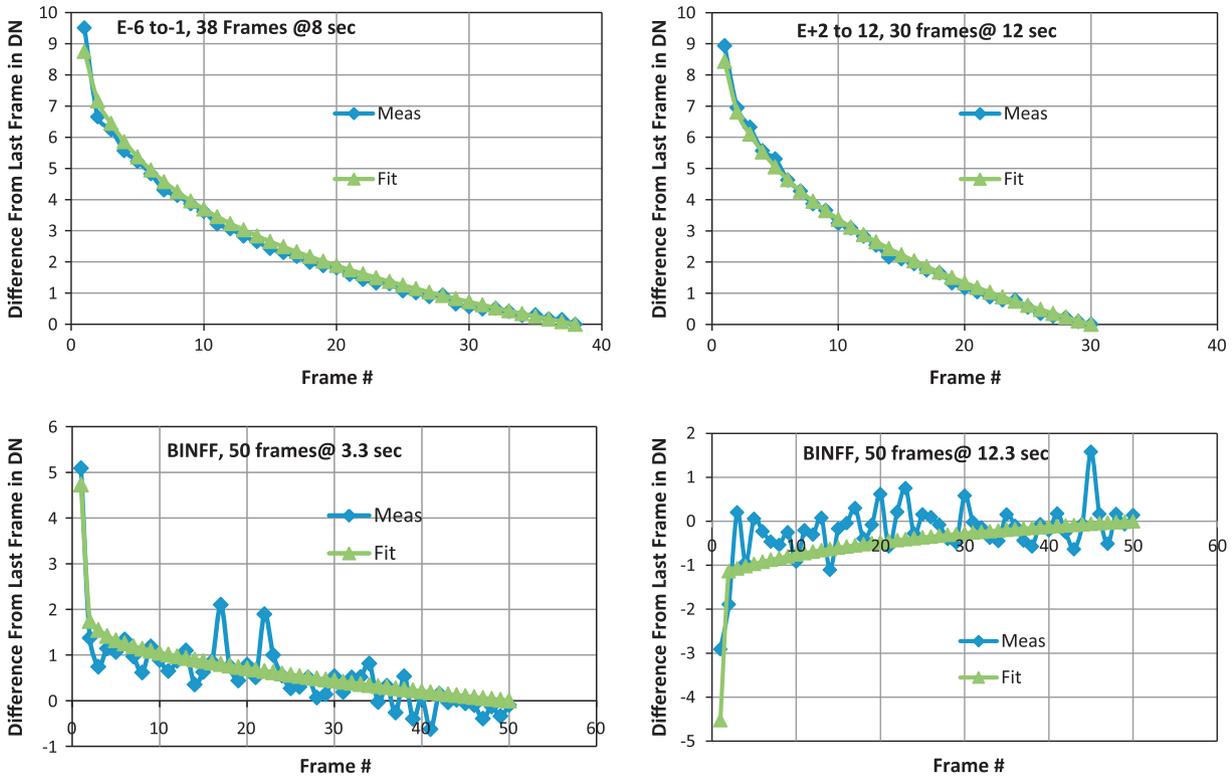


Fig. 20. Examples of background transient model fits to data. Top panels are for BINFF data taken immediately following four flush frames using the same integration times, beginning with a fully saturated detector. Bottom-left panel data were preceded by 50 frames with a 2.88-s integration time and a 1-s gap. Bottom-right panel data were preceded by 50 frames with a 10.6-s integration time and a 1-s gap.

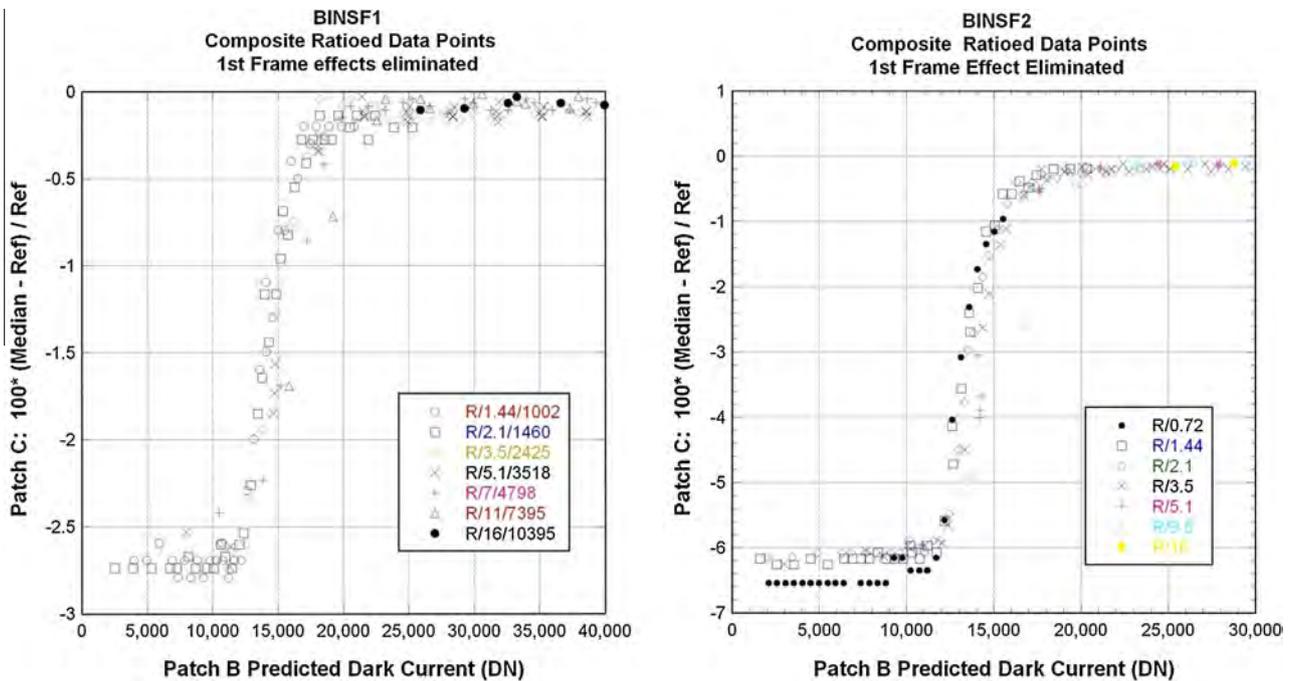


Fig. 21. Percent difference in background level in the IR subframe modes as a function of the amount of charge accumulated in the unreset border regions of the detector. Patch B is an area in the unreset region, and Patch C is in the read-out subframe area.

background rate due to thermal emission from the bench or to the signal rate from an external source being observed. Thus, it cannot be correctly modeled as either a fixed signal offset (its magnitude is integration time dependent) or a gain factor (its magnitude does

not depend on either the thermal background rate or the scene signal rate). The glow signal only looks like a gain factor in the linearity data because the thermal background rate is nearly the same within and between calibrations. Correcting the subframe glow

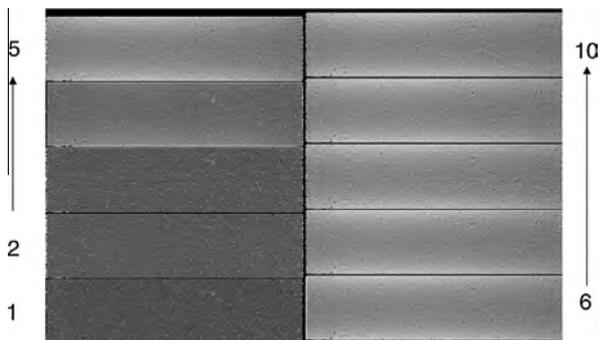


Fig. 22. Ratio of background signals for a sequence of BINSF1 subframes divided by the 10th BINFF full-frame image in a sequence. Frames 1–3 do not show any glow from the unset regions because their accumulated charge is still <11,000 DN. Frame 4 is a transition frame with unset region signal between 11,000 and 20,000 DN showing the onset of the glow pattern. Frames 5–10 show the full stable glow pattern that results when the charge in the unset regions exceeds 20,000 DN.

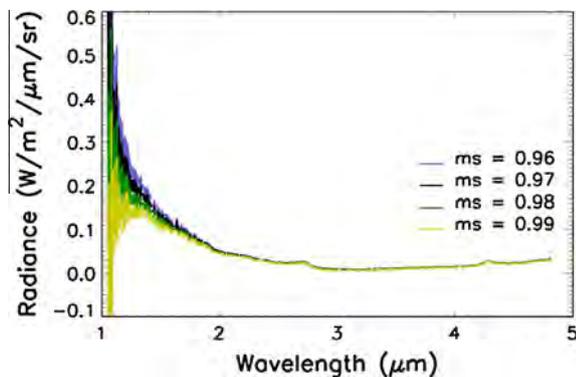


Fig. 23. Comet Hartley 2 example spectrum showing the effects of improper IR spectrometer background removal. Depending on the dark scaling (ms = scaling factor), here only varying by $\sim 2\%$, the short-wavelength continuum radiance goes from positive to negative.

effect by subtracting a mode-dependent master dark is acceptable provided the SIM temperature of the frame being calibrated is close to that when the frames used to create the subframe master dark were taken (~ 137.5 K). This is the approach used in the pipeline processing for PDS Version 2.

4.3.3.7. Background removal. The ultimate goal in the background analysis was to understand how to remove it from individual images. As long as the dark current contribution remains negligibly small, the background rate varies uniformly across the field of view as the SIM bench temperature varies. So theoretically a master background frame can be created and scaled up or down on the basis of SIM temperature and integration time to match the level in the image of interest. Tests showed that, under normal circumstances, subtracting this scaled master dark from an image worked well for removing the background, provided the scaling could be accurately determined. Initially, the master frame was simply scaled automatically using the relevant exposure time and bench prism temperatures to create a background subtraction frame for each observation. It was determined after many tests that the temperature-dependent function for scaling the master dark (Klaasen et al., 2008 and Fig. 13) was not sufficiently accurate for the Hartley 2 data (i.e., the data were never fit well using scaling based on either individual frame temperatures, smoothed temperatures, or daily average temperatures) without additional scaling and that background uncertainties of a few percent were much too large to be useful for the data analysis that was being done (Fig. 23).

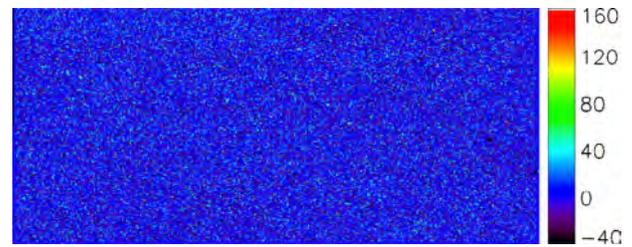


Fig. 24. Representative dark data frame from a Hartley 2 observation scan on October 17, 2010 (E-34d, DOY 290, hi0340609368_4000001_017.fit) with pixel-by-pixel non-linearity correction and last-frame dark subtraction applied. In comparison to Fig. 6, using this calibration procedure removes the instrumental artifacts in the frame, making this method the preferred calibration for the low-signal data acquired of Hartley 2's coma several days from closest approach. As in Fig. 6, the color bar stretch is linear from -40 to 160 DN. Unlike Fig. 6, the signal levels in the frame are greatly reduced and evenly bound 0 DN.

Because of the accuracy that was needed, it was determined that generating a sufficiently accurate background frame strictly from recorded prism temperatures and integration times would not be possible. Instead, two strategies were employed that incorporated information from the images themselves to help constrain the background level: (1) use an in-scene dark frame or (2) create a manual scaling factor to apply to the mode-dependent dark frames unique for each scan to produce the most physically plausible spectrum. These are similar methods to those applied to Tempel 1 data analysis in 2005 and are described in more detail in the following paragraphs.

The first method is applicable to data scans with at least one frame free from target-generated signal and takes advantage of the fact that most of the cometary sequences were spatial scans with at least one frame situated far from the nucleus, meaning that each scan contains a region that is essentially just background. This method is used in PDS Version 2 of the calibrated data set for the Hartley 2 approach data acquired from October 1, 2010 through November 3, 2010 and then after closest approach from November 10, 2010 through November 26, 2010. Fig. 24 shows the resulting data frame after the new pixel-by-pixel non-linearity correction and last-frame dark subtraction were applied to the Hartley 2 data in Fig. 6. Since all images in a scan were typically obtained within a few-minute time span, it can be assumed that the conditions remain constant throughout, and thus, to first order, the background level of the last frame represents that for all the frames in the scan.

An independent test validating the non-linearity correction and dark-subtraction procedure was also conducted on a lunar scan acquired with the spectrometer in the smallest sub-frame mode (Mode 3) on December 5, 2009. These data include frames with dark sky at the end of the scan. Fig. 25 shows that, similar to the Hartley data, the pixel-by-pixel non-linearity correction is improved over the quadrant-averaged non-linearity correction, especially at the edges of the detector in the dispersion direction. Sub-frame glow structure remains in the top three panels of Fig. 25, but the mode-dependent subtraction treatment (C) is better than the two that ignore (A) or model (B) the sub-frame glow. The in-scene dark frame produced the most uniform background with signal levels near zero DN. When an in-scene dark frame is available, as it is for the Hartley approach and departure frames, it is the best dark to use.

However, when an in-scene dark frame is not available, the preferred method to use is the pixel-by-pixel linearity treatment with manual scaling of the mode-dependent master dark. For the data acquired near closest approach to Hartley 2, from November 4, 2010, through November 9, 2010, there is not a “clean” background region due to the bright coma present in each frame of the scans. Optimal manual scaling factors for the mode-dependent master

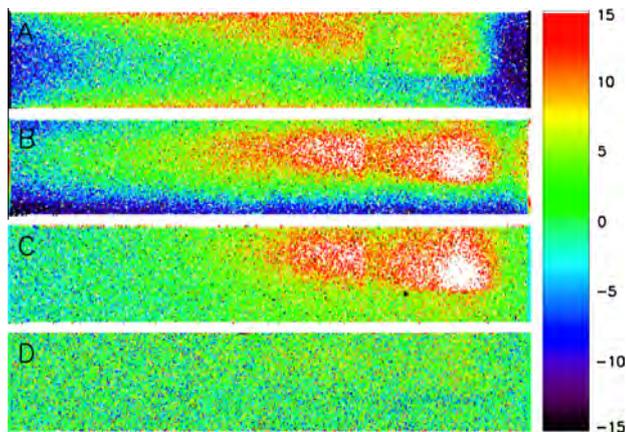


Fig. 25. Representative IR ALTF data frame of dark sky from lunar scans on December 5, 2009 with various non-linearity correction and background subtraction procedures applied: (A) quadrant-averaged non-linearity correction and temperature-dependent functional scaling of the mode-independent master dark, (B) quadrant-averaged non-linearity correction with in-scene scaling of mode-independent master dark with analytical sub-frame glow correction, (C) pixel-by-pixel non-linearity correction with in-scene scaling of the mode-dependent master dark (which includes the sub-frame glow correction), (D) pixel-by-pixel non-linearity correction with in-scene dark frame subtraction. The linear color bar stretch from -15 to 15 DN is identical in all four panels.

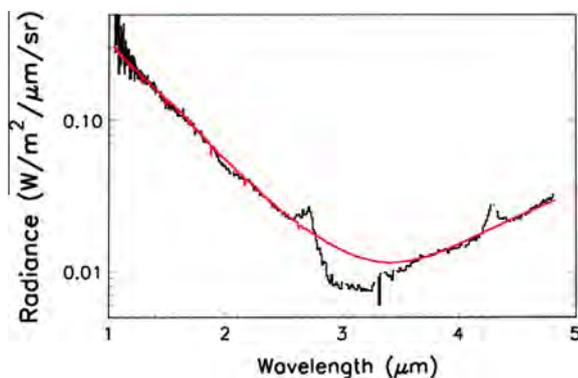


Fig. 26. Illustration of the model (red line) used to reproduce the continuum of a representative HRIIR (black line) Hartley 2 spectrum. The spectrum shows clearly the water ice absorption at $3 \mu\text{m}$ and water vapor, organic, and carbon dioxide emissions at $2.7 \mu\text{m}$, $3.3 \mu\text{m}$, and $4.3 \mu\text{m}$, respectively, which are not fit by the model, as expected.

dark need to be determined by examining each scan. The technique for determining the manual scaling factor to use in each scan is as follows.

A range of manual scaling factors is applied to the appropriate mode-dependent master dark for a given scan, the scaled dark frames are subtracted from each frame of the data, and then a continuum is fit to the dark-subtracted data. The nominal scaling factor is selected based on the ability to achieve a best-fit to the data in the continuum region of the spectra. The cometary continuum in the $1\text{--}5 \mu\text{m}$ region, which underlies several absorption and emission features of dominant cometary molecules, is composed of scattered solar radiation at short wavelengths ($<2.5 \mu\text{m}$) and a thermal contribution due to the hot dust at longer wavelengths. The scattered flux is estimated using a reddened solar spectrum normalized to the data in the wavelength range $2.20\text{--}2.25 \mu\text{m}$, where there is a good SNR, no expected absorption features, and negligible thermal emission. For the thermal infrared contribution the product of the infrared emissivity and the Planck function is

employed. The continuum level of the spectral radiance is the sum of the scattered light and thermal infrared emission (see Fig. 26). The free parameters in the modeling are the spectral slope of the solar spectrum and the temperature in the Planck function. The best fit is obtained using a least squares method. This approach is commonly used to model the continuum spectra of comets. Hartley 2 spectra display strong water ice absorptions, centered at 1.5 , 2.0 , and $3.0 \mu\text{m}$, together with water gas, organics, and carbon dioxide emission bands. In order to avoid these absorption and emission bands, the continuum modeling is performed in the wavelength ranges $1.20\text{--}1.30$, $1.70\text{--}1.80$, $2.25\text{--}2.50$, $3.50\text{--}4.20$, and $4.40\text{--}4.70 \mu\text{m}$.

Fig. 27 shows an example of the master-dark scale factor determination process for a typical Hartley 2 coma measurement. The best-fit master-dark scale factor in this case is determined to be 0.97 .

4.3.3.8. Sequencing techniques to stabilize background. After lessons learned from the Tempel 1 flyby observations, the Hartley 2 observation sequences were purposely planned to avoid single-frame observations, to avoid sub-frame observations where possible, and to read and reset the detector four times prior to the start of each scan to reduce the effects of the overshoot in the first frame described above and in Klaasen et al. (2008). We do not apply any second-order correction for the monotonic falloff (or sometimes the increase) as a function of time since the beginning of each sequence; however, the background transients remaining after the four unsaved reset frames are relatively small (Fig. 20).

4.3.4. Wavelength map

The DI wavelength map was obtained from laboratory measurements, using argon and krypton lamps and absorption lines of CH_4 and CO (Klaasen et al., 2008). From these ground-based measurements, the wavelength map accuracy had been estimated to be ~ 2 physical pixels over the entire detector (one binned pixel). Between 2007 and 2010, we performed several new calibration runs using the star Beta Hydri and the Moon in order to test and increase the accuracy of the DI wavelength map. Some of these calibrations were acquired in the full-resolution UBFF mode, providing data with the highest spectral resolution.

The spectra of β Hyi and of the Moon can be approximated by a blackbody continuum on which several stellar lines are clearly visible. Their spectra are well tabulated in the literature. More than 30 lines have been identified in the spectra of β Hyi (see Fig. 32 in Section 4.3.5 for an example) and the Moon, allowing a detailed investigation of the wavelength calibration accuracy from 1.2 to 3.1 microns.

Some differences from the DI wavelength map have been observed. Fig. 28 presents the differences between the expected locations of the stellar lines (tabulated in the literature by Lançon and Rocca-Volmerange (1992) and Wallace et al. (2000) for instance) from the DI wavelength map and the actual locations measured from the EPOXI data across the entire field of view (FOV). Differences are shown in terms of wavelength (in μm) and in terms of number of pixels. The differences reach -0.5 pixels on the top and $+0.6$ pixels on the bottom of the FOV. From the top to the bottom, this difference has been fitted by a linear regression in order to obtain its intercept and its slope.

Since the differences in terms of pixels are similar at all wavelengths, we suspected a difference in the spectral smile correction from that derived during DI. The smile correction compensates for wavelength-map changes along the spatial axis (see Klaasen et al. (2008) for a complete report of the smile effect). We compared the smile's shape for different spectral lines obtained from lunar data with the expected shape computed from the DI calibration. Measured and theoretic smile shapes agreed within a few tenths of a

pixel from the bottom to the top of the FOV at different wavelengths, a much smaller difference than the $\pm 0.5/0.6$ pixels observed in Fig. 28. This result confirms the accuracy of the smile shape correction between the pre-launch (in 2005) and the EPOXI (in 2009) calibrations. We identified the origin of the errors as a displacement of the center of the smile in the FOV. In the DI smile correction, the center of the smile was close to line 262 along the spatial direction (6 pixels above the center of the FOV in full-frame mode). A shift in the center of the smile of several pixels in the spatial direction reduces the mismatches. In order to identify the proper number of pixels to shift the smile center (which corresponds to a displacement of the wavelength map along the spatial direction), we used the lunar calibration data, for which the Moon light has been dispersed along the entire slit. We identified several solar lines in the 1.2–3.0- μm range, and we shifted the wavelength map along the spatial direction in order to minimize the mismatches along the entire axis. Fig. 29 gives the resulting slope in the top-to-bottom wavelength fits, evaluated from the linear regression presented previously, as a function of the wavelength map shift in pixels. The correct center of the smile has been found when the slope of the linear regression equals zero (i.e., when the wavelength differences from top to the bottom are exactly identical).

The wavelength map discrepancies are minimized when the DI wavelength map is shifted by 15 ± 2 unbinned pixels along the spatial direction. The variation of the best-shift value to apply depends on several parameters, such as the date of the calibrations (tested between December 2007 and December 2009) and the spectrometer temperature (tested from 137.2 to 138.7 K) as well as uncertainties in the regression fits. This variation is small, however, so we decided to adopt a single shift for all data in the new calibration (15 pixels). The new wavelength map has been tested on several data sets, including spectra of Vega and Sirius, and the accuracy is now better than one physical pixel in the 1.3–2.5- μm range and better than two physical pixels over the entire detector (unbinned mode). For spectrometer temperatures above 138 K, the wavelength calibration is slightly degraded by a few tenths of a pixel. However this temperature is almost never reached in our scientific data.

4.3.5. Absolute spectral response

In late 2010, it was discovered that the IR spectrometer radiometric calibration curves were in error by a factor of 2. This error was uncovered by comparing stellar irradiances generated by processing IR spectrometer star scans through the standard calibration pipeline (Klaasen et al., 2008) with irradiances published in the lit-

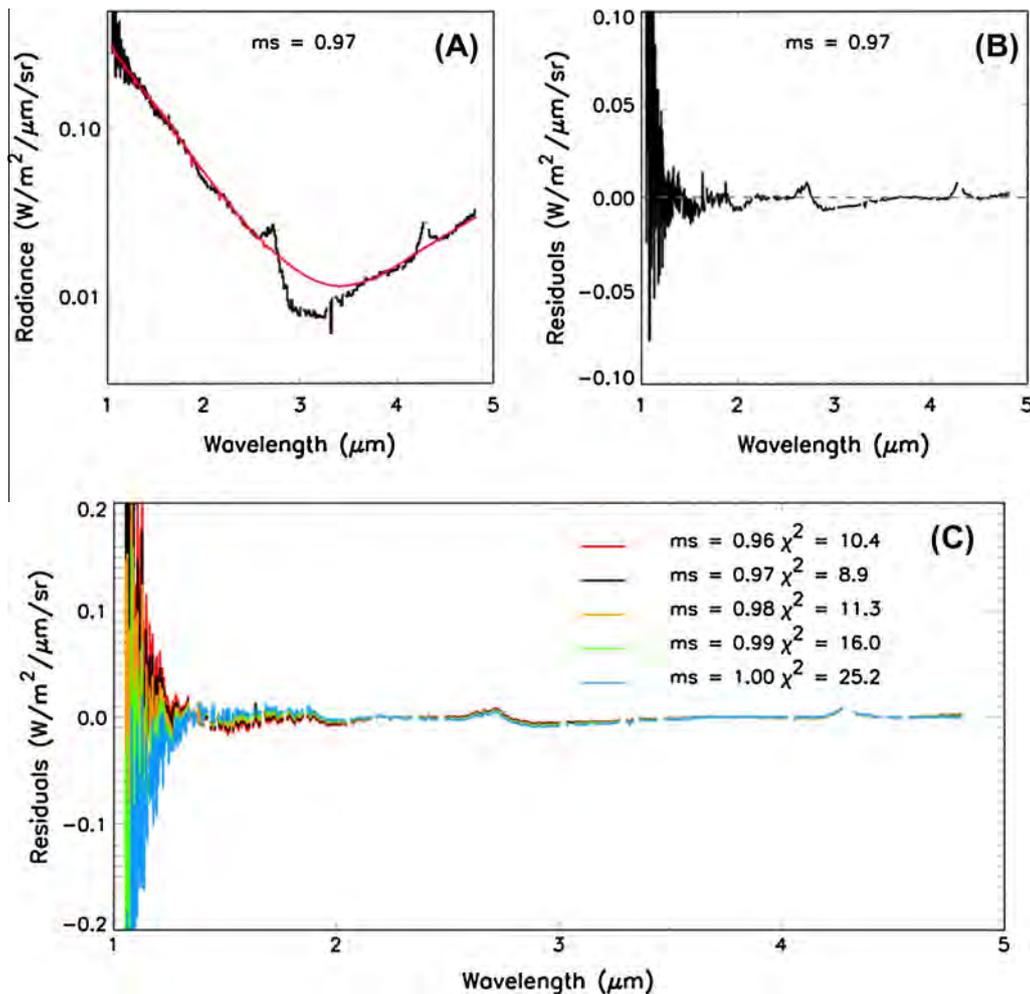


Fig. 27. Determination of the manual scaling factor to apply to the master dark frame subtracted from representative Hartley 2 coma spectral data; the best-fit factor results in small continuum residuals across the entire spectrum. Panel (a) shows the model (red line) used to reproduce the continuum of a representative HRIIR (black line) spectrum. The spectrum has been obtained using the optimum manual scaling factor ($ms = \text{scaling factor in plot legend}$) of 0.97. Panel (b) displays the residuals between the model and the data. Panel (c) shows the effects of an improper IR spectrometer background removal. The residuals between the measured spectrum for different manual scaling factors and the corresponding continuum best-fit modeling are shown.

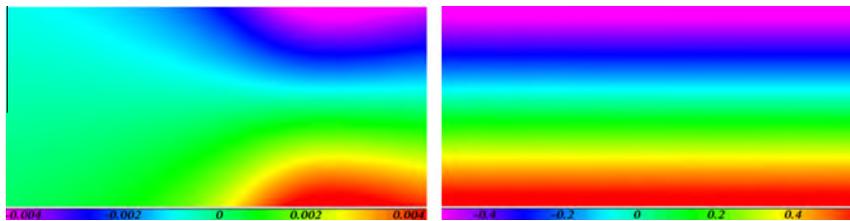


Fig. 28. Differences between the DI IR spectrometer wavelength map and that measured during EPOXI across the detector FOV; in units of μm (left) and physical (unbinned) pixels (right).

erature. The error was introduced by using the wrong pixel solid angle in the derivation of the original calibration curves (a pixel solid angle of $5 \times 5 \mu\text{rad}$ was assumed, but the correct value is $5 \times 10 \mu\text{rad}$ with the larger dimension being set by the slit width, not the pixel size on the detector). Therefore, the DI calibration curves in Klaasen et al. (2008) should be divided by 2. An independent validation of this correction is reflected in the work of Livenood et al. (2011).

In addition, the accuracy of the DI calibration curve was less than desired for the science analyses being undertaken. The noise was not entirely removed, the calibration longward of $4.8 \mu\text{m}$ was suspect, and the short-wavelength end contained ripples due to the internal beamsplitter that were not adequately characterized. As a result, additional calibration data of the photometric standard star β Hvi were taken in 2009 and 2010. β Hvi was chosen because it is a close solar analog, has a well-measured spectral irradiance in the shortwave IR, and is bright enough to yield good SNR in our IR spectrometer scans.

A total of 155 scans of β Hvi outside the anti-saturation filter (ASF), which covers the central $\sim 1/3$ of the spectrometer slit, were taken at a scan rate of one slit width per integration period with 50-frames per scan. The first 5 frames of each scan were discarded due to the first-frame transient effect. The data were initially quad-average linearized and flat-fielded using the flat-field file derived using quad-averaged linearization. The total dark-subtracted flux for each spectral column in each scan was obtained by summing over the 21 frames centered on the one containing the peak star signal and over 25 spatial pixels along the slit. The dark frame for each scan was normally derived from the resistant mean of the six frames on each side of the 21 containing star signal. The measured fluxes from all 155 spectra were combined using their resistant mean value at each column, and each column was assigned its proper wavelength and wavelength interval based on the row location of the spectral signal and the spectrometer wavelength map. The spectrometer bench temperature and the row locations of the star spectra varied only slightly over the entire data set, so any wavelength misregistration to the columns was reduced to a few tenths of a pixel. The resulting spectrum is given in Fig. 30. The residual uncertainty in the spectrum based on the standard deviation across the data set is plotted in Fig. 31.

From the spectrum shown in Fig. 30, the spectrometer radiometric calibration curve was derived at high spectral resolution (0.5 nm) as follows. The irradiance spectrum of β Hvi as measured by ISO (Vandenbussche et al., 2002) was taken as our reference. However, since the ISO spectrum is not complete across the entire 1–5 μm range, we replaced it with the MODTRAN[®] solar irradiance spectrum (1-nm resolution; http://redc.nrel.gov/solar/spectra/am0/other_spectra.html) scaled to fit the ISO spectrum in the 2.3–4.8 μm range. β Hvi is a G2IV star, a bit colder and older than the Sun, and a comparison of the scaled MODTRAN[®] and β Hvi spectra showed a slight difference in the continuum, in agreement with the temperature difference of the two stars. The continuum slope difference was estimated to be $-0.8\%/micron$ – consistent with the expected value from J, H, and K photometry measure-

ments (SIMBAD JP11 catalog <http://simbad.u-strasbg.fr/simbad/sim-fid>, Campins et al., 1985). The scaled MODTRAN[®] spectrum was then corrected using this continuum slope difference. A preliminary spectrometer calibration curve was produced by dividing the resulting irradiance curve at the corresponding wavelength of each column by the β Hvi spectral response of the IR spectrometer and multiplying by the wavelength interval of each column.

The resulting calibration curve showed some features at high spectral frequency that are correlated with stellar absorption lines and are clearly not instrumental. A correction for these differences between β Hvi and the Sun in terms of metallicity was made by averaging across these features in the calibration curve (see Fig. 32). The final calibration curve was tested on lunar data, and the check validates both the continuum slope and the stellar absorption line corrections.

The ratio between the new calibration curve and that from Klaasen et al. (2008) corrected for its $2\times$ error is shown in Fig. 33, and the old and new curves are both plotted in Fig. 34. This new calibration curve was used in producing PDS Version 1 of the IR data set.

Longward of $4.3 \mu\text{m}$, the β Hvi spectrum has low SNR. To improve the accuracy of the spectral calibration at the long wavelength end, 72 additional scans of the bright star Vega were obtained and analyzed in 2011. To maximize the SNR, the Vega data were acquired with 2×2 -pixel binning. The pixel binning places a lower limit on the wavelength registration accuracy of these data. The final calibration curve to be used for PDS Version 2 uses the Vega-derived calibration curve longward of $4.3 \mu\text{m}$.

In addition, the spectral calibration curve was re-derived using all of the pixel-by-pixel linearization based files. No significant differences in the resulting calibration curve were observed except for a minor ($\sim 1\%$) decrease in response rate calibration between 2 and $2.5 \mu\text{m}$. The ratio between the Version 1 and Version 2 calibration curves is shown in Fig. 35. For consistency, the new calibration curve derived using pixel-by-pixel linearization based files will be used for the PDS Version 2 delivery.

Validation of the IR spectrometer radiometric calibration accuracy was performed by comparing calibrated radiances to those obtained from simultaneous VIS images. Fig. 36 shows the results. At 0.52 and $0.74 \mu\text{m}$, where the continuum dominates, the visible data deviate from the continuum modeling by 9% and 1%, respectively.

The stability of the IR spectrometer absolute spectral response was monitored throughout the mission using the same technique and the same set of standard calibration stars (see Table 1 for a list) as was used in the primary DI mission (Klaasen et al., 2008). Examples of the measured total star signals vs. detector column from the November 2010 calibration are shown in Fig. 37 both outside and behind the ASF. Comparisons with similar measurements from the primary mission are shown in Fig. 38. No changes were observed to within the single-scan measurement accuracy, which we estimate to be about $\pm 10\%$ absolute and about $\pm 5\%$ for a given star.

The stellar photometry measurements suggest an increase in the ASF transmission of about 10% at all wavelengths for data taken

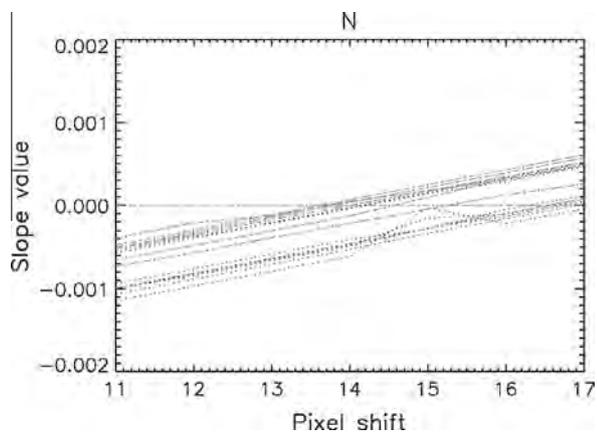


Fig. 29. Slope of the wavelength differences from top to bottom of the IR spectrometer FOV for a range of spectral lines vs. the applied shift in the assumed line number of the center of the spectral smile; zero slope value occurs when the smile is perfectly centered. Dotted lines: mean of all reference spectral lines; dashed lines: 1.282 μm reference line only.

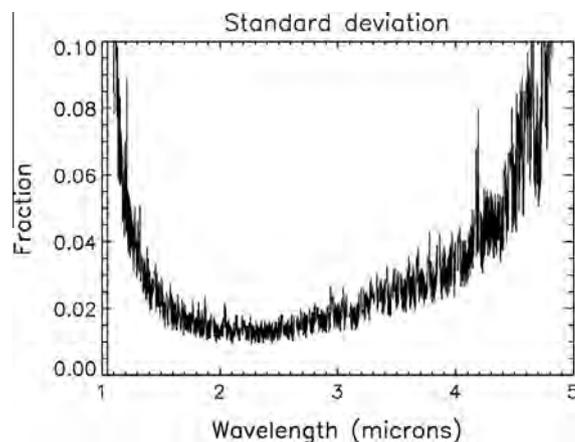


Fig. 31. Uncertainty in Beta Hydri spectrum measured with the IR spectrometer.

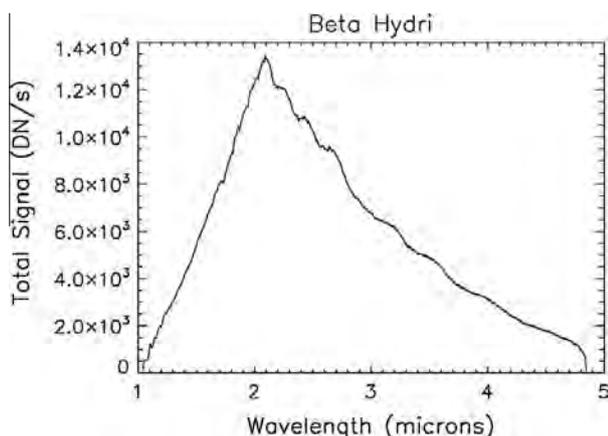


Fig. 30. Spectrum of Beta Hydri observed with the IR spectrometer outside the anti-saturation filter.

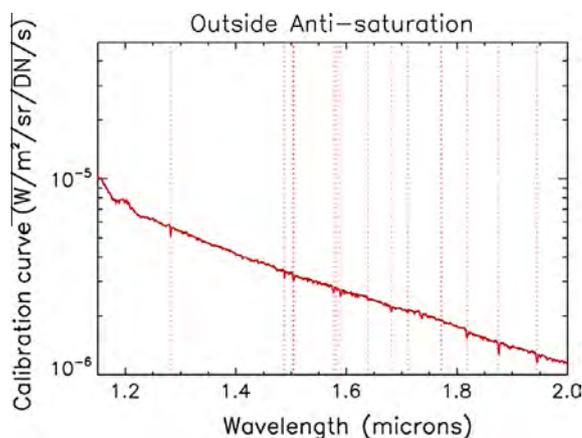


Fig. 32. Section of smoothed calibration curve (black) and original curve with effects of differing stellar absorption lines between β Hyi and the Sun (red). The features expected from the metallicity differences are identified by the vertical dashed lines and correlate well with the features seen.

in 2010 and 2011. This change is consistent with the increase in ASF transmission derived as part of the flat-field determination (see Section 4.3.7) as compared to that determined from stellar photometry during the prime mission (Klaasen et al., 2008). An additional increase is suggested longward of about binned column 425 ($\sim 3.8 \mu\text{m}$), but the filter transmission is dropping rapidly with wavelength in this region, and the SNR is quite poor. The ASF transmission in this region derived from the flat-field determination also shows a progressive increase with wavelength compared to the ASF transmission curve derived during the primary mission. Further refinement in the determination of the ASF long wavelength transmission is discussed in Section 4.3.7.

4.3.6. Bad pixels

Given current technology, IR detector arrays inherently have more defective and non-uniformly responsive pixels than CCD chips. Therefore, these pixels must be identified so that the calibration pipeline can interpolate over the defective areas to meet the scientific objectives. During most of EPOXI, we continued to apply the same algorithm developed for DI (which applied a quadrant-averaged linearity correction; Klaasen et al., 2008) to identify bad pixels using sets of dark frames taken at increasing exposure times. For our analysis, we looked for several types of defective or bad pixels in the active area of the array (reference rows and columns were excluded from the analysis): hot pixels that rapidly saturate

or that are always saturated, cold pixels that always have a low signal compared to neighboring pixels, and pixels with atypical or very erratic responses, including pixels with a response that could not be linearly fit adequately. Table 4 provides the number and percentage of bad pixels in the active image area using unbinned, full-frame dark data from the two pre-flight calibrations and the 10 inflight science calibrations from April 2005 through November 2010. The table includes focal-plane and optical-bench (prism) temperatures for each calibration.

The gradual increase in the number of bad pixels during DI and EPOXI is not unexpected, because IR detectors degrade over time. The additional bad pixels are dispersed randomly across the active array in both the spatial and spectral directions; there is no concentration at longer wavelengths and no clustering near “always bad” pixels. There are 4401 of the same pixels that are always flagged as bad for all 10 inflight calibrations. For the six calibrations during EPOXI, 5469 of the same pixels are always bad. Isolating the last two EPOXI calibrations in 2010, 6900 of the same pixels were bad.

We used the same process as was used on DI to make separate bad-pixel maps for each binned mode for each inflight calibration for EPOXI. Table 5 provides the number of bad pixels for each map implemented in the pipeline for EPOXI and the effective dates over which each map was used for PDS Version 1 (McLaughlin et al., 2011a,b,c).

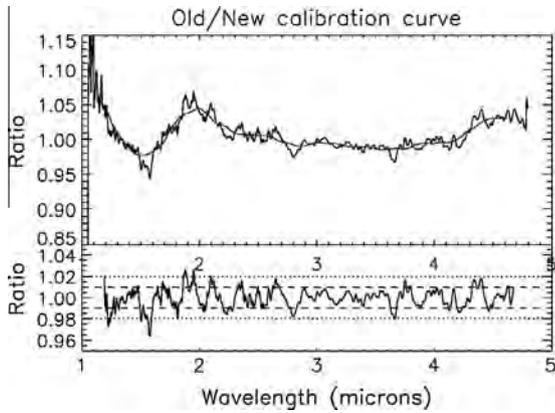


Fig. 33. Ratio of the IR spectrometer radiometric calibration curve of Klaasen et al. (2008) divided by 2 (see text) to the newly derived curve based on extensive β Hyi measurements. The smoothed curve shows the low-frequency difference. The bottom plot shows the remaining noise (not removed in the final calibration curve) after removing the low-frequency variations. Corrections of up to 5% to the old curve appear warranted with residual random noise of about 2% or less over the range from 1.2 to 4.6 μm .

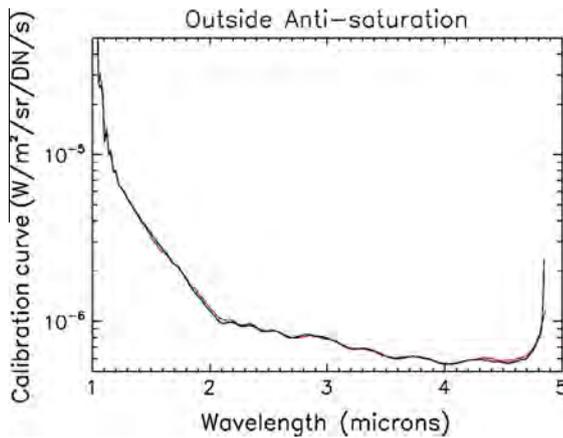


Fig. 34. IR spectrometer curves outside the anti-saturation filter– old from Klaasen et al. (2008) divided by 2 (red) and new from this work (black). The calibration curve below 1.0425 μm and above 4.83 μm is undetermined.

The unbinned, full-frame (Mode 4) darks taken at increasing integration times during the February 2011 non-linearity calibration provided a data set that allowed an improved screening method for defining bad pixels. As for the linearity analysis and master-dark creation described above, the last 40 frames (of 50) of each of the 11 integration times <16,200 ms were used. The bad pixels were defined concurrently with the pixel-to-pixel non-linearity analysis.

We developed the following criteria and accompanying algorithm to identify bad pixels in the inflight data.

1. Pixels are flagged as bad if they are outside of the active area of the chip. The active area of the chip is defined with $[x,y]$ coordinates in a zero-indexing system as $[5:1018,6:510]$ (Hampton et al., 2005).
2. Pixels that have too much variation within the pixel for a similar set of exposure parameters across the set of 40 frames are flagged as bad. A resistant mean is run for each pixel through the 40 frames having identical exposure times with a cutoff of 3-sigma. Any pixels with more than 25% of the data points rejected by the resistant-mean test are considered to have too much variation and are flagged as bad pixels.

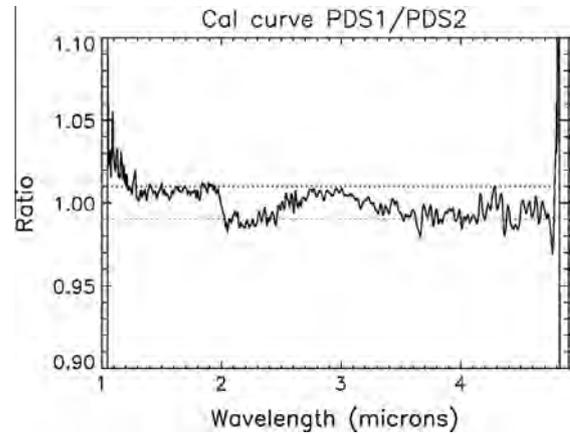


Fig. 35. Ratio of the PDS Version 1 (quad linearization) to the PDS Version 2 (pixel-by-pixel linearization) IR spectrometer absolute spectral calibration curves.

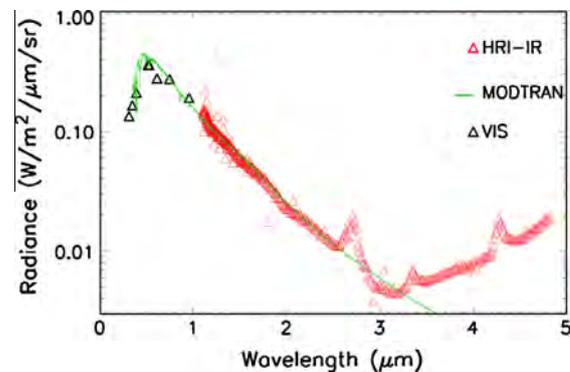


Fig. 36. Example spectrum of Comet Hartley 2 covering the wavelength range 0.3–4.9 μm . The spectrum corresponds to the jet region off the end of the small lobe facing the Sun observed 7 min after closest approach. The scattered flux (green line) is estimated using a sloped MODTRAN solar spectrum. The spectral slope is a free parameter iteratively modified by means of a χ^2 minimization algorithm until a best-fit to the observations is achieved. The best-fit is performed in the wavelength ranges 1.20–1.30, 1.70–1.80, and 2.25–2.50 microns where there is good SNR, no expected absorption features, and negligible thermal emission. The visible (black) and IR (μm) radiances show good agreement.

3. Flag pixels as bad if they respond with a resistant mean greater than 12,000 DN at any exposure time <16,200 ms. The product of the average temperature-dependent dark rate and these exposure times is well below 12,000 DN; therefore, if a pixel responds with a resistant mean >12,000 DN for any exposure time used in this analysis, the pixel is considered a hot pixel.
4. A pixel in a region with too much scatter that cannot be linearly fit with a y-intercept of zero is determined to be a bad pixel. The resistant mean, with a 3-sigma cutoff, for a 7×7 -pixel box is calculated for each pixel location for each exposure time (49 pixels, 40 frames). An intermediate flag is assigned where more than 25% of the data points, i.e. >490 pixels, are rejected by the resistant-mean test because there is too much variation within the region for a given exposure time. A linear fit of the resulting mean response with time for the unflagged exposure times (provided there are at least 3) is then made, using the form $y = mx$, because the IR signal should be 0 DN for an integration time of 0 ms. If the linear fit cannot be made (i.e., there are <3 data points to fit), then the pixel is determined to be bad.
5. A linear fit is also calculated for each individual pixel's response with time, requiring at least three non-flagged values and a y-intercept of 0. Again, if <3 data points are not flagged, that pixel's linear response cannot be fit, and it is defined as a bad pixel.

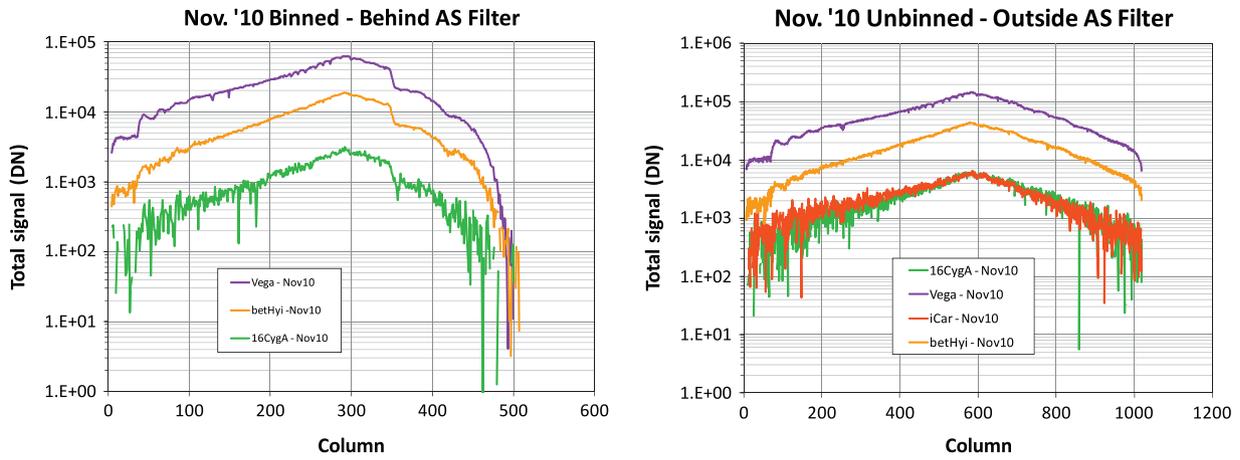


Fig. 37. Representative IR spectrometer stellar photometry results.

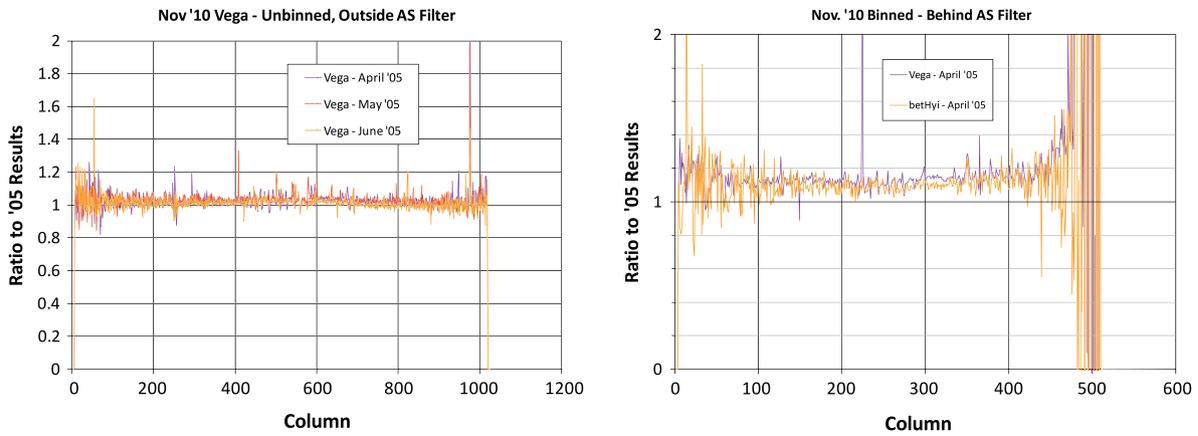


Fig. 38. Comparison of IR spectrophotometry to that from the DI prime mission.

Table 4

Number of IR spectrometer bad pixels identified using unbinned, full-frame, dark data from pre-flight calibrations (TV2 and TV4) and inflight calibrations performed during Deep Impact and EPOXI.

Mission	Calibration set	IRFPA temp (K)	Optical bench temp (K)	# Bad pix	% Bad pix	Mode
Deep Impact (pre-flight)	TV2	82.6	130	5036	0.98	6 (DIAG)
	TV2	85.1	136	4940	0.96	6 (DIAG)
	TV2	87.2	141	4259	0.83	6 (DIAG)
	TV4	106.3	139	37,482	7.31	6 (DIAG)
Deep Impact (in-flight)	April 2005 ^a	84.97	139.52	5518	1.08	4 (UBFF)
	May 2005 ^a	84.28	137.80	5823	1.14	4 (UBFF)
	June 2005 ^a	84.07	137.05	5463	1.07	4 (UBFF)
	July 2005 ^a	84.37	136.93	5870	1.15	4 (UBFF)
EPOXI	January 2008 ^a	84.4	140.4	7030	1.37	4 (UBFF)
	June 2008 ^a	84.1	137.1	7596	1.48	4 (UBFF)
	June 2009 ^a	84.1	137.2	7352	1.44	4 (UBFF)
	February 2010 ^a	84.3	137.5	7617	1.49	4 (UBFF)
	September 2010 ^b	83.6	137.3	7616	1.49	4 (UBFF)
	November 2010 ^b	83.6	136.5	7835	1.53	4 (UBFF)

^a Linearity constants derived from the same calibration set were applied.

^b Linearity constants derived from the February 2010 calibration set were applied, because the IR teams determined there was no appreciable difference in the linearity since early 2010.

6. A pixel is bad if its response is greatly different from that of its nearest neighbors. The linearity response curve derived in criterion 5 is compared to that in criterion 4. If the slope for an individual pixel is more than a factor of 1/3 different than the slope for the mean of the pixel's local region, then that pixel is considered bad.

Once these criteria have been imposed, a comprehensive bad-pixel map is made setting all of the selected bad pixels using the criteria listed above to a value of 1. The pixels in the reference rows and columns (i.e. non-active area on the chip) are also set to 1. All other pixels in the active image area are set to a value of 0, indicating good pixels. The bad-pixel maps for binned and sub-frame

modes are derived from this map. If any one or more of the four unbinned pixels associated with a single binned pixel is bad, then the binned pixel is also defined as a bad pixel.

Table 6 provides the number and percentage of bad pixels in the active image area using the unbinned, full-frame dark data from the February 2011 non-linearity calibration. For comparison, the table also includes the number of bad pixels identified in July 2005 and used in the pipeline for the Tempel 1 encounter data and those in November 2010 used in the pipeline for PDS Version 1 of the EPOXI data (McLaughlin et al., 2011a,b,c) along with the focal plane and optical bench temperatures. Fig. 39 presents an image of the latest bad pixel locations. This new bad-pixel map will be used for PDS Version 2 of the EPOXI IR data.

Once bad pixels have been identified, they are reclaimed or recovered in the non-reversible pipeline data product by linearly interpolating between adjacent, non-saturated, good pixels – one pixel above and below the bad pixel. This interpolation is performed only in the spatial dimension, as most of the interesting science measurements are performed on an extended object (the comet) where high spectral resolution is important for identifying species. Finally, these pixels are flagged as bad in the quality map extension of the non-reversible products so that the user knows that their values are interpolated and are not true measurements.

4.3.7. Flat field

As discussed by Klaasen et al. (2008), no complete flat-field calibration of the IR spectrometer was obtained pre-launch or during the primary DI mission. Therefore, no flat-field corrections were applied to the DI IR spectrometer data.

Several attempts to collect suitable flat-field calibration data were made during EPOXI using the technique of rapidly scanning the slit across the Moon in the along-slit direction (see Table 1). Various scan rates were used to adjust the total signal collected so as to provide unsaturated data at long wavelengths using fast scans and good SNR data at short wavelengths using slow scans. Multiple repeat scans at each rate across different size chords on the lit Moon were performed to build up the SNR. The final flat-

field calibration file was constructed using data from 174 separate scans acquired during 2009 (see Fig. 1).

The signals collected in the 8 or 10 frames from each chord/scan-rate combination were combined at each pixel using a resistant mean calculation after linearization (quad-average for PDS Version 1; pixel-by-pixel for Version 2), dark subtraction, and masking out of all saturated pixels. This approach removed the effects of cosmic ray events. The resistant mean values were then weighted by the total time that each pixel was on the lunar disk depending on the scan rate used and combined to give one frame for each chord. This process produced a set of frames with a high-SNR lunar spectrum in each image line.

To derive the spectrometer relative response at each pixel, it was necessary to remove the lunar spectrum for each chord. The lunar spectra were obtained by using the rows outside of the ASF, removing the spectral smile, and taking resistant means along columns. An image containing the resulting spectrum in every row was then created for each chord, and the spectral smile was reinserted in these images. The lunar spectrum images were then divided into the corresponding original images to produce a flat-field frame for each chord including the ASF profile.

Because the registration of the ASF shifts slightly in column location with changes in bench temperature (-0.7 pixels/K), we elected to separate the ASF transmission function from the flat-field correction file. The spectral transmission of the ASF in each column was derived by ratioing the average of 20 rows just inside the ASF boundary to the average of 20 rows just outside the ASF boundary and averaging the result for the upper and lower boundaries for the flat-field images of each chord. The resulting curves were shifted slightly in their column positions to adjust for the known shift in the ASF transmission function with bench temperature (137.7 K was used as the reference temperature). The ASF transmission functions from all seven chords were then averaged to produce the final best-estimate ASF transmission function, which is shown in Fig. 40 for the quad-averaged linearization case.

A master flat-field calibration file that represents the pixel-to-pixel variations in detector response was then produced by dividing the rows lying behind the ASF by the ASF transmission function. The 14 rows at each boundary of the ASF region were handled by dividing by the best-fit fraction of the ASF transmission curve for each row. Certain special pixels had their values set to 1.0 (known bad pixels, pixels at the wavelength extremes lost in removing and adding back the smile, and reference pixel regions at the edges of the frames).

The resulting file using quad-averaged linearization is shown in Fig. 41. This flat-field correction file was used for PDS Version 1, since linearization was done using quad-averaged corrections. The longwave end behind the ASF has such low signal that an accurate flat-field correction is not possible (or necessary). Some of the spectral response variations due to the beamsplitter at short wavelengths have not been entirely removed because they vary slightly with row number, while the reference lunar spectra were averaged over all rows.

For calibration of IR frames in the pipeline, the ASF profile is shifted to the correct wavelength based on the bench temperature and reinserted into the appropriate master flat rows (including the edge-effect model at the boundary rows). The image is then calibrated by dividing by the combined master flat and ASF profile. Tests of this procedure have shown that corrections are good to 1% or better.

For PDS Version 2, a revised flat-field file based on pixel-by-pixel linearization was derived (Fig. 42).

In addition to a revised flat for Version 2, continuing radiometric analysis of the Hartley 2 nucleus revealed problems at long wavelengths for data located behind the ASF in that the calibrated nucleus radiance was too low to be physically real compared to the

Table 5
IR spectrometer bad-pixel maps used in the pipeline for EPOXI PDS Version 1.

Calibration set	Mode	# Bad pix	% Bad pix	Effective date for pipeline
January 2008	BINFF, ALTF	5286	4.15	2007-10-04
	BINSF1	2489	3.97	
	BINSF2	1133	3.77	
	UBFF, DIAG	7030	1.37	
June 2008	BINFF, ALTF	6040	4.74	2008-05-28
	BINSF1	2888	4.60	
	BINSF2	1305	4.30	
	UBFF, DIAG	7596	1.48	
June 2009	BINFF, ALTF	6170	4.84	2009-06-01
	BINSF1	2985	4.76	
	BINSF2	1356	4.47	
	UBFF, DIAG	7352	1.44	
February 2010	BINFF, ALTF	6483	5.08	2010-02-01
	BINSF1	3105	4.95	
	BINSF2	1427	4.70	
	UBFF, DIAG	7617	1.49	
September 2010	BINFF, ALTF	6522	5.11	2010-09-01 ^a
	BINSF1	3130	4.93	
	BINSF2	1447	4.99	
	UBFF, DIAG	7616	1.49	
November 2010	BINFF, ALTF	6491	5.09	2010-09-28
	BINSF1	3120	4.97	
	BINSF2	1432	4.72	
	UBFF, DIAG	7835	1.53	

^a Not presently used in the pipeline.

Table 6

IR bad pixel comparison between DI and EPOXI. The February 2011 bad-pixel map will be used for the PDS Version 2 dataset.

Calibration set	IR FPA temp (K)	Optical bench temp (K)	No. of bad pixels	% of bad pixels
July 2005	84.37	136.93	5870	1.15
November 2010	83.82	136.46	7835	1.53
February 2011	84.17	137.57	10,021	1.96

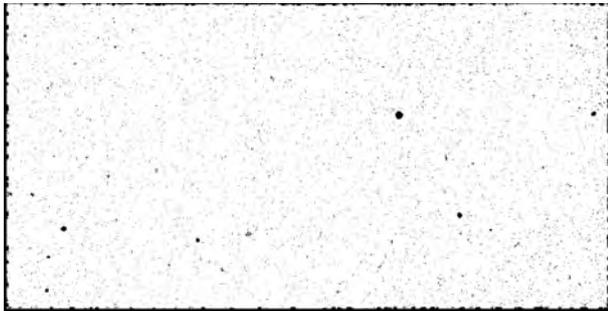


Fig. 39. Bad-pixel map for IR spectrometer unbinned full-frame Mode 4, based on data acquired during the February 2011 inflight linearity calibration. Bad pixels and reference row pixels are indicated in black. Good pixels are white.

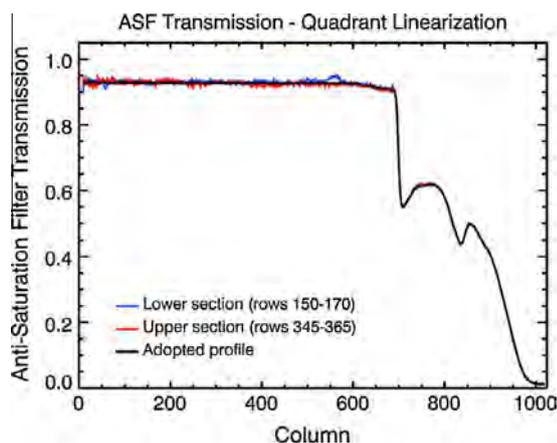


Fig. 40. The transmission of the ASF filter at each column position used for the PDS Version 1 dataset.

radiance at shorter wavelengths. Preliminary work with comet data revealed that there is a slight contamination, due to light from outside the ASF filter being scattered into the central region through the edge of the filter. This slight increase in the ASF signal is negligible for most of the spectrum, but at the long wavelength end where the transmission is very low, the scattered light contribution becomes comparable to the transmitted light signal. If not accounted for, this effect produces a significant overcorrection when applied to observations that do not have the scattered light component (e.g., when the object is entirely within the ASF region). Due to the low SNR of calibration targets in this region, the exact contributions of the transmitted and scattered light could not be independently separated, so an empirically derived value of 0.6% of the signal from outside the ASF was adopted as representing the scattered light contribution. This value produces acceptable results for both the standard star observations and for comet spectra at the long wavelengths. (Values between 0.4% and 0.8% produce comparable results, but deviations become noticeable outside this range.)

The first step in re-deriving the ASF transmission was to subtract 0.6% of the spectrum outside the ASF from that inside. The profile across the ASF region at long wavelengths shows a slight curvature, with the brightness near the edge of the filter being

higher than that at the center (Figs. 41 and 42 show this effect). This curvature was removed by fitting a parabola to the profile at column 1000 and then applying this shape at all wavelengths. After removal of the scattered light, the ASF spectral transmission was re-computed in the manner described above. Fig. 43 shows the resulting ASF transmission curve and its ratio to the previously derived curve. Note that the calibration pipeline does not attempt to correct for light scattered behind the ASF; we have just derived a more nearly correct ASF transmission curve for calibrating directly transmitted signal behind the ASF. Therefore, caution must be used when analyzing data at long wavelengths ($>4.5 \mu\text{m}$) behind the ASF when bright objects are located outside the ASF.

4.3.8. Noise

The random noise in the IR spectrometer detector readout was evaluated using the photon transfer technique (see Section 4.3.2). The minimum instrument bench structure temperature and the minimum available integration time precluded gathering data with signal levels below about 1500 DN; the level of photon shot noise in this signal meant that the read noise levels could only be approximated, but also that knowledge of this value is not very important. In the unbinned mode, the read noise was estimated to be between 5 and 6 DN. No significant differences were observed between the two detector quadrants. In the binned modes, the read noise was estimated to be between 3 and 3.5 DN. The read noise levels are about three times higher than those measured during the prime mission. No systematic changes were observed in the read noise over the course of the EPOXI mission.

4.3.9. Charge bleeding/residual

While no evidence of charge bleeding or residual image as a result of over-saturating the IR detector had been previously detected during DI (Klaasen et al., 2008), our more recent understanding of the cause of the first-frame effect arrived at during EPOXI (see Section 4.3.3.5) revived the expectation that we should see a residual image effect. The most sensitive test for this effect is with a frame that has a highly saturated local region and is immediately followed by a dark frame in which some residual signal might be detected. This type of frame pair exists in the IR spectrometer data set when a very bright cosmic ray signature is seen in an image from a set of otherwise dark frames. Six such frame pairs were identified for study from the first day of the Hartley 2 departure sequence.

Residual signal from the cosmic rays was in fact observed at a level of about 0.5% of the cosmic ray signal seen in the previous frame. For cosmic ray signals below about 4000 DN, no residual could be detected above the background noise level. For pixels saturated by cosmic ray signal, the residual is about 50 DN. This result supports the model of charge traps (Smith et al., 2008a,b) causing the first-frame effect.

5. VIS calibrations

5.1. Geometric

The two central rows of the CCDs in each camera are 1/6 of a pixel smaller vertically than a normal row. Therefore, recon-

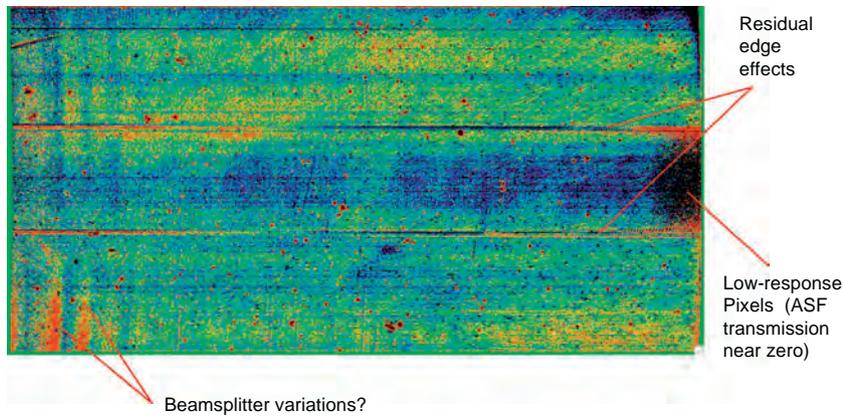


Fig. 41. IR spectrometer flat-field response based on quad-averaged linearization (PDS Version 1) stretched $\pm 2\%$ and color coded. Some residual errors at the boundaries of the ASF filter persist.

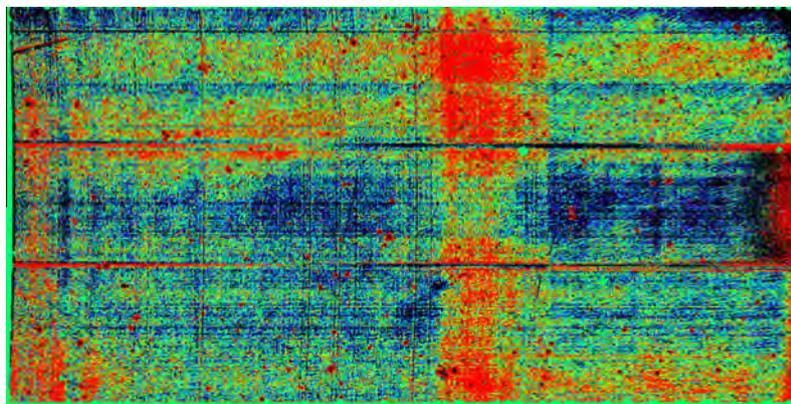


Fig. 42. IR spectrometer flat-field response based on pixel-by-pixel linearization (PDS Version 2) stretched $\pm 2\%$ and color coded.

structured images, which have uniform row spacing, have a $1/3$ -pixel extension introduced at the center of the array. Thus for two features on either side of the midpoint line, the vertical component of the *actual* angular separation between those features is one-third of a pixel less than their measured difference in vertical pixels in the image. As for all geometric distortions, correction of this distortion will require resampling of the image and an attendant loss in spatial resolution. The standard pipeline process does not perform this correction so as to preserve the best spatial resolution.

Determinations of the relative alignments of the boresights of the three EPOXI instruments were made using a variety of data sets spaced throughout the extended mission. Simultaneous images of a star target with both cameras and images taken while scanning the IR spectrometer slit across a star target were used. Stars observed were Vega in June 2008 and February 2010, β Hyi in October 2009 and May 2010, and a background star when observing Hartley 2 in 2010. Variations in the relative alignments determined were small, typically $<6 \mu\text{rad}$, and no systematic trends were observed. The variations are most likely due to differences in the thermal drifts between the instruments and random measurement error. The mean offsets observed are shown in Fig. 44. Changes from the DI-measured offsets are less than one MRIVIS pixel (or one IR spectrometer binned pixel).

5.2. Spatial resolution

For the EPOCH component of the EPOXI mission, carefully measured, detailed HRI-Visible Point Spread Functions (PSFs) were

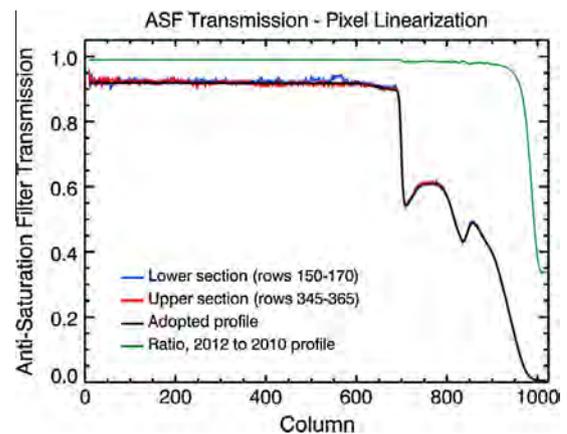


Fig. 43. The transmission of the ASF filter at each column position used for the PDS Version 2 dataset and its ratio to that of Fig. 40.

needed in white light for each of the stars observed for characterization of transiting planets, because the PSF is a function of the specific spectral characteristics of each star. Also, carefully measured PSFs were required in all filters used for Earth observing to permit good image deconvolution. The PSF files developed under EPOXI are archived at the Planetary Data Systems Small Bodies Node, and described in detail at <http://pdssbn.astro.umd.edu/holdings/dif-cal-hriv-6-epoxi-stellar-psfs-v1.0/>.

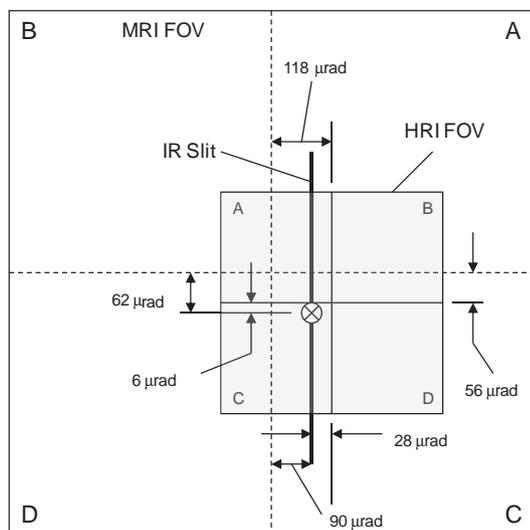


Fig. 44. Relative location of the IR spectrometer slit with respect to the MRIVIS and HRIVIS FOVs.

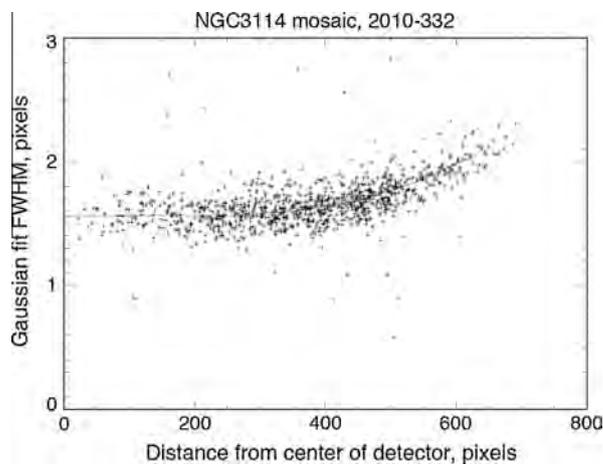


Fig. 46. FWHM of MRIVIS star images as a function of distance from center of the detector for images of NGC 3114 observed on 2010-332.

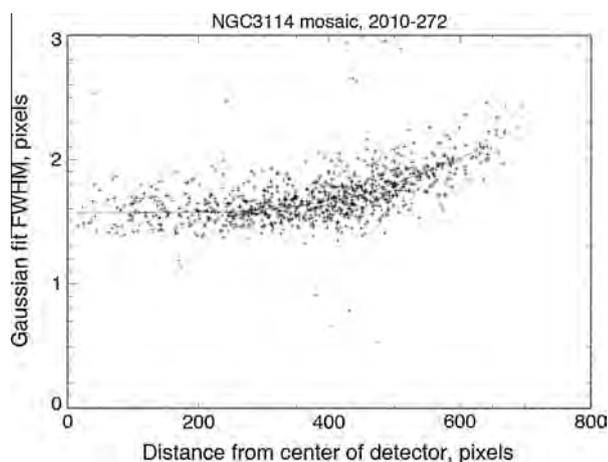


Fig. 45. FWHM of MRIVIS star images as a function of distance from center of the detector for images of NGC 3114 observed on 2010-272.

Before and after the flyby of Comet Hartley 2, CLEAR1 3×3 MRI mosaics of open cluster NGC 3114 were taken on DOYs 272 and 332 of 2010, similar to those taken before and after the bake-out of the MRI during the Deep Impact mission. In these exposures, NGC 3114 was moved to nine different locations in the field of view to provide many star images scattered over the entire field of view of the MRI visible detector. On each day, three of the exposures were found to have significant smear and were eliminated from the analysis presented here. Figs. 45 and 46 show the fitted Gaussian FWHM sizes of star images as a function of distance from the center of the detector for DOYs 272 and 332 of 2010, together with a fitted curve derived from the median values in bins 0–200, 200–300, 300–400, 400–500, 500–600, and 600–800. The distributions are not significantly different on the two days of observation, and there appears to have been no significant change from the distribution of MRI PSF sizes seen after the bake-out of the MRI in 2005.

5.3. Radiometric calibration

5.3.1. Linearity

The response linearity of the charge-coupled device (CCD) detectors to varying light input was monitored throughout the

EPOXI mission using data from the standard calibrations (Table 1). The internal stimulus (STIM) lamps were used as the light source. The light pattern on the detectors from this source shows a gradient in the sample direction, most pronounced on the side of the detector toward lower sample numbers. Therefore, linearity was measured in 200-line by 50-sample subareas spaced across the detectors (14 areas per quad). The STIM lamp output is not actively controlled, and it was seen to vary by up to 0.8% in a few image sets. These variations are thought to be a result of changes in the S/C bus voltage, since the STIM is powered directly off the heater bus, which comes directly from the S/C bus without any power conditioning. Dark subtractions and corrections for frame-transfer smear were done prior to performing linear fits to the response vs. integration time for each subarea. Fits were derived to minimize the percentage errors in a least-squares sense.

Linearity remained quite good throughout the mission; however, the residuals did increase slightly with time. Prior to 2010, the linearity residuals were typically <0.5% and always <1% for signal levels in the range from 70 to 14,300 DN above the bias; residuals grew to as much as 1.5% below 70 DN for MRIVIS. Starting in 2010, the residuals increased to <1% typically and always <1.5% for signals in the 100–14,300 DN range with residuals as large as 1.9% below 100 DN. Fig. 47 shows a typical linearity calibration result from September 2010. The signal levels are systematically below the linear fit at the low and high signal levels and above the linear fit for mid-range signals. These results are consistent with those from the primary mission. The residual outliers at about 4000 DN signal are attributed to a random fluctuation in STIM output for one image in this set.

5.3.2. Gain/full well

The gain and full-well levels of the VIS cameras were tracked throughout the EPOXI mission by applying the photon transfer techniques to dark and STIM frames gathered during each standard calibration activity (Table 1). Signal and random noise levels were determined for the same subareas as were used for the linearity determinations. A typical photon transfer curve result is shown in Fig. 48.

The gain and full-well levels for both cameras have remained quite stable throughout the mission. The gain factors for both cameras have increased slightly from about $28 \pm 2 e^-/DN$ during the prime mission to about $29 \pm 2 e^-/DN$ during the EPOXI mission. Knowledge of the true gain is not very critical; the only use of this value is in estimating the SNR in each pixel. For absolute radiometric calibration, the gain term is folded in with the CCD quantum

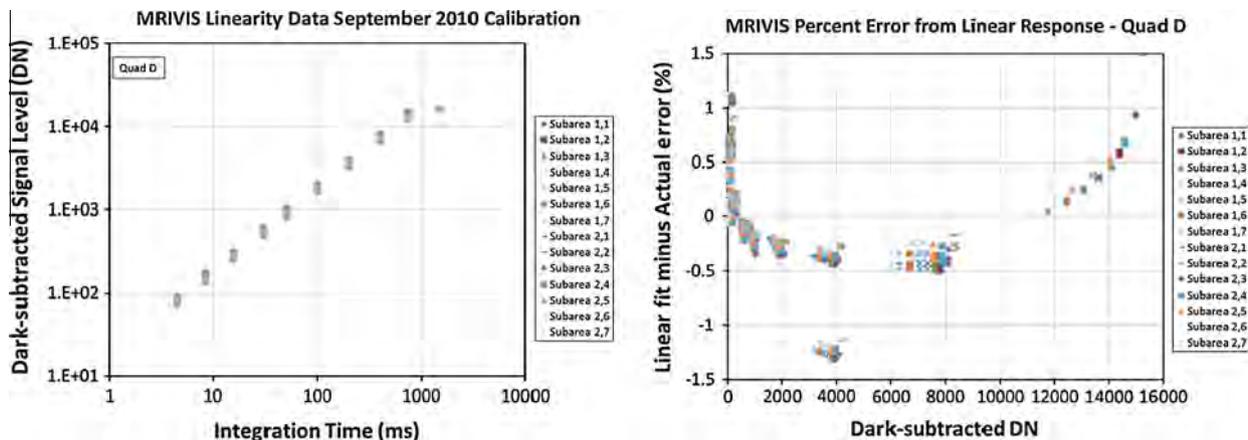


Fig. 47. Typical VIS CCD response linearity and residuals.

efficiency and the pixel area, and it is only the product of these that needs to be known accurately.

The conservatively defined full well of the VIS CCDs has remained constant at about 14,000 DN ($\sim 400,000 e^-$) during EPOXI, which is slightly higher than the 12,000 DN level measured during the prime mission.

5.3.3. Zero exposure background

The bias levels of both VIS cameras, as determined for each quadrant from the serial overclocked pixel values, remain consistent with those from the DI primary mission for comparable CCD and pre-amp temperatures. The dark current rates also remain the same – on the order of 0.01 DN/s or less at the normal CCD operating temperatures. Bias and dark current subtraction routines remain unchanged in the processing pipeline from those used for the DI primary mission.

However, the bias and dark current subtraction of single values for each quadrant used during DI leaves multi-row offsets (stripes) with an amplitude as high as 4 DN per pixel in the backgrounds of MRIVIS and HRIVIS images. The stripes are local to the quadrants of each image, i.e., they do not cross quadrant boundaries. The stripes can be corrected leaving an improved background with an amplitude of a few tenths of DN per pixel, and a routine has been added to the pipeline to make this correction. In the most direct approach, which can be used when no external source signal has been produced in a frame, the average of each row of each quadrant can be subtracted, resulting in a statistically smooth background. However, real sources (e.g., stars and comets) complicate the row-by-row averaging and affect the end result. Therefore, the de-stripe procedure in the VIS pipeline first tests the image for bright sources and then continues along one of three branches depending on if the image is background dominated, source dominated, or an intermediate case.

A background pixel mask is used to define areas that are likely to be purely background. Those pixels that are >1.6 DN from the bias level are assumed to be sources. For a given row, the residual bias is measured with respect to the mean value of the overclocked pixels of the five closest rows (i.e., ± 2 rows). Fig. 49 shows the result of the mask procedure on an HRIVIS star field.

Once the mask is generated, the second phase of the de-stripe procedure computes the strength of the stripes based on the unmasked pixels. However, faint sources or extended surface brightness may still be present. The procedure takes a conservative approach and considers the unmasked pixels to be background dominated if their mean is <1.6 DN, on a row-by-row basis. The de-stripe procedure averages the entire quadrant background, row-by-row, to determine a column profile of stripe offsets.

If the column profile mean is >1.6 DN, the procedure assumes that a significant portion of the image is still contaminated by sources. In this second case, the procedure takes the minimum of the columns close to the image edges, on a row-by-row basis, to determine the stripe offsets within each quadrant. The number of columns examined depends on the size of the image, and has been conservatively chosen to yield useful statistics (e.g., 32 columns for a 1024×1024 pixel image, 16 columns for 256×256).

If the columns close to the image edges are also dominated by source flux (again, mean pixel value >1.6 DN), the stripes cannot be easily and quickly computed in the pipeline, and the image is returned unmodified. However if the columns close to the image edges are not significantly dominated by source flux (mean pixel value <1.6 DN), then the pipeline computes the stripes and subtracts them from the image.

In the background-dominated and source-contaminated cases, the computed offsets are subtracted from the image, excluding the serial overclocked columns. After computing the stripe correction, the de-stripe process attempts to preserve the original image background level by averaging all the quadrant column profiles together and adding that sum back into the image.

When executed, the de-stripe removal is applied to all rows of the image, including, if present, the parallel overclocked rows at the top and bottom of the image. This step affects the readout smear correction procedure (Klaasen et al., 2008), as it depends on the magnitude of the parallel overclocked pixels. If the parallel overclocked rows were not modified, the readout smear correction procedure, executed after the de-stripe procedure, would propagate any quadrant-by-quadrant bias residuals back into the image. Removing the residual bias (including putative striping) in the parallel overclocked rows greatly reduces quadrant offsets at the readout smear correction step.

The de-stripe correction in the pipeline is recorded in a $2 \times N$ pixel FITS extension, where N is the number of rows in the image (e.g., 1024 or 512). The extension has units of DN. The first column of the de-stripe extension is the 1-dimensional array of corrections subtracted from image columns $\leq M$, where M is the middle column of the image (i.e., 512 for a 1024-column image, 1-based indexing). The second column of the de-stripe extension was subtracted from image columns numbered $>M$. As mentioned above, the serial overclocked columns are not modified by the de-stripe procedure. If the de-stripe correction was not applied, then all values in this extension are set to 0.

An example of a stripe-corrected image from DOY 260 of 2010 is provided in Fig. 50. In this CLEAR filter image, the standard deviation of the original image column profiles (after masking sources) is 0.61 DN. After the stripes are removed, the standard deviation is

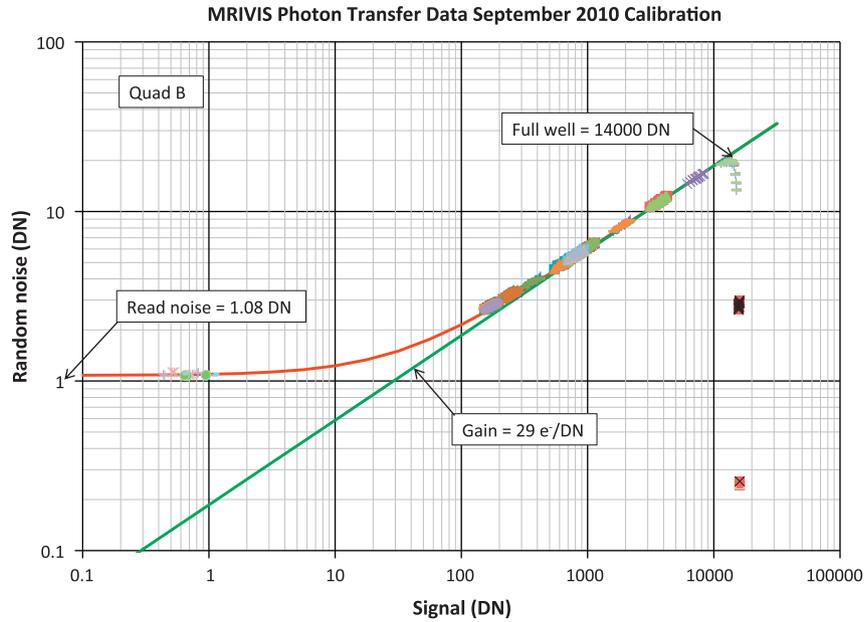


Fig. 48. Typical VIS photon transfer curve.

reduced to 0.18 DN. Due to the selection criteria for destriping, only a relatively small percentage of the images had this process applied to them: 17% of MRI images and 31% of HRI images from the Hartley 2 encounter (2010 DOYs 246–332).

A significant anomaly in the MRI background level was encountered during the approach to Hartley 2. From encounter (E)-60d to E-40d, MRI images of Hartley 2 were being taken through the clear and the CN (387-nm) filters. Between E-56d and E-42d, the observed signal level in the CN filter gradually increased to about 8 DN above normal and then decreased to its original level again. The enhanced signal was nearly uniform across the entire frame (out to >150,000 km from the nucleus) at every time step, which is not explained by any plausible natural phenomenon related to comets. The clear-filter images showed a similar increase peaking at about 1–2 DN. The CN exposure times were 480s compared to the clear exposures of 60s, roughly the same ratio as the peak signal levels. While at first it was speculated that this observation reflected an unknown cometary process (A'Hearn et al., 2010, 2011), further analysis led us to suspect that the signal increase was due to some instrumental effect.

Three hypotheses for the cause of the “CN anomaly” were offered: (1) an electronic effect in the signal chain, (2) stray light reaching the detector from some source behind the filter wheel, and (3) spurious glow from the STIM lamp LEDs. During the anomaly period, the S/C was observing Hartley at an atypical attitude with respect to the Sun illumination direction. The instruments were not fully shaded by the S/C solar panels during this time. The geometry was gradually changing throughout the period, with the solar phase angle of Hartley increasing from about 36° at E-60d to 50° at E-40d. Between E-58d and E-42d, the HRI was powered off and then back on due to elevated temperatures in some spacecraft telecomm components. Several MRI electronics temperatures shifted by about 7°C from their usual values (3 measured colder, one measured warmer) and then back again in conjunction with the HRI power switches and correlated with the “CN anomaly” period. However, we would expect that any electronically induced variation in background signal due to changes in electronics temperatures would also affect the overclocked pixel values used in bias subtraction. So while something instrumental was certainly

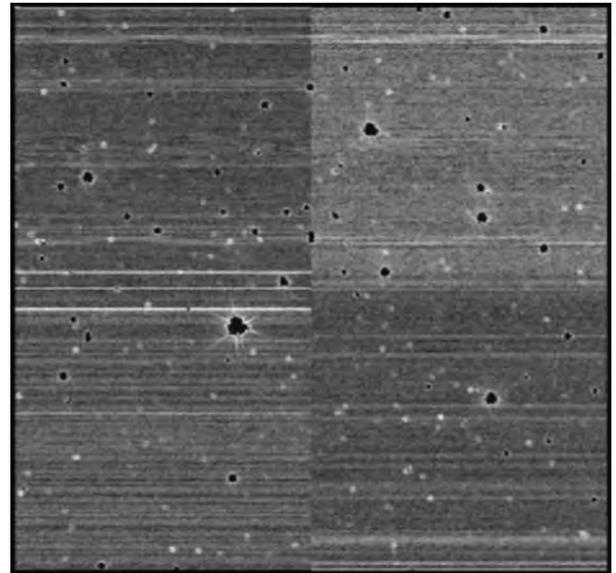


Fig. 49. Residual bias in an HRIVIS frame after the mask generation process showing the horizontal stripes that remain after quad-level bias subtraction. Areas in black (previously filled by stars or bad pixels) would not be used for background computation in the stripe removal routine.

affecting the MRI output levels, the normal calibration process should correct for any electronically induced effects, which did not happen. So hypothesis 1 is likely not the direct cause of the “CN anomaly.”

Hypothesis 3 also seems unlikely. Besides not having a very plausible cause for the STIM lamp LEDs starting to glow, the most direct path for STIM light to reach the CCD is to be reflected off the filter (it is normally reflected off the closed light blocker when intentionally activated, but the light blocker is, of course, open during these exposures; see Hampton et al., 2005). The CN filter has a “shiny” side with broadband reflectance and a violet side, so any STIM light would be reflected back onto the CCD more by the CN filter than the clear. But the excess signal scales by the integration

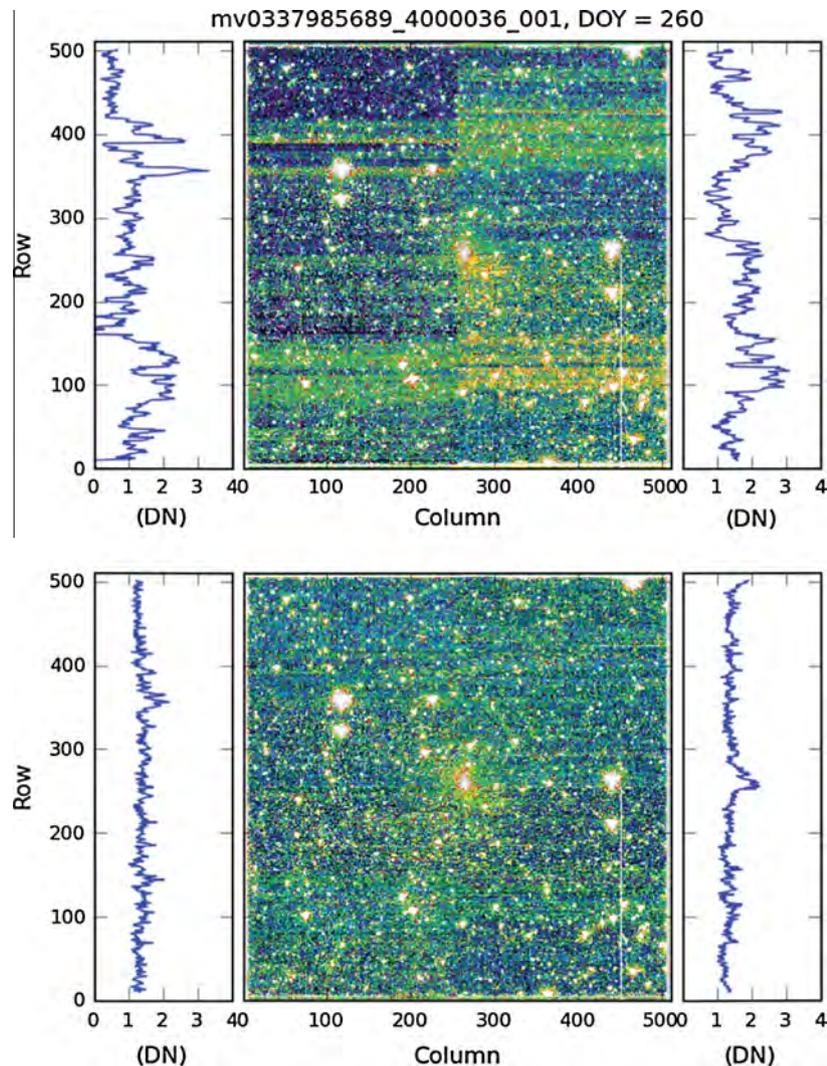


Fig. 50. MRIVIS image before (top) and after (bottom) stripe removal. The image is color scaled from 0 to 4 DN. The profiles on the left and right of each image are row averages for their respective image halves after a simple source masking. The de-stripe process reduces the standard deviation of the image profile from 0.61 to 0.18 DN. Also, note the improved quadrant background continuity.

time ratio between the two filters suggesting that the light falling on the CCD is the same independent of the filter being used.

That leaves hypothesis 2 – stray light. It seems plausible that a stray light path to the MRI CCD could exist (if not originally, perhaps now after 6 years in flight and one close comet flyby exposing the instrument to micrometeoroid impacts). Such a path could have a narrow entrance aperture so that light can reach the CCD from behind the filter only for a very narrow range of illumination conditions and that its intensity could gradually increase and then decrease as the illumination geometry changed over the 38–48° range of phase angles that applied during the “CN anomaly.”

In October and November 2011, we conducted three diagnostic tests to try to isolate the cause of the “CN anomaly.” One test involved imaging with the STIM lamp turned on and the light blocker held open, a second involved imaging while the MRI pre-amp temperature set points were adjusted to allow the pre-amp temperature to rise and then fall as it did during the anomaly, and a third conducted imaging with the spacecraft attitude modified to provide the same Sun illumination geometries as applied during the anomaly to check for light leaks. Neither of the first two tests showed any difference in signal level between images taken in the CN and clear filters. However, the light leak test reproduced the signal levels observed in both filters during the CN anomaly

providing convincing evidence that the anomaly was caused by a small light leak that allowed some sunlight to enter the instrument for a very specific illumination geometry not normally used for science observations.

Bottom line – the “CN anomaly” was caused by an instrumental effect rather than a real observation of transient enhanced cometary CN (A’Hearn et al., 2012).

5.3.4. Absolute spectral response

Reduction of data obtained during the primary and extended missions revealed inadequacies and errors in the calibration of the absolute spectral response of the three shortest-wavelength MRIVIS filters (309, 345 and 387 nm). The CCD quantum efficiency was not well constrained at wavelengths shorter than 400 nm, and this probably affected the earlier calibration results. The MRIVIS narrowband filters were patterned after the Hale-Bopp (HB) narrowband comet filters (Farnham et al., 2000), and this fact was used to provide an updated calibration. The MRIVIS narrowband filter characteristics are very close to those of the HB filters, with bandpasses within the range seen in different batches produced in the various HB sets. Tests performed during the HB filter calibration process showed that the differences between filters would

produce errors in the calibration of less than 10%, and it is expected that the MRIVIS narrowband filters will have errors at this level.

To update the absolute fluxes for the calibration stars viewed by EPOXI, the HB standard star measurements were used. These standard stars consist of late O- or early B-type stars (to minimize the effects of strong emission lines) and solar analogs (selected for removing the solar reflectivity from comet observations). These stars were intensely observed during the HB filter calibration, and the fluxes captured by the narrowband filters were well characterized. Some of these stars were observed as EPOXI calibration stars, while others with the same, or close, spectral types can be used by scaling their fluxes via their V magnitudes. Table 7 lists the EPOXI calibration stars, along with the HB standard stars that were used in the calibration described here.

For each of the EPOXI calibration stars, the expected flux in each filter was computed using the information presented by Farnham et al. (2000) and, if necessary, scaled by the V magnitude. These computations were done for all five of the MRIVIS narrowband filters, to keep the process consistent and to provide a check on the calibrations where the CCD QE was well determined.

The calibration stars listed in Table 7 were observed in each standard EPOXI cruise calibration; the derivation of revised MRIVIS calibration constants presented here is based on the September 2010 calibration. Photometry measurements were obtained for each observation and normalized to total DN/s. Comparing the photometry to the expected fluxes in each filter produced a conversion coefficient for each measurement as shown in Fig. 51.

Some scatter exists between the measurements of different stars in the 309- and 345-nm filters, though the measurements of each individual star are very consistent (there are eight measurements in each star/filter combination, and often they fall within the point size). This scatter probably arises from differences in the filter transmissions from those used in the HB calibrations (including the effects of small red leaks) and from differences in the stars (which may have different levels of reddening from the HB standards). The 387-nm filter shows significantly more scatter, which comes primarily from the fact that the filter straddles the Balmer discontinuity, and so is very sensitive to differences in the filter bandpass. Because of this scatter, it was debated whether the data should be weighted in some manner to improve the results (e.g., give lower weight to stars that are not an exact spectral match, etc.). Tests of different weighting scenarios showed that even in the extreme cases, the calibration coefficients would decrease by at most 4% or increase by at most 6%, depending on which stars were de-emphasized. Since it was not clear what the best weighting should be, and since it had little effect in any case, a straight average of all of the measurements was used to derive the calibration coefficients.

The adopted gray radiance conversion constants are given in Table 8. It should be noted that the 514- and 526-nm filter coefficients are almost identical to the coefficients derived in the earlier calibration (Klaasen et al., 2008). This result provides confidence that the measured QE at longer wavelengths is accurate.

Having the improved coefficients for the three short wavelength filters provided an independent means of measuring the CCD QE at those wavelengths, which in turn, would improve the calibration of the Clear-1 and Clear-6 filters. Adopting the measured optical properties for the narrowband filters and the procedures described in Section 5 of Klaasen et al. (2008), it was possible to derive the effective QE for each of the narrowband filters. Because the filters are narrow, the QE can be assumed to be constant across the band and adjusted until the radiance conversion for each filter matches the value derived using the procedures described above.

With the QE constrained at three points in the UV, it was assumed that the QEs at wavelengths longer than 500 nm were correct and that the QE at 200 nm was zero, and a smoothed spline

curve was fit through the data to update the QE at all wavelengths. With this new QE, the procedures described in Klaasen et al. (2008) were used to re-derive the gray radiance conversion constants, as well as the I/F calibration constants for the narrowband and the Clear-1 and Clear-6 filters. The calibration constants for all of the MRIVIS filters are given in Table 8. Although the three short-wavelength filter calibrations changed significantly, the Clear filter constants only changed by about 3%, and the constants for the filters above 500 nm remained essentially unchanged. The system response based on these new calculations is shown in Fig. 52.

Finally, using the new calibration coefficients, the ratios of fluxes measured from a solar spectrum between the different filters were computed. These ratios can be used as scaling factors for using the Clear and continuum filters to remove the continuum flux contribution from the gas filters. These coefficients are provided in Table 9.

The stability of the VIS cameras' absolute sensitivities through each spectral filter was monitored during EPOXI at each standard calibration (Table 1) using the same set of standard calibration stars that was used during the primary mission (Klaasen et al., 2008). Figs. 53 and 54 plot the ratios of the total star signals in the various DI and EPOXI calibrations to those of the May 2005 calibration. The results for each standard calibration are offset by one unit for visibility. In all cases, dark subtraction, crosstalk removal, readout smear removal, and flat-fielding have been done, and obvious cosmic rays within the photometry aperture have been removed.

The figures show that the absolute radiometric spectral response of the cameras remains nearly unchanged from that measured during the primary DI mission. The scatter in the measurements is at about the 10% level, as for DI. The scatter within the 5- to 8-frame repeat measurements for a given star in a given filter for a given calibration is typically <1% for HRIVIS and <2% for MRIVIS. The differences from the May 2005 measurements are less than the absolute calibration uncertainties for each filter claimed in the DI calibration paper (Klaasen et al., 2008).

The only systematic changes observed are:

- An increase in HRIVIS response in the 950-nm filter of about 6% for the July 2005 and both 2008 calibrations, which was replaced by a decrease in this filter's response of about 8% for the 2010 calibrations. As a result, the radiometric conversion constant for this filter has been changed in the processing pipeline from the DI value of $1.931 \text{ (W m}^{-2} \mu\text{m}^{-1} \text{ sr}^{-1})/\text{DN/ms}$ (Klaasen et al., 2008) to 1.822 for data acquired between 2007 and 2009 and to 2.085 for data acquired in 2010.
- A systematic increase in response with wavelength to star HD60753 for the HRIVIS in June 2008. The HRIVIS response to this star only (not any of the others) is elevated 3–5% for filters with effective wavelengths below 600-nm. The response rate increase grows to 8% at 645 nm, 14% at 743 nm, 21% at 842 nm, and 37% at 948 nm. It seems as if this star's output changed for this calibration – its magnitude increased slightly and it became more red. With the

Table 7
DI standard stars and their associated Hale–Bopp standards.

EPOXI star	Spec type	V mag	HB star	Spec type	V mag
Achernar	B6V	0.50	HD 37112	B6V	8.00
i Car (HD79447)	B3III	3.90	HD 209008	B3III	5.99
HD 60753	B2III	6.68	HD 120086	B2V	7.82
β Hyi	G0V	2.80	16 Cyg A	G1.5V	5.96
16 Cyg A	G1.5V	5.96	16 Cyg A	G1.5V	5.96
16 Cyg B	G2.5V	6.20	16 Cyg B	G2.5V	6.20

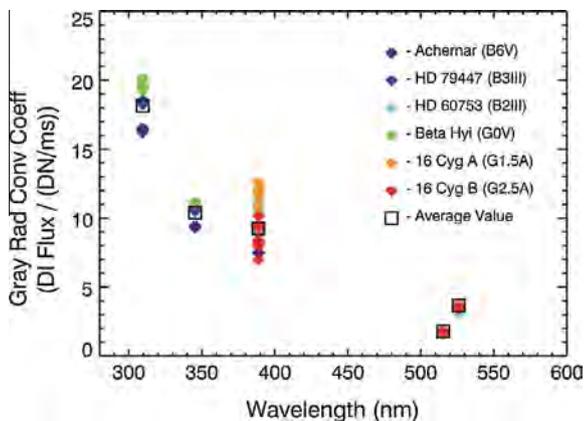


Fig. 51. Computed gray radiance conversion coefficients, computed from EPOXI observations and the expected fluxes as computed from the HB standard star measurements. The averaged coefficient for each filter is shown by the black box.

Table 8
Gray radiance and *I/F* conversion coefficients for MRIVIS filters.

Filter	Gray radiance ^a	<i>I/F</i> ^b
Clear-1	0.03527	7.722×10^{-5}
Clear-6	0.03531	7.730×10^{-5}
309-9	18.14	9.197×10^{-2}
345-8	10.40	3.587×10^{-2}
387-7	9.239	2.810×10^{-2}
514-2	1.789	3.005×10^{-3}
526-3	3.666	6.174×10^{-3}
750-4	0.2125	5.223×10^{-4}
950-5	0.5879	2.258×10^{-3}

^a (W/m²/μm/sr)/(DN/ms).

^b *I/F* at 1 AU/(DN/ms).

exception of the June 2008 calibration, the response rates for this star have always been consistent with those from the other reference stars we are using. None of the star catalogs include any information on the variability of this star.

- (c) Beginning in 2010, the MRIVIS response is consistently about 5% lower than seen previously across all stars and all filters. Therefore, all the radiometric conversion constants for the MRIVIS have been decreased by 5% in the processing pipeline from the EPOXI values for all data beginning with calendar year 2010. Table 8 conversion constants are 5% lower for data taken prior to 2010.
- (d) Beginning in September 2010, the HRIVIS response is consistently about 3% lower than seen previously across all stars and all filters (excluding the 950-nm one discussed above). Therefore, all the radiometric conversion constants for the remaining HRIVIS filters have been increased by 3% in the processing pipeline from the DI values for all data beginning with September 1, 2010.

5.3.5. Flat field

Flat-field corrections are required in both VIS cameras to correct for optical and electronic effects. A different flat field is required for each filter and each mode of the instruments. Flat fields were created for the DI mission (Klaasen et al., 2008). During the EPOXI standard cruise calibrations, we took several images with the internal calibration lamp (“STIM images”) that we compared to similar images taken in June 2005, during the last pre-encounter Deep Impact cruise calibration phase. The STIM images provide direct information on the pixel-to-pixel variations of the CCD, since the lamp is located behind the filters in the optical path.

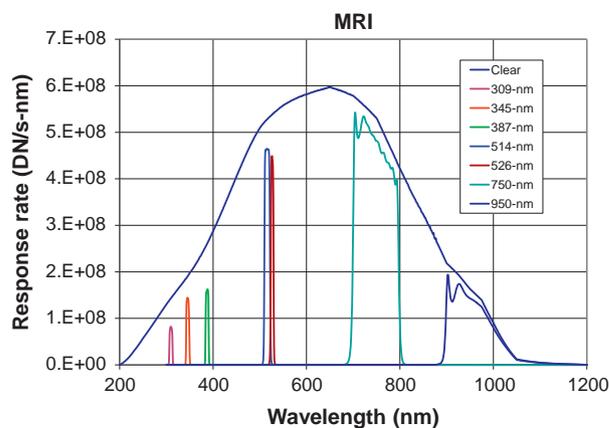


Fig. 52. MRIVIS system response, based on the quantum efficiency computed from the HB standard star measurements.

Table 9
MRIVIS continuum removal coefficients.

	OH	CN	C2
CLEAR1	0.432	0.720	1.281
UC	0.680	1.134	–
GC	0.332	0.554	0.986

Under the assumption that all CCD pixels respond linearly with increasing light input, we noticed that the response of 6.7% of the HRIVIS pixels (4.1% for MRIVIS) had changed by more than 1%, and 0.2% (none for MRIVIS) by more than 3% between June 2005 and January 2008. In the June 2008 calibration, which involved only HRIVIS, 7.5% of the HRIVIS pixels had changed by more than 1% from June 2005. The same analysis was repeated again in February 2010, at which time the response of 9.0% of the HRIVIS pixels and 7.7% of the MRIVIS pixels had changed by more than 1% from June 2005. Only 118 HRIVIS pixels and 163 MRIVIS pixels (over 1024² pixels) had changed by more than 5%. This level of change is regarded as small; however, we decided to implement updates to the DI flat-field corrections in the calibration pipeline that take into account these pixel-to-pixel variations since June 2005.

The flat-field calibration files were updated by multiplying each file by a “corrective image.” The “corrective image” is the change in pixel-to-pixel response between February 2010 and June 2005, which is the normalized ratio of the response vs. integration time for each pixel derived from STIM images taken at these two different times, such that:

$$\text{New_File} = \text{Old_File} * (\text{STIM_February_2010}/\text{STIM_June_2005})_{\text{norm}}$$

In creating these “corrective images” 32 frames from each calibration were used – 4 taken at each of 8 different integration times. They were dark subtracted, and pixel-by-pixel median values were computed over each 4-frame set to reduce random noise and eliminate any cosmic ray effects. Linear slopes were fit for each pixel across the median values for the 8 integration times (excluding any with saturated values). The slope value at each pixel forms the new flat-field map. We have created “corrective images” for each mode: 1024 × 1024, 512 × 512, 256 × 256, 128 × 128 and 64 × 64. Use of these updated flat-field files, as demonstrated in Fig. 55, indeed provides a better correction of the pixel-to-pixel variations.

This EPOXI update was then tested successfully in September 2010 and November 2010, during the last two standard calibrations. Although the response of a few pixels had changed by more

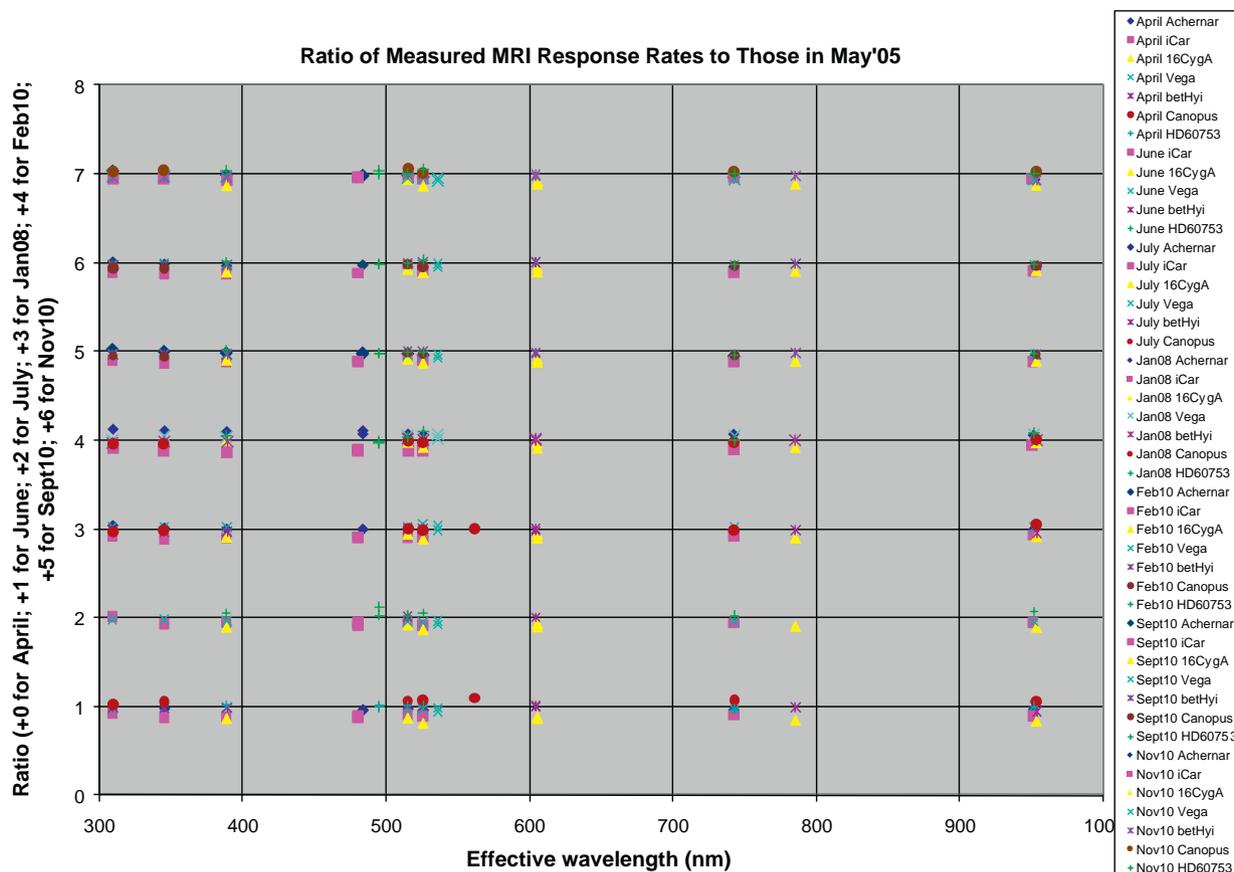


Fig. 53. Ratio of total star signals in the MRIVIS camera through each filter for the standard calibration stars over all calibrations to those from May 2005; plotted vs. filter effective wavelength.

than 1% since February 2010 (2.0–2.1% of the pixels for HRIVIS, 2.1–2.6% of the pixels for MRIVIS), we do not consider this change significant and have retained the February 2010 corrected flats in the calibration pipeline. This decision is further reinforced by the fact that only 8 HRIVIS pixels and 21 MRIVIS pixels had changed by more than 5% since February 2010.

Finally, the EPOCH team derived a modification to the HRIVIS flat-field function over the central 128×128 -pixel area of the detector (Ballard et al., 2010; Christiansen et al., 2010,2011) consisting of a two-dimensional spline fit. While this second-order correction provided a significant reduction in the noise in the EPOCH photometry measurements, its magnitude is quite small (<0.5%), the fits vary with the spectral type of the star used to create them and with time, and the corrections are not significant for most other types of photometric data reductions. Therefore, this correction has not been incorporated into the EPOCH data processing pipeline.

The two 1/6-pixel narrower central rows collect only 5/6 of the charge of a normal row. This effect is corrected by the flat-field division so that the pixels in these rows have the correct scene radiance assigned to them. However, point-source or disk-integrated photometric measurements using aperture photometry areas that include these central rows will be slightly distorted unless special adjustments are made (e.g., an extra 1/6-pixel worth of signal is added to the pixels in each of these two rows in the reconstructed images, as described in Appendix A of Belton et al., 2011).

5.3.6. Noise

The random noise in the VIS cameras' CCD readout was determined using the photon transfer technique (see Fig. 48). The

read-noise levels have remained stable since before launch at roughly 1 DN (1σ). Since the read-noise level is comparable to the encoding quantization step size, its measured value can vary by several tenths of a DN as the mean bias level on which the noise rides shifts slightly. Measured read-noise values have ranged from 0.6 to 1.8 DN over time with no systematic trends apparent.

The VIS cameras exhibit electronic cross talk, in which images with significant and structured signal in one quadrant show mirrored, low-intensity images in other quadrants (Klaasen et al., 2008). In order to monitor this effect, we included a specific observational sequence in each of the cruise calibrations. The sequence is an observation of Canopus saturated about 100 times above full well with Canopus located in each quadrant in turn. From these observations, we were able to derive the fraction of signal (called the gain) that is reflected from one quadrant to the other for every combination of quadrants. The results are given in Table 10 for HRIVIS and Table 11 for MRIVIS for each cruise calibration since 2005.

The variations in measured gain between the different cruise calibrations are not significant, and for the calibration pipeline we decided to use the average value over the different EPOCH cruise calibrations. Depending on the combination of quadrants and instruments, the gain varies from 0.02% to 0.09%, which corresponds to 2–9 DN for a 10,000-DN source.

Fig. 56 illustrates the results before and after the cross-talk correction. Before the correction, we see Canopus (saturated) in the upper right quadrant and an electronic ghost in the other three quadrants. After the correction, the electronic ghosts nearly disappear from these three quadrants. The ghost signals were reduced in this case from between 8 and 14 DN to 1–1.5 DN.

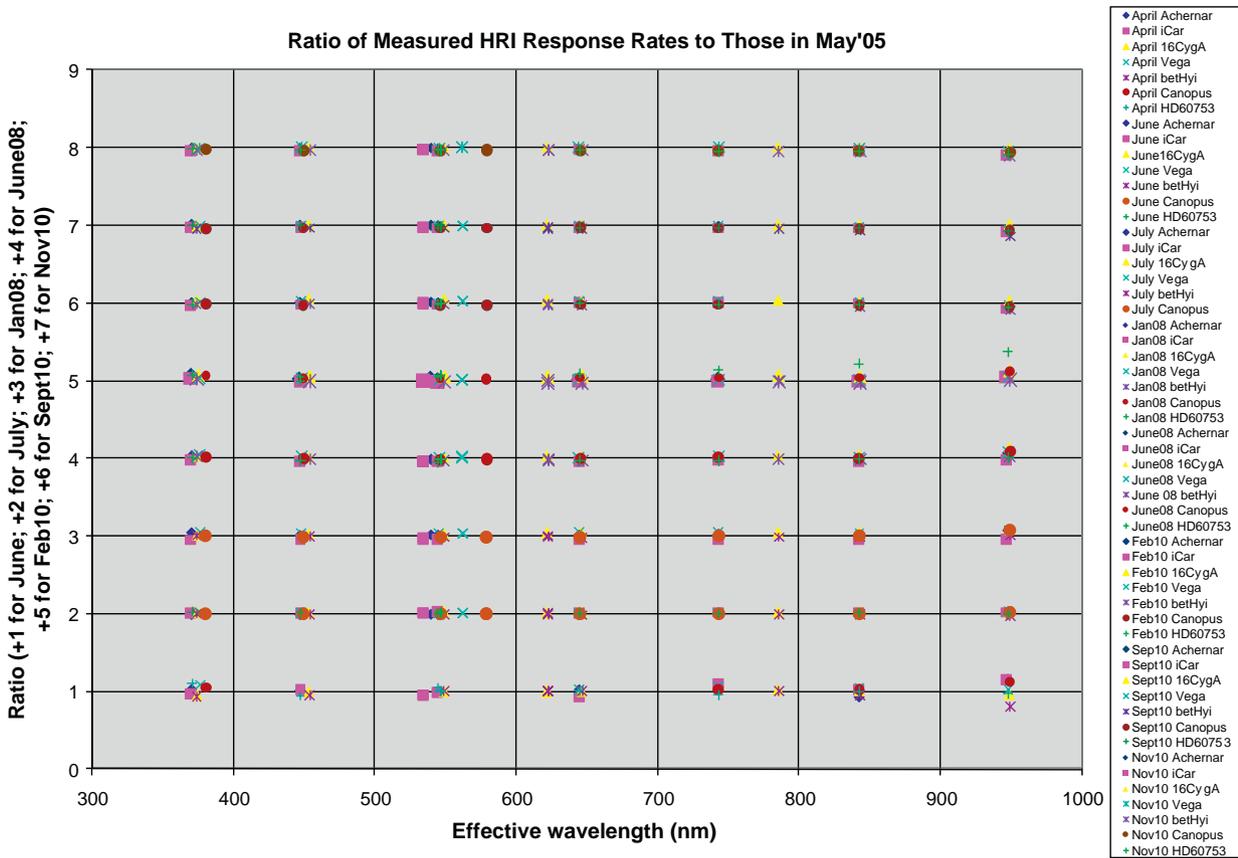


Fig. 54. Ratio of total star signals in the HRIVIS camera through each filter for the standard calibration stars over all calibrations to those from May 2005; plotted vs. filter effective wavelength.

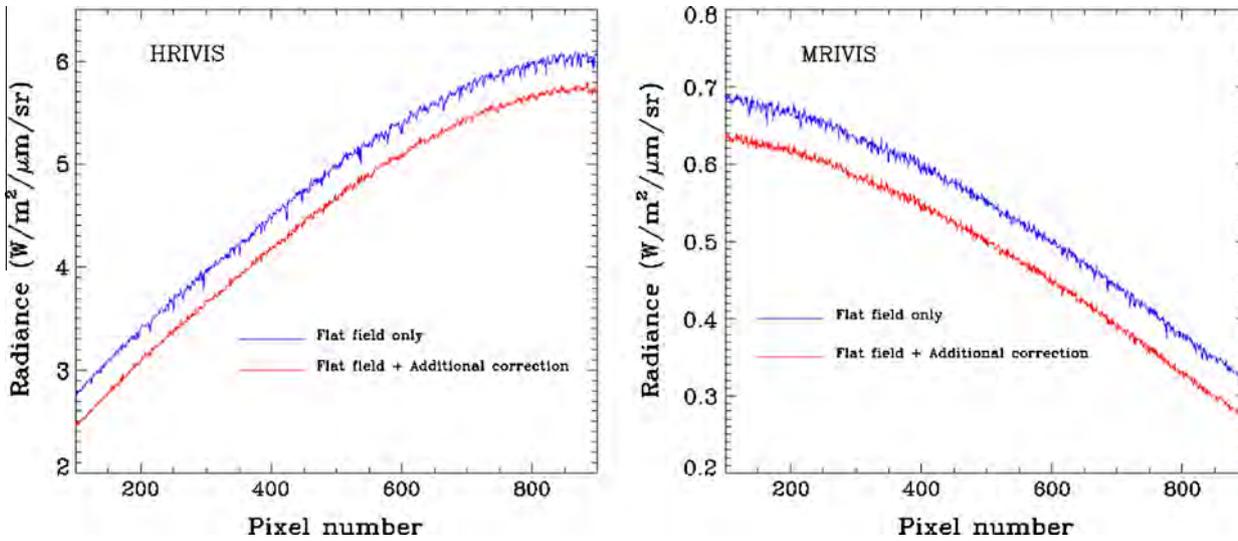


Fig. 55. Comparison of a horizontal profile (row #500) of a STIM image for HRIVIS and MRIVIS, with the flat-field correction as implemented for the Deep Impact calibration pipeline in blue and the updated correction incorporating the changes in pixel-to-pixel response between June 2005 and February 2010 in red (as implemented for the EPOXI calibration pipeline). The update improves the data reduction, since it reduces the pixel-to-pixel variations. The blue and red curves are offset vertically for clarity.

Another issue arose in dealing with compressed data. Because of the voltage bias added to the detector electronics after pre-launch calibration, the lowest 14-bit value expected from the ADC in any VIS quadrant is a few DN over 350, and the zero entry of the 8-bit LUT on the spacecraft was set to represent 14-bit DN from zero to 350. That is, for compressed data, any 14-bit DN pixel reading between zero and 350 will be downlinked as an 8-bit zero.

During the EPOXI mission some 8-bit zeroes were downlinked, and the calibration software decompressed them to a value of 175 DN, the midpoint of the 0–350 DN range represented by LUT entry zero. This assignment had repercussions in the calculation and subtraction of bias background values: any 175-DN values in the bias pixels were discarded by the resistant mean bias calculation, which shifted the calculated bias upward; any 175-DN values in the ex-

Table 10

Cross-talk gains (10^{-4}) for HRIVIS. The gain is the fraction of the signal that goes from one quadrant (A, B, C or D) to the other.

HRIVIS	DI June 05 (10^{-4})	EPOXI January 08 (10^{-4})	EPOXI June 08 (10^{-4})	EPOXI February 10 (10^{-4})	EPOXI September 10 (10^{-4})	EPOXI November 10 (10^{-4})	Average EPOXI value (10^{-4})
B over A	4.0	3.5	3.5	3.5	3.5	3.5	3.5
C over A	3.0	3.5	3.0	3.0	3.5	3.0	3.2
D over A	3.5	3.5	3.0	4.0	4.0	3.0	3.5
A over B	3.5	3.5	3.0	3.0	3.0	4.0	3.3
C over B	8.0	6.0	7.0	7.0	7.0	9.0	7.2
D over B	4.0	3.5	3.5	4.0	3.5	4.0	3.7
A over C	3.0	3.0	3.0	3.0	3.0	3.0	3.0
B over C	8.0	8.0	7.5	7.5	8.0	8.0	7.8
D over C	5.0	8.0	5.5	5.0	5.0	6.0	5.9
A over D	3.3	3.3	2.0	2.5	3.5	1.0	2.46
B over D	4.0	4.0	4.0	4.0	4.0	4.0	4.0
C over D	4.5	6.0	4.0	5.0	5.0	4.0	5.0

Table 11

Cross-talk gain (10^{-4}) for MRIVIS. The gain is the fraction of the signal that goes from one quadrant (A, B, C or D) to the other.

MRIVIS	DI May 05 (10^{-4})	EPOXI January 08 (10^{-4})	EPOXI February 10 (10^{-4})	EPOXI September 10 (10^{-4})	EPOXI November 10 (10^{-4})	Average EPOXI value (10^{-4})
B over A	6.0	6.0	6.0	6.0	6.0	6.0
C over A	5.0	5.0	5.0	5.0	5.0	5.0
D over A	8.0	9.0	9.0	9.0	9.0	9.0
A over B	5.0	5.0	5.0	5.0	5.0	5.0
C over B	3.0	3.0	3.0	3.0	3.0	3.0
D over B	3.0	3.0	3.0	3.0	3.0	3.0
A over C	4.0	4.0	4.0	4.0	4.0	4.0
B over C	4.0	4.0	4.0	4.0	4.0	4.0
D over C	3.5	3.5	3.5	3.5	3.5	3.5
A over D	9.0	9.0	9.0	9.0	9.0	9.0
B over D	3.5	3.5	3.5	3.5	3.5	3.5
C over D	3.0	3.0	3.0	3.0	4.0	3.25

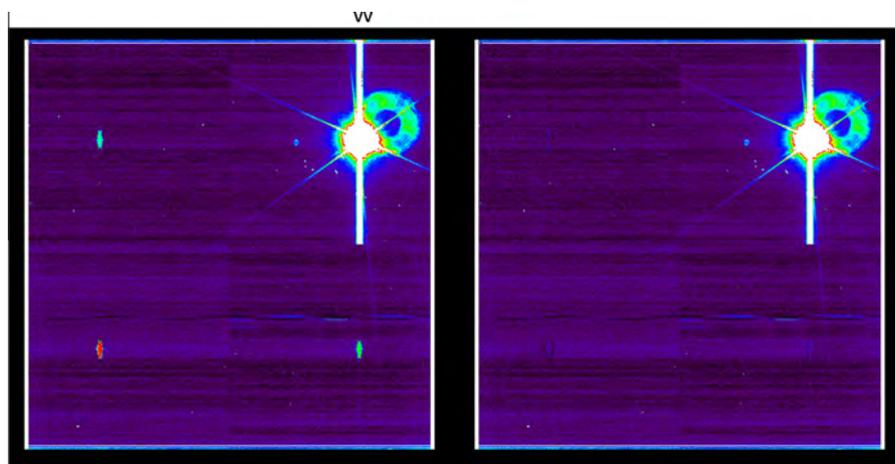


Fig. 56. Saturated HRIVIS image of Canopus (located in quadrant B) before (left) and after (right) cross-talk correction. The electronic ghosts almost disappear. Similar results are obtained for MRIVIS.

posed pixels were calibrated to large negative radiances. A fix for this effect was implemented in the data processing pipeline by decoding compressed 0 DN to 350 DN.

5.3.7. Scattered light

Klaasen et al. (2008) presented some results of scattered light calibrations of both the HRIVIS and MRIVIS cameras from DI using Canopus and the Moon as sources. Similar Canopus calibrations were obtained in each of the standard cruise calibrations during EPOXI; however, we have not yet analyzed these data. In addition, special HRIVIS scattered light calibrations using the Moon were conducted during the December 2007 Earth flyby. In one calibra-

tion (12/17/07), a 3×3 -frame mosaic was obtained in three filters (clear, 350-nm UV, and 950-nm IR). The Moon was fully contained in the frames at the center mosaic position but was just outside the FOV for the other eight bordering mosaic positions. Three exposure levels were used for the center position (on-scale signals plus two levels of saturation), and two different long-exposure levels were used for the border positions. In the second calibration (12/28/07), the Moon was imaged through all nine HRIVIS filters in conjunction with a radiometry calibration.

Quantitative models of the HRIVIS scattered light levels have not yet been derived from these data. We present here some qualitative results. Fig. 57 shows a bias-subtracted UV image of the

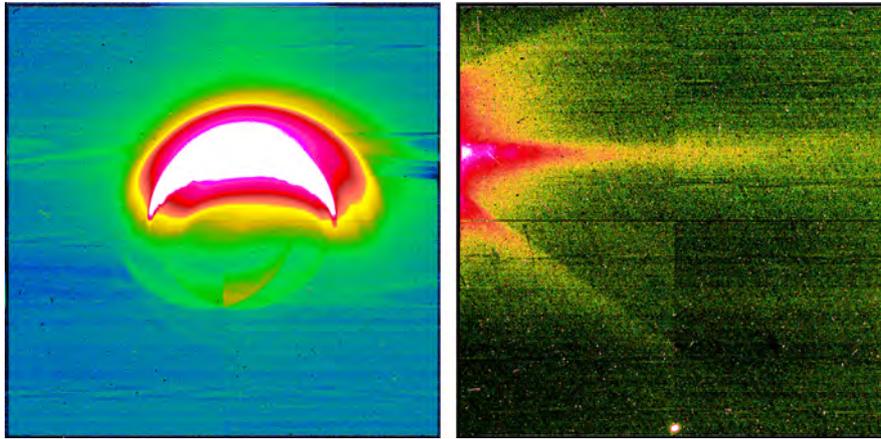


Fig. 57. Long exposures of the Moon through the HRIVIS UV filter logarithmically stretched and pseudo-color coded to show the scattered light with the Moon inside the FOV (left) and just outside the FOV (right).

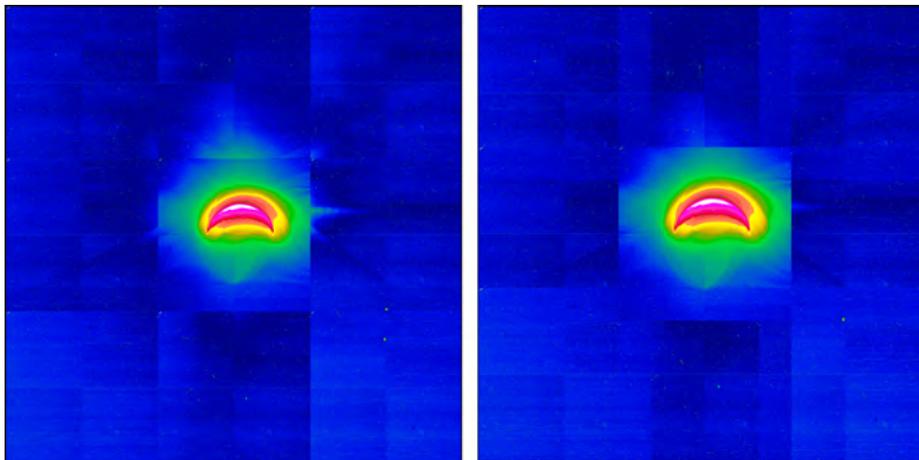


Fig. 58. 3 × 3-Frame mosaic of the Moon in the HRIVIS UV filter logarithmically stretched and pseudo-color coded to show the scattered light patterns. Left – frames laid edge to edge to show complete images; right – frames overlapped with the nominal 20% pointing overlap to position all the scattered light patterns correctly relative to the Moon (center frame is in front of the others).

Moon enhanced using a logarithmic stretch and pseudo-color coded to reveal the low-level signals surrounding the lit crescent. Also shown in this figure is a UV image with the Moon just outside the FOV processed similarly. The scattered light patterns for a source inside the FOV and just outside the FOV are readily seen. The faint mirror image of the lunar crescent in the lower half of the left image is due to low-level crosstalk (see Section 5.3.6). The ray patterns separated by $\sim 45^\circ$ in the right-hand picture are expected and are due to the structure holding the telescope's secondary mirror in place.

Fig. 58 shows the 9-frame mosaic of the Moon in the UV filter. On the left, the frames are positioned in the correct relative mosaic positions but separated so as to eliminate the planned frame overlaps and to allow visibility of the entirety of each frame. The center image is actually a composite of three different exposure levels – one taken with the normal on-scale exposure and the others taken with ratios of $10\times$ and $100\times$ the normal on-scale exposure. The composite replaces saturated pixels in the longest exposure with those from the middle exposure multiplied by $10\times$ and replaces saturated pixels in the middle exposure with those from the shortest exposure multiplied by $100\times$ to produce an image with all pixels having the correct relative on-scale signal levels and expanding the dynamic range of the measurements by $100\times$. The border frames all were taken with the same exposure time as the longest

exposure time used for the center frame. A non-linear contrast stretch was applied to all the frames that greatly enhances the contrast at low signal levels and compresses the contrast at high signal levels to allow the very low-level scattered light effects to be seen along with the very much brighter lunar disk itself. Finally a pseudo-color code was assigned to the output DN levels to enhance the visibility of low-light-level signals. On the right are the same data but now mosaicked with the nominal 20% frame overlap. The center frame lies on top of the edge frames. A comparison of the two mosaics illustrates that the scattered light pattern at a given angular distance from the Moon (i.e., in the frame overlap areas) differs depending on whether the Moon is in the FOV or not. The CCD quadrant boundaries can be clearly seen. Although a dark-frame subtraction was performed on all the frames, no flat-field correction was applied; flat fielding would likely eliminate those quadrant boundaries. However, the scattered light pattern will still cause brightness mismatches at the frame boundaries because the pattern in a given frame depends on the direction and magnitude of the off-axis Moon location, which differs for each frame.

We have quantified the level of scattered light in HRIVIS images with the Moon in the FOV in each filter by extracting brightness scans along columns that cross the lunar limb in the radiometry frames. The Moon's radius in the images was 894 pixels, and the solar phase angle was 99° . Fig. 59 shows the radiance measured

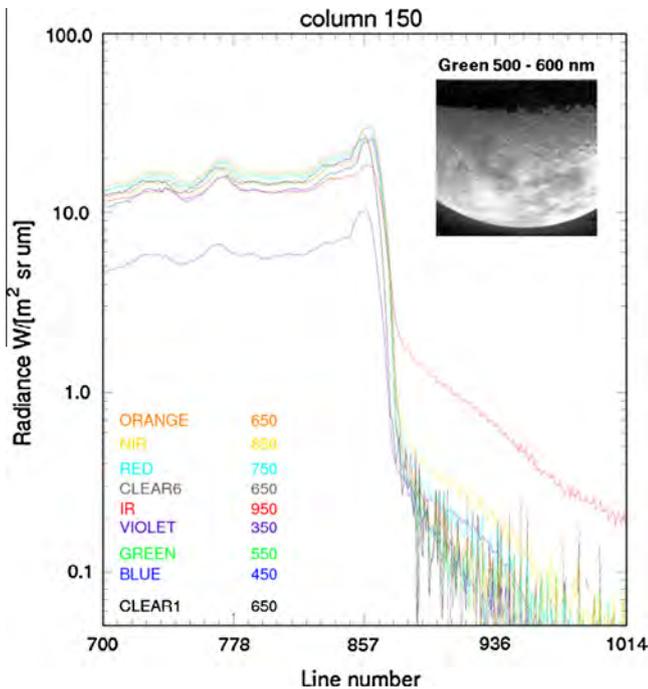


Fig. 59. HRIVIS radiance in each filter along column 150 crossing the lunar limb showing elevated scattered light in the IR, violet, and NIR filters.

along these limb scans for each filter plotted on a log scale. Notice that the 950-nm IR filter shows significantly higher off-limb scattered light than the other filters. The scattered light in the 850-nm NIR and 350-nm violet filter are also somewhat higher than that in the other filters. Fig. 60 plots the ratio of the observed radiance at locations 20, 25, and 50 pixels off the lunar limb to the radiance at the bright limb. At each location the IR, violet, and NIR filters show elevated scattered light levels.

6. Pipeline updates

The updates to the calibrations documented in this paper required changes to the data processing pipeline software installed on the project Science Data Center (SDC) computers at Cornell University, at the University of Maryland (UMd), and at the Jet Propulsion Laboratory (JPL). An updated flowchart illustrating the image processing pipeline used to calibrate EPOXI images is shown in Fig. 61. In addition, science team internal reviews and peer review of previous DI data slated for archive in the Planetary Data System (PDS) brought to light some errors in the existing calibration code that were corrected.

6.1. New files and constants for PDS Version 1

The simplest pipeline updates involved adding updated, or removing obsolete, calibration files and calibration constants that track the time-varying performance of the instruments. These changes were added to the database of previous calibration files along with the applicable time periods for each, and no software changes were usually required for the pipeline to be able to use these files. The IR spectrometer calibration file updates for Version 1 were derived using quad-averaged linearity corrections, as was done for the primary mission. In addition, the new calibration files, stored in FITS format, were given enhanced labels with comments describing provenance, pixel orientation, and how each file is meant to be used in DI data calibrations. These files are docu-

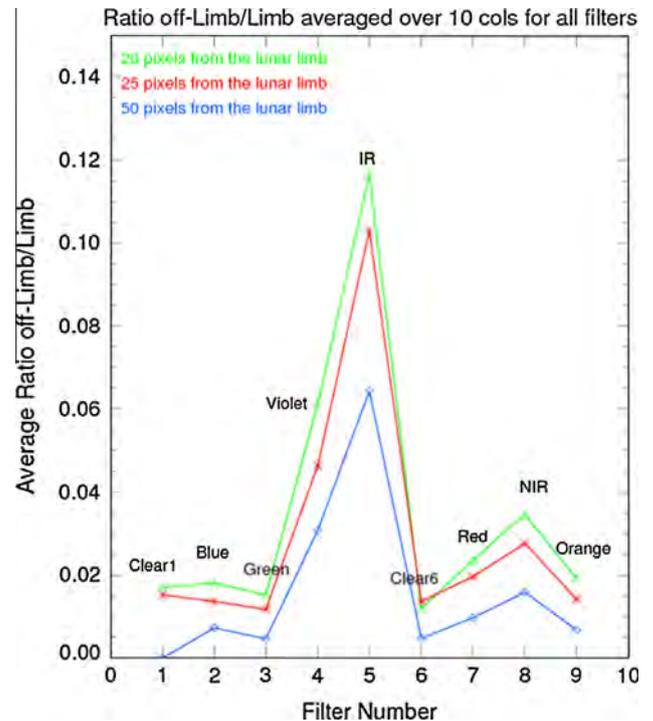


Fig. 60. Ratio of observed radiance at three distances from the lunar limb to the bright limb radiance for all HRIVIS filters.

mented in the PDS archive of the calibration and are summarized here by the first date (YYMMDD: YY = year; MM = month; DD = day of month) of data affected by each set of files:

For PDS Version 1:

- 050112 Revised all MRIVIS calibration constants
- 050112 Updated IR absolute calibration, including 2× correction
- 050112 Updated IR flat fields that include the ASF (Section 4.3.7) and removed previous IR flat fields from calibration database
- 050112 Updated IR spectral smile equation
- 071001 Updated IR quad-averaged linearization coefficients
- 071001 Updated IR master dark file
- 071004 Updated IR bad pixel maps from 01/2008 calibration
- 071004 Updated IR read noise values from 2008 calibration
- 071004 Updated HRIVIS 950-nm filter calibration constant from 2008 calibrations
- 071004 Updated VIS cross-talk files from 2008 and 2010 calibrations
- 080528 Updated IR bad pixel maps from 06/2008 calibration
- 080623 Updated IR quad-averaged linearization coefficients
- 090601 Updated IR bad pixel maps from 06/2009 calibration
- 090618 Updated IR quad-averaged linearization coefficients
- 100101 Updated HRIVIS 950-nm filter calibration constant from 2010 calibrations
- 100101 Updated all MRIVIS filter calibration constants by 5%
- 100201 Updated VIS flat fields
- 100201 Updated IR bad pixel maps from 02/2010 calibration
- 100216 Updated IR quad-averaged linearization coefficients
- 100801 Updated IR read noise values from 09/2010 calibration
- 100901 Updated all HRIVIS filter calibration constants (except 950-nm) by 3%
- 100928 Updated IR bad pixel maps from 11/2010 calibration using 100216 linearity constants

6.2. Compressed 0-DN fix

Because of the issue in decompressing 0 DN compressed VIS data (Section 5.3.6), the pipeline software was changed to decompress the VIS cameras' zero LUT entry, and only that entry, to the top of its range, i.e., to 350 DN. Also, because it could not be ascertained whether the original pixel's 14-bit DN value was actually 350 or a lower value, all such pixels were marked as SATURATED in the FLAGS FITS extension, the rationale being that just as the actual value of a saturated pixel at the upper end of the possible DN range (255 for compressed 8-bit DN's; 16,383 for uncompressed 14-bit DN's) cannot be determined, the same is true for a pixel at the lower end of the range. The user of the FITS file can easily determine the type of saturation of such a flagged pixel by examining its stored value.

6.3. IR transient background correction

A mathematical model for the transient behavior of the IR background signal has been implemented based on the effects of charge trapping (Section 4.3.3.5). Calibrating the IR background transients requires knowledge of the detector reset and read history. To minimize this effect for Hartley 2 encounter data, four IR frames were taken immediately before, and with the same mode and integration time as, each downlinked data sequence. However, these four-frame sets were never saved and downlinked. As such, the frame header information on the ground is incomplete for making the background transient correction. Therefore, planning data from sequencing spreadsheets are parsed by the pipeline software and

used to reconstruct IR detector reset/read history and perform the background transient calibration. The data critical to this analysis will be extracted from the sequencing spreadsheets and archived with the relevant PDS data sets along with an explanation of how to use them. This correction is not applied in the standard pipeline processing but is available as a separate user option.

6.4. VIS stripes removal

Several techniques were developed by the Science Team to deal with residual observation-dependent, semi-coherent stripes in the VIS background data (Section 5.3.3). Calibration procedures to deal with these effects were added to the pipeline software, though not always run in the pipeline by default. An additional extension that saves the stripe-removal adjustments for each line of each quad was added to retain the reversibility of the calibration.

6.5. IR flat-field update

Improved lunar calibration data acquired during EPOXI allowed construction of an IR spectrometer flat-field calibration file (Section 4.3.7). The new flat-field correction algorithm used for PDS Version 1 includes the correction for the ASF in the flat-field calibration step. In prime DI mission calibrations, on the other hand, the flat-field correction ignored the change across the ASF boundary, and the correction for the ASF was implemented using two different absolute calibration curves, one of which was applied outside the ASF and the other inside the ASF.

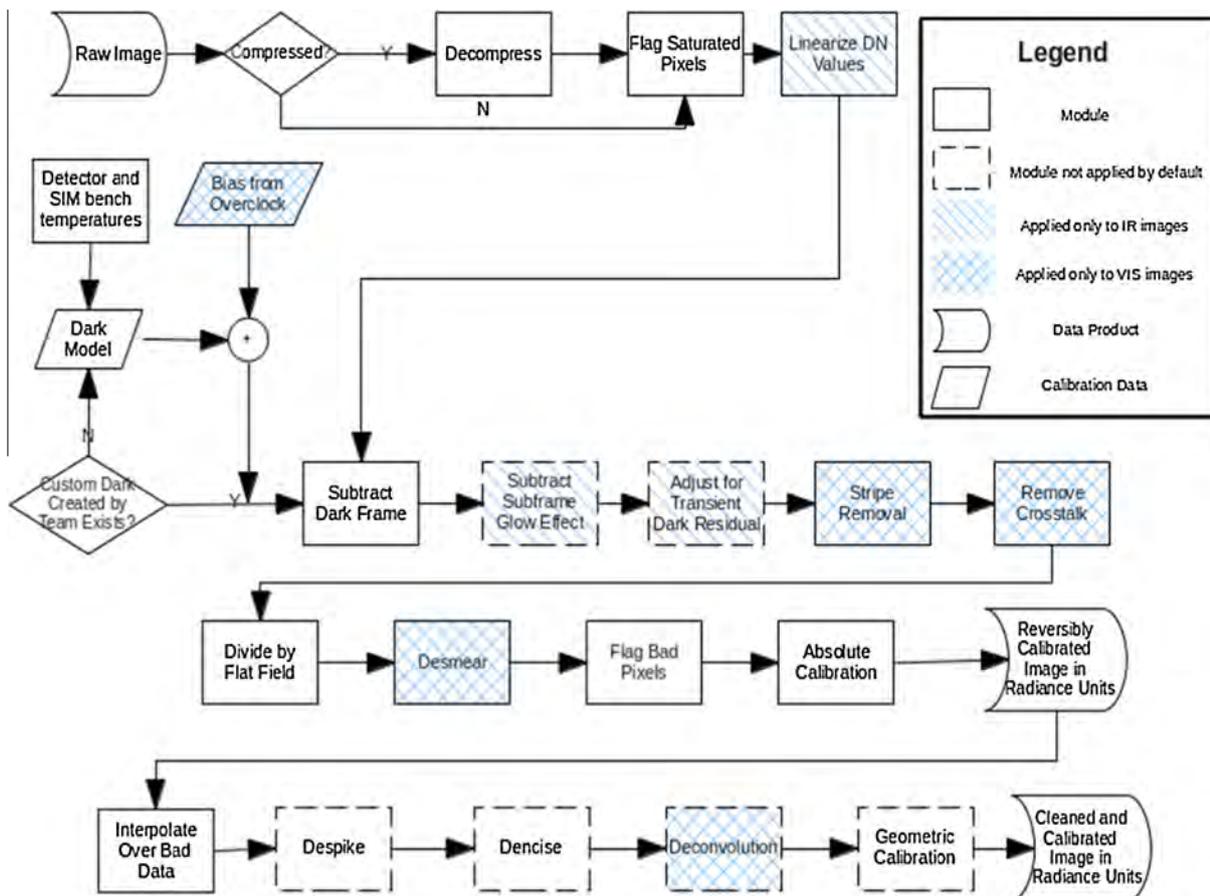


Fig. 61. Logic flow for the EPOXI image processing pipeline.

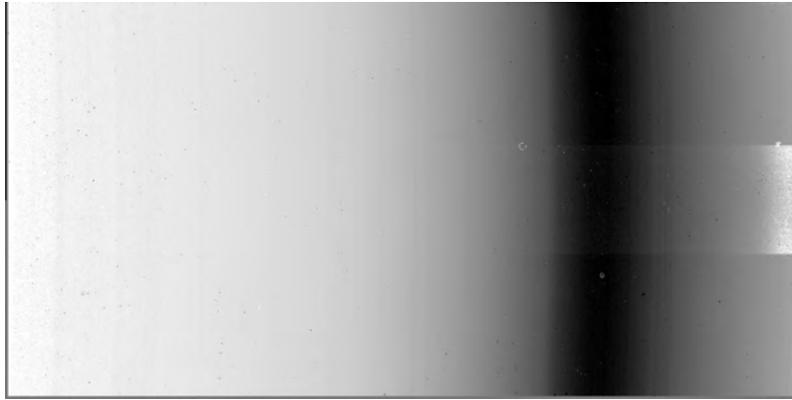


Fig. 62. The IR spectrometer frame from Fig. 1 processed using the calibration pipeline for PDS Version 2. The radiance range is from 0.1 to 1.5 W/m²/μm/sr. A logarithmic contrast stretch has been applied to allow visibility across the entire radiance range. The left side (shorter wavelengths) of the spectrum is dominated by reflected sunlight while the right side (longer wavelengths) is dominated by thermal emission. Interpolation over bad pixels has been performed, but some artifacts remain. The ASF signature is largely, but not entirely, removed.

6.6. IR ALTFP line-dependent integration time

Ongoing analysis of IR timing (Section 4.3.3.2) revealed that read and reset timings differ in the ALTFP mode, and the effective exposure time varies with line number, i.e., along the slit in the spatial direction. This variation has been modeled in the pipeline software, and line-specific integration times are used to convert from DN to DN/s in the ALTFP mode.

6.7. VIS flat-field file re-orientation

Code was modified to change the orientation of pixels in the VIS flat-field calibration files, including POCs and columns and rows just inside OCs, to match that in the flight images. Earlier calibration files interleaved pixels by quadrant and shifted the overclocked pixels to the end of the rows and columns of each quadrant, which is how calibration images were generated from pre-launch ground calibrations. A keyword in the FITS header designates the pixel orientation.

6.8. Header fixes and elimination of I/F files

The pipeline software changes described above generated changes in the calibrated data FITS headers, such as: whether or not a calibration step was applied; calibration files used; constants used; other parameters; and previous observation history. Where appropriate, new calibration steps are included in the SNR extension.

Also, since the VIS reflectance (I/F) calibrated data differ from the radiance data by only a constant factor, the pipeline no longer generates I/F files. It instead generates only the radiance data file and includes the conversion factor from radiance to I/F in the FITS header. For the IR spectrometer, the I/F version was judged not useful due to the presence of thermal emission from the source.

6.9. PDS Version 2 (IR spectrometer data only)

Version 1 of the HRI-IR Hartley 2 data archived with the Planetary Data System (McLaughlin et al., 2011a,b,c) does not include the pixel-by-pixel linearity correction treatment and its propagation through the calibration steps (i.e., bad-pixel map, flat-field file update, revised spectral calibration curve), the mode-dependent master darks, the optimized scaling factor for the master dark, or the revised ASF spectral transmission curve. A refinement in the absolute spectral calibration curve was also implemented. The

sub-frame master darks also now include the effect of glow from saturated border pixels (Section 4.3.3.6). These improvements are all included in Version 2 to be archived with the PDS. A version of the IR spectrometer frame in Fig. 1 processed using the calibration pipeline for PDS Version 2 is shown in Fig. 62.

7. Conclusion

Updates to the instrument calibrations achieved during the Deep Impact primary mission and continued monitoring of instrument performance during the EPOXI extended mission have led to results that will allow better interpretation of the scientific data collected during both of those missions. The instruments' performance has remained remarkably stable over the nearly 7 years of flight. Significant improvements in the understanding and calibration of the IR spectrometer response non-linearity, time-varying background level, flat field, wavelength map, and absolute spectral response have been achieved. Techniques for reducing some semi-coherent horizontal noise stripes in the visible cameras' readouts were developed, and some adjustments have been made to their absolute radiometric conversion constants. The data processing pipeline has been updated to incorporate the improvements in the instrument calibrations.

Acknowledgments

This work was supported by National Aeronautics and Space Administration's Discovery Program through Contract NNM07AA99C to the University of Maryland and by task order NMO711002 to the Jet Propulsion Laboratory, California Institute of Technology. The contributions of O. Groussin and F. Merlin to this project were funded by the Centre National d'Etudes Spatiales (CNES) in France. We appreciate the data processing contributions of David Brown, Peter Kim, Nick Mastrodomos, Nelly Mouawad, and Brian Prager and the Project flight team efforts of Tim Larson, Al Nakata, Steve Wissler, Greg LaBorde, Rich Rieber, Steve Collins, and Amy Walsh in collecting the data. Valuable review comments were provided by Dr. David Humm of the Applied Physics Lab.

References

- A'Hearn, M.F. et al., 2010. Comet 103P/Hartley. *IAU Circ.* 9177, 1.
- A'Hearn, M.F. et al., 2011. EPOXI at Comet Hartley 2. *Science* 332, 1396–1400.
- A'Hearn, M.F. et al., 2012. Corrections and clarifications. *Science* 335, 918.
- Ballard, S. et al., 2010. A search for additional planets in the NASA EPOXI observations of the exoplanet system GJ436. *Astrophys. J.* 716, 1047–1059.

- Barry, R.K. et al., 2010. Development and utilization of a point spread function for the Extrasolar Planet Observation and Characterization/Deep Impact Extended Investigation (EPOXI) mission. *Proc. SPIE* 7731, 77313D-1–77313D-7.
- Belton et al., 2011. Stardust-NExT, Deep Impact, and the accelerating spin of 9P/Tempel 1. *Icarus* 213, 345–368.
- Campins, H., Rieke, G.H., Lebofsky, M.J., 1985. Absolute calibration of photometry at 1 through 5 microns. *Astron. J.* 90, 896–899.
- Christiansen, J.L. et al., 2010. Studying the atmosphere of the exoplanet HAT-P-7b via secondary eclipse measurements with *EPOXI*, *SPITZER*, and *KEPLER*. *Astrophys. J.* 719, 97–104.
- Christiansen, J.L. et al., 2011. System parameters, transit times, and secondary eclipse constraints of the exoplanet systems HAT-P-4, TrES-2, TrES-3, and WASP-3 from the NASA *EPOXI* mission of opportunity. *Astrophys. J.* 719, 97–104.
- Farnham, T.L., Schleicher, D.G., A'Hearn, M.F., 2000. The HB narrowband comet filters: standard stars and calibrations. *Icarus* 147, 180–204.
- Hampton, D.L. et al., 2005. An overview of the instrument suite of the Deep Impact mission. *Space Sci. Rev.* 117, 43–93.
- Janesick, J., Klaasen, K., Elliott, T., 1987. Charge-coupled-device charge-collection efficiency and the photon-transfer technique. *Opt. Eng.* 26 (10), 972–980.
- Klaasen, K.P. et al., 2008. Deep Impact instrument calibration. *Rev. Sci. Instrum.* 79, 091301-1–091301-77.
- Lançon, A., Rocca-Volmerange, B., 1992. A library of near-IR stellar spectra from 1.328 to 2.5 microns. *Astron. Astrophys. Suppl. Ser.* 96 (3), 593–612.
- Livengood, T.A. et al., 2011. Properties of an Earth-like planet orbiting a Sun-like star: Earth observed by the *EPOXI* mission. *Astrobiology* 11 (9), 907–930.
- McLaughlin, S.A., Carcich, B., Sackett, S.E., Klaasen, K.P., 2011a. *EPOXI* 103P/Hartley2 Encounter – HRII Calibrated Spectra V1.0, DIF-C-HRII-3/4-EPOXI-Hartley2-V1.0, NASA Planetary Data System.
- McLaughlin, S.A., Carcich, B., Sackett, S.E., Klaasen, K.P., 2011b. *EPOXI* 103P/Hartley2 Encounter – HRIV Raw Calibrated V1.0, DIF-C-HRIV-3/4-EPOXI-Hartley2-V1.0, NASA Planetary Data System.
- McLaughlin, S.A., Carcich, B., Sackett, S.E., Klaasen, K.P., 2011c. *EPOXI* 103P/Hartley2 Encounter – MRI Calibrated Images V1.0, DIF-C-MRI-3/4-EPOXI-Hartley2-V1.0, NASA Planetary Data System.
- Smith, R.M., Zavodny, M., Rahmer, G., Bonati, M., 2008a. A theory for image persistence in HgCdTe photodiodes. *Proc. SPIE* 7021, 70210J-1–70210J-12.
- Smith, R.M., Zavodny, M., Rahmer, G., Bonati, M., 2008b. Calibration of image persistence in HgCdTe photodiodes. *Proc. SPIE* 7021, 70210K-1–70210K-12.
- Vandenbussche, B. et al., 2002. The ISO-SWS post-helium atlas of near-infrared stellar spectra. *Astron. Astrophys.* 390 (3), 1033–1048.
- Wallace, L., Meyer, M.R., Hinkle, K., Edward, S., 2000. Near-infrared classification spectroscopy: J-band spectra of fundamental MK standards. *Astrophys. J.* 535 (1), 325–337.