

Invited Article: Deep Impact instrument calibration

Kenneth P. Klaasen,^{1,a)} Michael F. A'Hearn,^{2,b)} Michael Baca,^{3,c)} Alan Delamere,^{4,d)} Mark Desnoyer,^{5,e)} Tony Farnham,^{2,f)} Olivier Groussin,^{2,g)} Donald Hampton,^{6,h)} Sergei Ipatov,^{2,i)} Jianyang Li,^{2,j)} Carey Lisse,^{7,k)} Nikolaos Mastrodemos,^{1,l)} Stephanie McLaughlin,^{2,m)} Jessica Sunshine,^{2,n)} Peter Thomas,^{5,o)} and Dennis Wellnitz^{2,p)}

¹*Jet Propulsion Laboratory, California Institute of Technology, Mail Stop 306-392, 4800 Oak Grove Dr., Pasadena, California 91109, USA*

²*Department of Astronomy, University of Maryland, College Park, Maryland 20742-2421, USA*

³*Science Applications International Corporation, 5180 Parkstone Drive, Suite 100, Chantilly, Virginia 20151, USA*

⁴*Delamere Support Services, 525 Mapleton Ave., Boulder, Colorado 80304, USA*

⁵*Cornell University, Space Science Building, Ithaca, New York 14853, USA*

⁶*Ball Aerospace and Technologies, 1600 Commerce St., Boulder, Colorado 80301, USA*

⁷*Johns Hopkins University Applied Physics Laboratory, SD/SRE, MP3/W-155, 7707 Montpelier Road, Laurel, Maryland 20723, USA*

(Received 9 August 2007; accepted 18 July 2008; published online 25 September 2008)

Calibration of NASA's Deep Impact spacecraft instruments allows reliable scientific interpretation of the images and spectra returned from comet Tempel 1. Calibrations of the four onboard remote sensing imaging instruments have been performed in the areas of geometric calibration, spatial resolution, spectral resolution, and radiometric response. Error sources such as noise (random, coherent, encoding, data compression), detector readout artifacts, scattered light, and radiation interactions have been quantified. The point spread functions (PSFs) of the medium resolution instrument and its twin impactor targeting sensor are near the theoretical minimum [~ 1.7 pixels full width at half maximum (FWHM)]. However, the high resolution instrument camera was found to be out of focus with a PSF FWHM of ~ 9 pixels. The charge coupled device (CCD) read noise is ~ 1 DN. Electrical cross-talk between the CCD detector quadrants is correctable to < 2 DN. The IR spectrometer response nonlinearity is correctable to $\sim 1\%$. Spectrometer read noise is ~ 2 DN. The variation in zero-exposure signal level with time and spectrometer temperature is not fully characterized; currently corrections are good to ~ 10 DN at best. Wavelength mapping onto the detector is known within 1 pixel; spectral lines have a FWHM of ~ 2 pixels. About 1% of the IR detector pixels behave badly and remain uncalibrated. The spectrometer exhibits a faint ghost image from reflection off a beamsplitter. Instrument absolute radiometric calibration accuracies were determined generally to $< 10\%$ using star imaging. Flat-field calibration reduces pixel-to-pixel response differences to $\sim 0.5\%$ for the cameras and $< 2\%$ for the spectrometer. A standard calibration image processing pipeline is used to produce archival image files for analysis by researchers. © 2008 American Institute of Physics. [DOI: 10.1063/1.2972112]

TABLE OF CONTENTS

I. OVERVIEW.	2	III. INSTRUMENT DESCRIPTIONS.	4
II. INTRODUCTION TO THE MISSION.	3	A. CCD.	5
		B. IR focal plane.	6

^{a)}Tel.: (818) 354-4207. FAX: (818) 393-3290. Electronic mail: kenneth.p.klaasen@jpl.nasa.gov.

^{b)}Tel.: (301) 405-6076. FAX: (301) 405-3538. Electronic mail: ma@astro.umd.edu.

^{c)}Tel.: (703) 227-8089. FAX: (703) 802-9515. Electronic mail: michael.w.baca@saic.com.

^{d)}Tel.: (303) 447-2780. FAX: (303) 448-0931. Electronic mail: alan@delamere.biz.

^{e)}Tel.: (607) 255-7453. FAX: (607) 255-6918. Electronic mail: md246@cornell.edu.

^{f)}Tel.: (301) 405-3856. FAX: (301) 405-3538. Electronic mail: farnham@astro.umd.edu.

^{g)}Tel.: (301) 405-3354. FAX: (301) 405-3538. Electronic mail: groussin@astro.umd.edu.

^{h)}Tel.: (303) 938-0490. Electronic mail: donhampton@earthlink.net.

ⁱ⁾Tel.: (301) 405-8058. FAX: (301) 405-3538. Electronic mail: ipatov@astro.umd.edu.

^{j)}Tel.: (301) 405-2103. FAX: (301) 405-3538. Electronic mail: jyli@astro.umd.edu.

^{k)}Tel.: (240) 228-0535. FAX: (240) 228-8939. Electronic mail: carey.lisse@jhuapl.edu.

^{l)}Tel.: (818) 393-2477. FAX: (818) 393-6388. Electronic mail: Nikolaos.Mastrodemos@jpl.nasa.gov.

^{m)}Tel.: (301) 405-3605. FAX: (301) 405-3538. Electronic mail: jyli@astro.umd.edu.

ⁿ⁾Tel.: (301) 405-1045. FAX: (301) 405-3538. Electronic mail: stefmcl@astro.umd.edu.

^{o)}Tel.: (607) 255-5908. FAX: (607) 255-9002. Electronic mail: thomas@baritone.astro.cornell.edu.

^{p)}Tel.: (301) 405-1546. FAX: (301) 405-3538. Electronic mail: wellnitz@astro.umd.edu.

C. Inflight stimulators.....	6
D. Quadrant nomenclature.....	6
E. Image compression.....	9
F. General flight performance.....	9
1. HRI focus.....	10
2. IR spectrometer operating temperature....	10
IV. CALIBRATION DATA COLLECTION.....	10
V. VIS CAMERA CALIBRATIONS.....	11
A. Geometric calibrations.....	11
1. Focal length and geometric distortion. . .	11
2. Relative boresight alignments.....	17
B. Spatial resolution.....	18
1. MRI and ITS point spread functions....	18
2. HRI point spread function.....	19
3. HRI deconvolution.....	21
C. Radiometric calibration.....	21
1. Linearity.....	21
2. Gain/full well.....	21
3. Zero-exposure level.....	23
4. Frame transfer smear.....	24
5. Absolute sensitivity.....	25
6. Individual pixel response.....	32
7. Noise.....	34
8. Scattered light.....	37
9. Charge bleeding/residual.....	39
10. Radiation noise.....	40
VI. IR SPECTROMETER CALIBRATIONS.....	43
A. Geometric calibrations.....	43
1. Focal length.....	43
2. Relative boresight alignments.....	44
3. Slit alignment.....	44
B. Spatial resolution.....	45
C. Radiometric calibration.....	46
1. IR linearity.....	46
2. Gain/full well.....	49
3. Zero-exposure level.....	51
4. Wavelength map.....	57
5. Spectral resolution.....	58
6. Absolute spectral sensitivity.....	58
7. Individual pixel response.....	63
8. Noise.....	68
9. Scattered light.....	69
10. Charge bleeding/residual.....	71
11. Radiation noise.....	72
VII. PIPELINE PROCESSING.....	73
A. Standard steps.....	73
B. Calibration quality map.....	74
C. Signal-to-noise ratio map.....	74
D. Spectral registration maps.....	75
E. Optional steps.....	75
APPENDIX A: DATA COMPRESSION LOOKUP TABLES PRIOR TO I-16 DAYS (18 JUNE 2005).....	75
APPENDIX B: ACRONYMS.....	76

I. OVERVIEW

In support of the U.S. vision for space exploration, which is to advance the nation's scientific, security, and economic interests through a robust space exploration program,

NASA's Science Mission Directorate carries out the scientific exploration of the Earth, Moon, Mars, and beyond; charts the best route of discovery; and reaps the benefits of Earth and space exploration for society. NASA conducts scientific exploration that is enabled by access to space, extending humankind's virtual presence throughout the solar system via robotic visitors to other planets and their moons, to asteroids and comets, and to icy bodies in the outer reaches of the solar system. The principal goals are to understand our origins, to learn whether life does or did exist elsewhere in the solar system, and to prepare for human expeditions to the Moon, Mars, and beyond. The particular goal in the area of planetary science is to advance scientific knowledge of the origin and history of the solar system, the potential for life elsewhere, and the hazards and resources present as humans explore space.

As one means of achieving this goal, spacecraft are sent to various solar system bodies from which to make scientific measurements not possible to make from Earth. These measurements are obtained using science instruments designed for that purpose and capable of performing reliably in the harsh space environment. Digitized data from such instruments are radioed back to Earth for analysis. However, in order for the measurements made by these instruments to be scientifically useful, the digitized engineering data they send back must be converted to physical units representing some aspect of the object that was observed. This conversion is enabled by accurately calibrating the output of the instruments in response to observations of other well-known source targets. These calibrations are typically carried out using both ground-based laboratory measurements and in-flight measurements using celestial or onboard targets. Equally important for scientific analyses is knowing the uncertainties in the calibrated measurement quantities and the possible sources of error that could invalidate the measurement interpretations.

One recent NASA mission called Deep Impact (DI) sent a robotic spacecraft to investigate comet Tempel 1. The spacecraft carried four scientific instruments: three digital cameras with identical detectors and electronics but different optical telescopes and one infrared (IR) imaging spectrometer that used as its fore optics the same telescope as the highest resolution camera. These instruments can be used to make measurements of the size, shape, and relative spatial locations of target objects, of the morphology of surface features, of the reflectance properties of objects as a function of wavelength, of the temperatures of target objects, and of any time variability in these properties. Thus, it is essential to calibrate the instruments in the areas of the geometric fidelity of their images, the spatial resolution of the images, their absolute spectral response rates to source radiance at each pixel as a function of wavelength, the precise wavelengths of the light being measured, as well as the accuracies of each of these calibrations and a quantification of the possible sources of error that could degrade the data quality. The techniques for calibrating these instruments are presented here along with the results of the calibrations. This information will be critical for any researcher intending to do scientific research using the DI instrument data. The general calibration tech-

niques and processes should also be of interest to anyone using similar instruments to do scientific research.

We begin with a brief description of the DI mission followed by a high-level description of the various science instruments. We include a summary of the calibration data available for these instruments—both ground-based and in-flight. The results of the calibrations are then presented. Results for the three cameras are discussed together since many of their performance parameters are similar. The calibration of the IR spectrometer is presented next. We discuss geometric calibration, spatial resolution calibration, and radiometric calibration for each instrument. Finally we present the standard data processing steps the DI science team used to create calibrated data sets for science analysis.

II. INTRODUCTION TO THE MISSION

The primary scientific goals of NASA's DI mission¹ are (1) to determine the differences between a comet nucleus' surface, with its ambient outgassing, and its interior, which might contain enhanced volatiles or otherwise different compositions, and (2) to determine the structural properties of the nucleus surface layers, such as whether those layers are strong or strengthless. These goals were addressed by creating a high-energy impact into a comet nucleus and observing from close range the resulting crater and the impact ejecta outflow characteristics. The comet selected for investigation was the short-period comet 9P/Tempel 1, whose nucleus has a mean radius of about 3 km.² The DI spacecraft was launched from Cape Canaveral on 12 January 2005 and arrived at Tempel 1 on 4 July 2005 (GMT) at a range of about 1.5 AU from the Sun and 0.9 AU from the Earth.

The flight system consisted of two fully functional spacecraft. The main spacecraft, which was designed to make a close flyby of the nucleus of Tempel 1, carried an impactor spacecraft, which was intended to impact the nucleus to create a crater opening up the subsurface of the nucleus for investigation. The impactor weighed 364 kg and

carried a charge-coupled device (CCD) camera (described in Sec. III) for both science and autonavigation, an S-band radio link to the flyby spacecraft, a thruster-based attitude control system with both a star tracker and a gyro for attitude measurement, and a complete autonavigation system designed to analyze the images onboard and target the impactor to a site on the nucleus that would be in sunlight and visible from the flyby spacecraft. The impact speed was 10.2 km/s, which resulted in a kinetic energy of impact of about 19 GJ, corresponding to the explosive power of 4.5 tons of TNT.

The flyby spacecraft with the impactor mated to it was initially placed on a trajectory that would result in an impact with the nucleus. The impactor then separated from the flyby spacecraft 24 h before impact, at which point the flyby spacecraft diverted to miss the nucleus by 500 km and slowed down by 100 m/s to provide an 800 s viewing window after impact. The instruments on the flyby spacecraft (described in Sec. III) are body mounted, so the flyby spacecraft had to be rotated as it passed by the nucleus in order to keep the instruments pointed at the comet. At 500 km prior to closest approach, the flyby spacecraft froze in an attitude that kept its dust shields in the direction of its relative velocity vector to protect the spacecraft from damage due to possible particle impacts near closest approach. After passing through the innermost coma, the spacecraft turned and looked back at the comet to take additional data.

Both spacecraft performed their mission as planned. The impactor selected and achieved an impact site that was sunlit and viewable from the flyby spacecraft throughout the available viewing period. The impact event produced a spectacular explosion of material from the nucleus subsurface. The flyby spacecraft kept its instruments pointed at the nucleus prior to and after the impact event. The instrument pointing after impact turned out to be biased somewhat further away from the impact site than planned because the plume of impact ejecta was much brighter, optically thicker, and more persistent than expected. This bright plume caused the autonavigation system to interpret the plume as part of the

TABLE I. Deep Impact instrument suite summary.

	HRI-IR	HRIVIS	MRI	ITS
	Telescope			
Aperture diameter (cm)	30	30	12	12
Focal length (m)	10.5	10.5	2.1	2.1
$f/\#$	12 ^a	35	17.5	17.5
	Imaging			
Largest format	512×256 ^b	1024×1024	1024×1024	1024×1024
Pixel size (μm)	36	21	21	21
I FOV (μrad)	10	2	10	10
FOV (mrad)	2.5	2	10	10
FOV (deg)	0.15	0.118	0.587	0.587
Scale per pixel (m)	7 at 700 km	1.4 at 700 km	7 at 700 km	0.37 at 37 km
Mass ^c (kg)		56.8	19.4	17.9
Imaging power ^d (W)		40	35	35

^aWithin the spectrometer.

^b1024×512 pixels at 18 μm each, a majority of the data taken during the encounter was rebinned 2×2.

^cIncludes optical sensor, cables, and electronic box.

^dDoes not include mechanism power, light blocker or filter wheel.

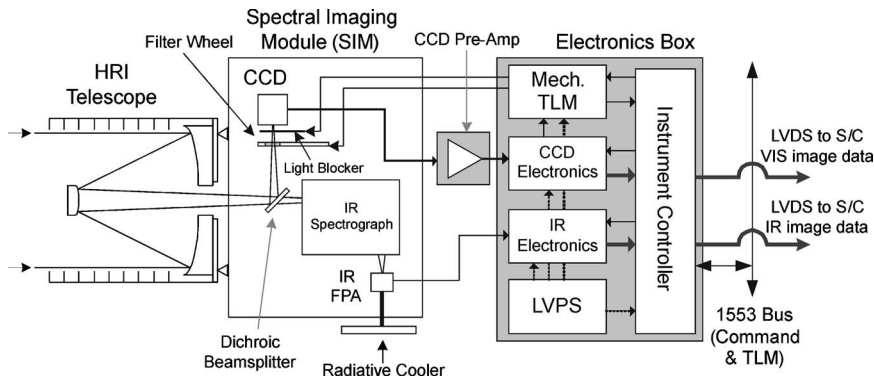


FIG. 1. HRI block diagram (Ref. 3, used with permission); FPA, focal plane array; TLM, telemetry; LVPS, low-voltage power supply; LVDS, low-voltage differential signaling interface.

nucleus resulting in an offset in its calculation of the direction toward the nucleus' center of brightness. However, despite this pointing bias, excellent science data were collected throughout the encounter.

III. INSTRUMENT DESCRIPTIONS

The four DI instruments are described in this section. A more detailed description can be found in Ref. 3. All four instruments are remote sensing imagers; they vary in their spatial resolutions and in their spectral ranges and spectral resolution capabilities. In addition to the instrument descriptions, we describe the standard orientation of the data in the archived image files, the available methods for compressing the data within the instruments, and an overview of how the instruments actually performed in flight.

The DI science instrument suite consisted of three individual instruments—a high resolution instrument (HRI) and a medium resolution instrument (MRI) on the flyby spacecraft, and an impactor targeting sensor (ITS) on the impactor.³ Fundamental parameters of the instruments are shown in Table I. Figures 1 and 2 show block diagrams of the HRI and MRI/ITS configurations, respectively, showing all the major system components.

The HRI consists of a 10.5 m focal length (f) Cassegrain telescope, with a 30 cm aperture, feeding both a filtered CCD camera and a long-slit imaging spectrometer. The two channels are separated by means of a dichroic beamsplitter, reflecting visible (vis) light from 0.34 to 1.05 μm wavelength, and transmitting IR light from 1.05 to greater than 4.8 μm . The vis light passes through one of nine filters and is collected by a 1024 \times 1024 pixel split-frame-transfer CCD. The distance between the closest filter surface and the CCD is 58.5 ± 0.5 mm for both HRI and MRI. The CCD is

divided into four independent 512 \times 512 pixel quadrants, each with a separate readout chain, in which the signal is converted to 14 bit digital data. HRI filter characteristics are shown in Table II.

The IR light is focused onto the slit of a two-prism spectrometer. The wavelength range of the spectrometer is 1.05–4.8 μm , with spectral resolving power of greater than 200 over the entire spectrum, and as high as 700 at 1.05 μm . The spectra are collected on one-half of a 1024 \times 1024 pixel HgCdTe midwave IR (MWIR) detector, for an effective 1024 spectral and 512 spatial pixels. Similar to the CCD, the IR detector is divided into four quadrants, two of which are converted to 14 bit digital data for storage.

To handle the large dynamic range between the warm nucleus and dim gas emission lines, an antisaturation filter was placed over the center third of the slit. The filter was chosen to attenuate the spectrum at wavelengths longer than 2.7 μm , and thus the thermal spectrum of the nucleus. Detector rows 172 through 329 are located behind the antisaturation filter.

The pixel scale for the HRI vis channel is 2 μrad , and for the HRI IR channel is 5 μrad per physical pixel. The slit width of the spectrometer is set to span two physical pixels, and most of the encounter data were taken in binned modes, with an effective pixel scale of 10 μrad .

The MRI and ITS are structurally very similar. Both are based on a 2.1 m fl Cassegrain telescope, with 12 cm aperture. The MRI has a nine-position filter wheel, while the ITS light is not filtered. MRI filter characteristics are shown in Table II. This paper considers radiometric calibration only to the level of average flux within a filter bandpass. Calibration techniques for separating the continuum- and emission-band fluxes in the cometary coma can be found in Ref. 4 in the

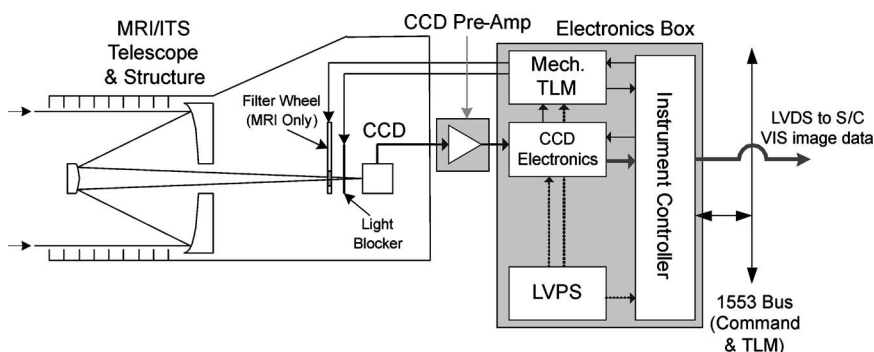


FIG. 2. MRI/ITS block diagram (Ref. 3, used with permission); acronyms are defined in Fig. 1 caption.

TABLE II. Deep Impact filter characteristics (Ref. 3, used with permission).

Filter wheel position	MRI center wavelength (nm)	MRI filter width (nm)	MRI filter target measurement	HRI center wavelength (nm)	HRI filter width (nm)
1	650	>700 ^a	Context	650	>700 ^a
2	514	11.8	C2 in coma	450	100
3	526	5.6	Dust in coma	550	100
4	750	100	Context	350	100 ^b
5	950	100 ^c	Context	950	100 ^c
6	650	>700 ^a	Context	650	>700 ^a
7	387	6.2	CN in coma	750	100
8	345	6.8	Dust in coma	850	100
9	309	6.2	OH in coma	650	100

^aFilters in positions 1 and 6 are uncoated fused silica and not band limited.

^bThe 950 nm filter is longpass.

^cThe coating on the 350 nm filter is shortpass; the substrate limits the short wavelength to ~ 320 nm. The dichroic beamsplitter limits the short cutoff to ~ 300 nm (for HRI only).

context of ground-based narrow-band comet observing filters. Both MRI and ITS have the same four-quadrant frame transfer CCD as the HRI, and are both converted to 14 bit digital data. The pixel scale for both MRI and ITS is $10 \mu\text{rad}$.

The two instruments on the flyby spacecraft (HRI and MRI) are body mounted on the antisunward side of the spacecraft (shaded by the solar array) to enable passive radiative cooling of the HRI IR spectrometer. During the encounter the optical bench of the spectrometer maintained a temperature of ~ 137.5 K, producing a dark signal rate of ~ 750 DN/s (DN=data number, the digitized output of detector counts from the instrument). The IR detector itself was cooled to 84 K with a two-stage radiative cooler. Two-dimensional IR images are built up by slewing the spacecraft at a rate that moves the slit by one slit width during a frame time.

A. CCD

The design of the DI CCD is common to all three vis instruments (HRI vis, MRI, and ITS) and is a split-frame-transfer device with an active imaging area of 1024×1024 pixels. The CCD is read out via four independent amplifiers each fed by a symmetric image area quadrant. The devices are thinned and back-side illuminated to improve overall quantum efficiency. The architecture of the DI CCD is shown in Fig. 3. The timing pattern generator was constructed to produce several readout modes including full 1024×1024 pixel frames and subframe images reduced by factors of 2, down to a minimum of 64×64 pixels, with subframe readout times approximately a factor of 2 shorter for each factor of 2 in frame size dimension. Figures 1 and 2 show that each instrument includes a light blocker in front of its CCD. This blocker is used simply to shield the image area while the image that has been transferred to the storage areas is being read out. This precludes a bright scene from saturating within the image area during the 1.4 s readout time and bleeding charge into the storage area, possibly corrupting the image being clocked out. This light blocker does not control the integration time; that is set by the timing of the split-

frame-transfer architecture. The addition of the light blocker does result in a bit more complexity in the number of image modes that are possible. For smaller subframes, where images are clocked out as fast as possible, the light blocker is not needed. Figure 4 shows typical readouts for modes using and modes not using the light blocker between each exposure. Exposure times are selectable, with a minimum commanded time of 0 ms and a maximum of 1 048 576 ms (1048 s), with a 1 ms resolution. Due to timing pattern generator overhead, a commanded time of 0 ms results in an effective 3.5 ms of exposure time. For modes that do not use the light blocker for each image [see Fig. 4(b)], the minimum exposure time is governed by the readout time of the image size, and the commanded exposure time is in addition to that readout time. Table III lists the available CCD imaging modes. Detailed




FIG. 3. CCD architecture (Ref. 3, used with permission).

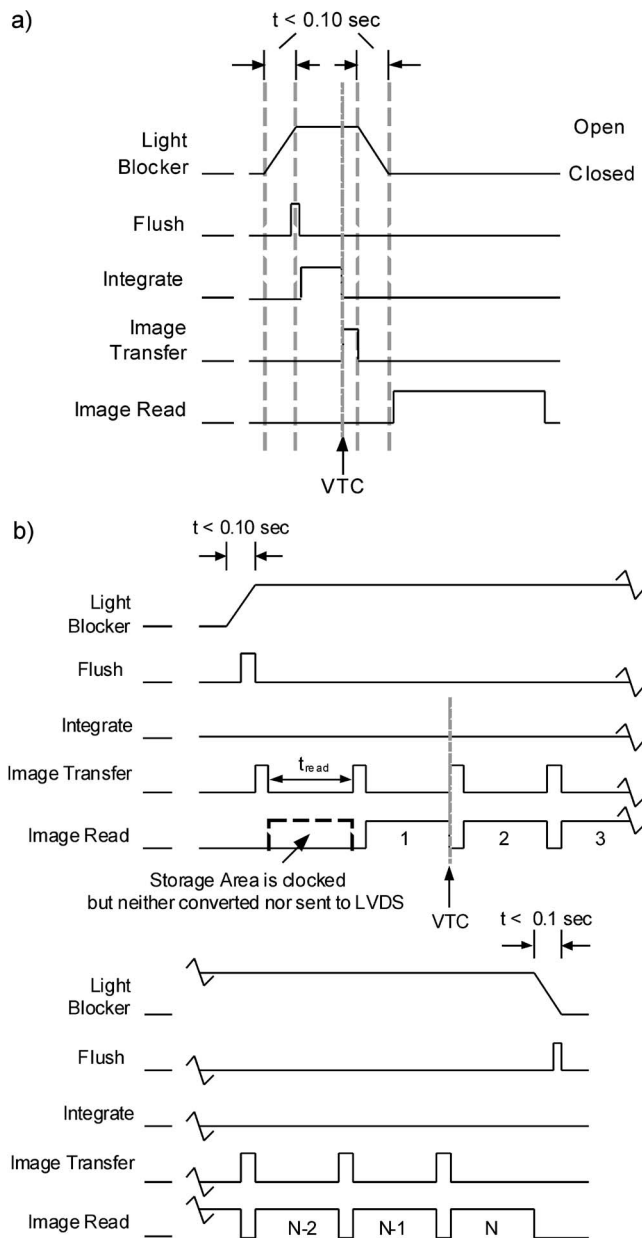


FIG. 4. Generalized timing patterns (Ref. 3, used with permission). (a) Timing pattern for a single exposure when the light blocker is used between each exposure. If N exposures are commanded, then this pattern is repeated N times. (b) Timing pattern when the light blocker is not used between each exposure and there is no additional commanded integration time beyond the minimum. In the second case the number of exposures, N , can be commanded as 1. Both figures show when the time stamp saved in the image header is collected with the arrow and label “VTC.” In both cases the time stamp is collected at the end of the exposure period.

information on the camera modes and timing can be found in Ref. 3.

B. IR focal plane

The IR detector used for the DI IR spectrometer is a hybrid array with a mercury cadmium telluride (HgCdTe) IR sensor indium bump bonded to a Rockwell Science Center HAWAII-IR multiplexer.⁵ The readout circuit was originally developed for the HST (Hubble Space Telescope) Wide-Field Planetary Camera III instrument. For DI, a custom MWIR

HgCdTe substrate with a cutoff wavelength at $4.8 \mu\text{m}$ was developed. The complementary metal oxide semiconductor (CMOS) readout array has 1024×1024 pixels that are $18 \mu\text{m}$ square, and has four electrically independent quadrants with separate readout amplifiers. The optical system was designed such that only two of the quadrants are illuminated. (Thus while each IR image only has two-halves, these halves are still called quadrants). The detector orientation was chosen such that the fast-readout direction is along the spectral direction, while the slow readout is along the slit, as seen in Fig. 5. Therefore the spectrum of each spatial sample has the same time history.

The array is read out in a ripple mode with each pixel addressed directly. Therefore the time history of each pixel is slightly different. Figure 6 shows a diagram of the IR focal plane array (FPA) timing for two readout modes, alternating and interleaved. Most encounter data were taken in interleaved mode. The imaging command accepts an exposure time parameter (termed the integration delay time) that ranges from 0 to 65 535 ms, in 1 ms steps. As seen in Fig. 6, the total integration time is the sum of the integration delay time and the minimum frame-to-frame time for the given readout mode. Table IV lists the available IR spectrometer operating modes.

C. Inflight stimulators

The three CCD channels included an in-flight light emitting diode (LED) stimulator. The LEDs were mounted within the light blocker structure and shine on the back of the light blocker blade. They produce a broad region of illumination on the CCD with a minimum to maximum brightness ratio of greater than 3. The LED was driven directly by spacecraft bus power (tied directly to decontamination heater circuits), and the illumination was not guaranteed to be constant either in a single calibration session (short term) or from session to session (long term). The stimulator images were used to detect any change in high-spatial-frequency performance of the CCD—e.g., dust redistribution, as well as gross changes in quadrant-to-quadrant performance—e.g., gain changes between quadrants.

The instrument design did include an IR internal stimulator designed to shine on the back surface of the dichroic beamsplitter, which would reflect the light into the IR slit. However, the sensitivity of the stimulator to electrostatic discharge events caused it to fail after a single ground test, and it was not repaired prior to any further ground test or flight.

D. Quadrant nomenclature

In order to understand the data from the instruments at the level of calibrations, it is important to understand both the way in which pixels are read out from the detector and the way in which they are stored in the resultant Flexible Image Transport System (FITS)/Planetary Data System (PDS) images.⁶ Throughout this paper we identify the four physical quadrants of the detectors as A–D (or just A and B in the case of the IR detector, which only uses two of the quadrants on the physical detector). The nomenclature in Figs. 7–9 assumes the standard convention for displaying

TABLE III. Visible (CCD) imaging modes (Ref. 3, used with permission).

No.	Mode	Stored size (X and Y)	Serial o'clocked size (X)	Serial o'clocked co-add	Parallel o'clocked size (Y)	Parallel o'clocked co-add	Light blocker used on each image?	Dither?
1	FF	1024	8	0	8	4	Y	N
2	SF1	512	4	0	4	4	Y	N
3	SF2S	256	4	0	4	4	Y	N
4	SF2N	256	4	0	4	4	N	N
5	SF3S	128	2	0	2	4	Y	N
6	SF3N	128	2	0	2	4	N	N
7	SF4O	64	0	0	1	2	N	N
8	SF4NO	64	0	0	0	0	N	N
9	FFD	1024	8	0	8	4	Y	Y




FIG. 5. IR focal plane array architecture (Ref. 3, used with permission). (a) Architecture of the full IR FPA; (b) spectrometer mapping onto the FPA. Only Quads II (A) and IV (B) are used; Quads I and III are not used.

FITS files: the faster-varying index in the data file (samples) is displayed to the right, and the slower varying index (lines) is displayed up (in PDS images the directions are controlled by keywords, which for our images are set to match the standard FITS display). Thus, the first byte of the FITS/PDS file appears in the lower left corner of the window, and the last byte in the FITS/PDS file appears in the upper right. All FITS/PDS archival images are structured to display a true image of the sky, with arbitrary rotation about the center of the image, rather than a mirror image of the sky. The header information in the downlinked data is always written in the first 100 bytes of quadrant A.

Figure 7 shows an in flight, vis image from HRI, in which the directions in the labels are referred to by the order of the bytes (pixels) in the archived data files. The images from the thermal-vacuum (TV) calibrations have the same orientation. For MRI and ITS, the different numbers of reflections in the optical path of the instruments lead to a right-left mirroring between the physical quadrants and the image of the sky and a mirroring between the TV calibrations and the in-flight data. Since the quadrant labeling refers to physical quadrants, the TV calibrations have the same orientation of the quadrants for all three instruments (A in upper left and D in lower right), but they have different orientations for in-flight data, i.e., the in-flight data for MRI and ITS have quadrant A in the upper right and quadrant D in the lower left for normally displayed FITS images. Thus the quadrants for in-flight images from MRI and ITS are shown in Fig. 8.

The readout order of the pixels is different from the order of bytes in the images since each quadrant is read out independently in parallel, and the bytes are then rearranged into an image. The direction of the split-frame rapid transfer is up and down in Figs. 7 and 8, symmetric about the centerline. This affects the smear of bright sources in short exposures (Sec. V C 4). After shifting to the shielded region of the detector, the top and bottom rows are read out first (top and bottom of the relevant subframe when only a subframe is read), and in each of these rows the outermost pixels are read out first. The rows immediately above and below the centerline are read out last, and within these two rows, the pixels immediately adjacent to the centerline are read out last. The header information is overwritten on the first 100 bytes of

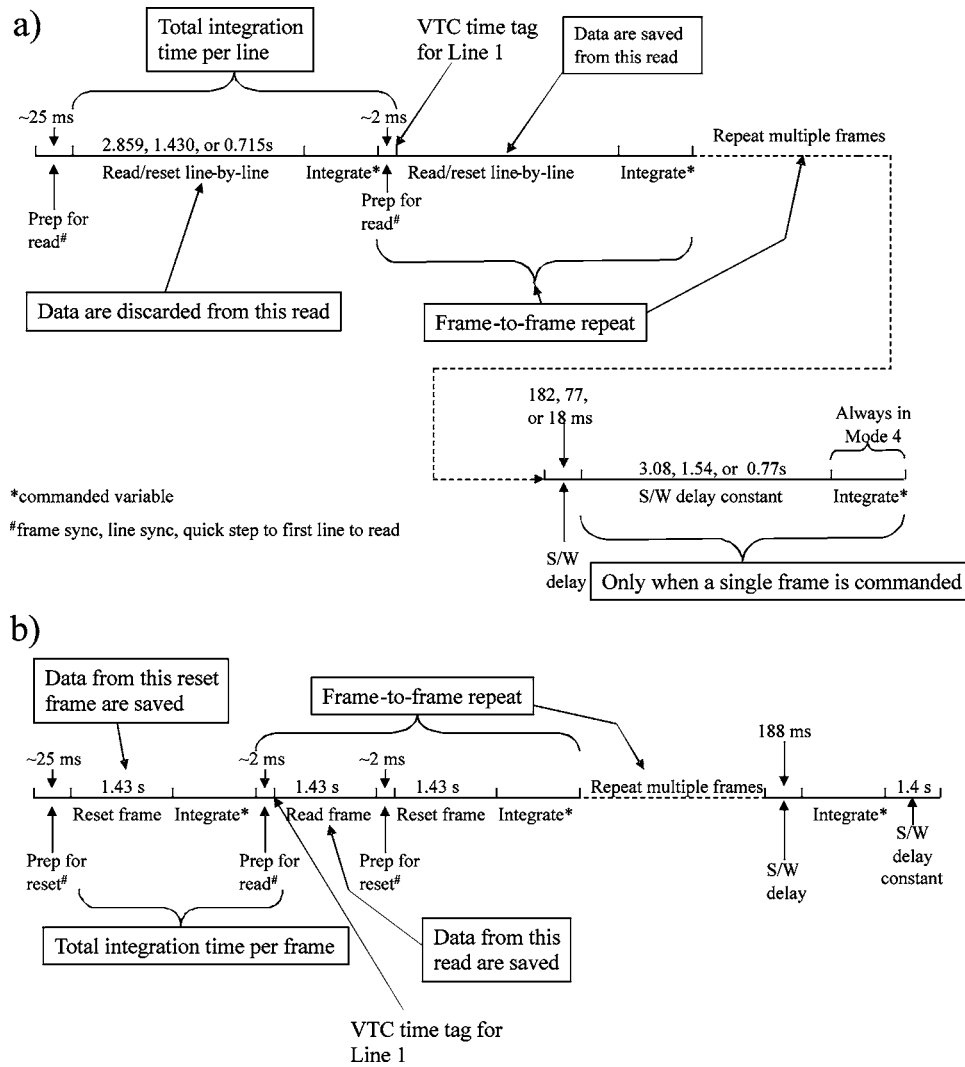


FIG. 6. IR focal plane readout timing. Timing for two different readout modes, (a) interleaved and (b) alternating, are shown.

quadrant A image data (upper left quadrant for HRI in-flight images and upper right quadrant for in-flight images with MRI and ITS) after the image is constructed. Overclocked pixels and rows are read out after the image pixels, but they are moved to the outside of the FITS/PDS image to preserve the contiguity of the image in normal displays.

The situation for the IR spectrometer is shown in Fig. 9. The normally displayed image, whether using the FITS stan-

dard display convention or displaying via the relevant PDS keywords, has wavelength increasing from left to right and the long spatial dimension of the slit oriented vertically. The vertical spatial direction in the spectrometer image is the same as in the HRI vis image, terminator at the top and limb at the bottom for a spectrum at the time of Fig. 7. There are only two quadrants used, although the actual detector has two additional quadrants that are not exposed to light and are

TABLE IV. IR spectrometer operating modes (Ref. 3, used with permission).

No.	Mode mnemonic	Mode	Stored image size	Minimum exposure time (s)	Frame-to-frame time for minimum exposure (s)
1	BINFF	Binned full frame	512 × 256	2.86	2.862
2	BINSF1	Binned sub-frame 1	512 × 128	1.43	1.432
3	BINSF2	Binned sub-frame 2	512 × 64	0.71	0.717
4	UBFF	Unbinned full frame	1024 × 512	2.86	2.862
5	ALTF	Alternating binned full frame	512 × 256	1.43	2.868
6	DIAG	Diagnostic	1024 × 512	1.43	1.432
7	MEMCK	Memory check	1024 × 512	N/A	2.862




FIG. 7. (Color) A full-frame, HRI vis image taken shortly before impact, displayed with the FITS convention. This orientation reproduces a true sky image (ecliptic north is to the right). The first and last bytes are those read from the FITS file and are not connected with the order of detector readout. Quadrants A–D noted throughout this paper are labeled in the image.

not read out. The orientation is the same both for in-flight data and for thermal-vacuum calibrations—A is on the left in a standard FITS/PDS display and B is on the right.

Since the IR detector is reset and read out on a pixel-by-pixel basis, the readout order affects the actual time at which a pixel is exposed, unlike the situation for the vis CCDs. Each pixel has the same exposure duration, but the exposure of the last pair of pixels read out does not start until one integration delay time plus ~ 2 ms before the first pair of pixels is read out (see Fig. 6). As with the lower half of the vis images, the bottom row is read out first, and within that row the outermost (leftmost and rightmost) pixels are read out first. The spectral row at the upper end of the slit in this

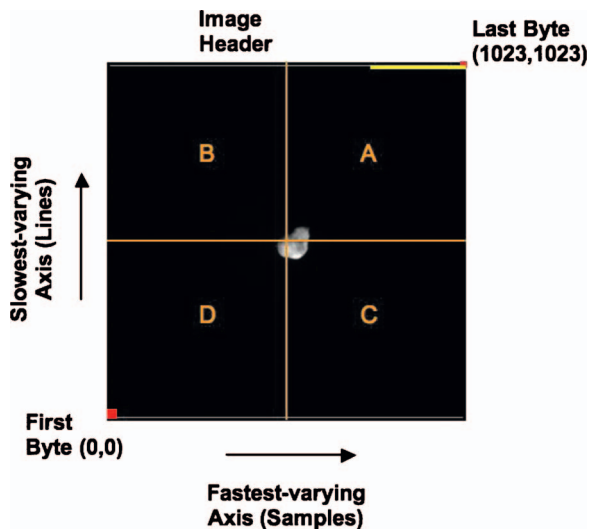


FIG. 8. (Color) A full-frame, MRI image taken at nearly the same time as the HRI vis image in Fig. 7. Displayed with the FITS convention, a true sky image is reproduced (ecliptic north is to the right). The first and last bytes are those read from the FITS file and are not connected with the order of detector readout. Quadrants A–D noted throughout this paper are labeled in the image.




FIG. 9. (Color) A full-frame, HRI-IR image taken shortly before impact, displayed with the FITS convention. For this FITS display, the wavelength increases as the fastest-varying axis increases to the right. The slowest-varying axis is the spatial direction along the slit. The first and last bytes are those read from the FITS file and are not connected with the order of detector readout. IR quadrants A and B noted throughout this paper are labeled in the image.

standard display is read out last, and within that row the two pixels on either side of the centerline are read out last. The header information is again written over the first 100 bytes of quadrant A, now in the lower left of a normal display.

E. Image compression

Image data can be compressed by means of a simple look-up table (LUT), which converts the 14 bit data to 8 bit values. To accommodate the desired volume of images required during the encounter, nearly all of the prime science data were stored as compressed images. Each detector (including the IR FPA) had four dedicated LUTs to choose from, which allowed for some optimization of the compression based on the expected image characteristics. The LUT can be selected to emphasize low-signal or high-signal images. The specific LUT to route the image through was specified in the image command; commands specifying multiple images result in all of those images having the same LUT. Figures 10–13 show the LUTs used during the encounter. Prior to June 18, 2005, the LUTs on the flyby spacecraft were different, and are shown in Appendix A.

F. General flight performance

While, in general, the instruments performed as expected, as described in detail in the rest of this paper, there were two aspects of the HRI performance that did not meet expectations; the focus of the HRI telescope and the operating temperature of the spectrometer at encounter.





FIG. 10. HRI CCD encounter LUTs for converting 14 bit DNs to 8 bit DNs.



FIG. 11. MRI CCD encounter LUTs for converting 14 bit DN₁₄ to 8 bit DN₈s.

1. HRI focus

The HRI focus failed to meet a mission requirement. Calibration images of stars early in the mission showed that the HRI telescope was significantly out of focus. A combined JPL, Ball, and U. of Maryland tiger team was formed to determine the cause of the defocus, as well as any possible methods to recover the focus. The cause was traced to a reference flat mirror used in ground TV tests that distorted more than expected at cryogenic temperatures. Once the focus was determined to be forward of the CCD (secondary despace was away from the primary), a planned bakeout of the graphite epoxy structure took place that reduced the defocus from 1 to 0.6 cm, changing the spot size from effectively 12 pixels to slightly less than 9 pixels. To gain back a significant fraction of the expected resolution, deconvolution algorithms were applied to many HRI vis images, as discussed in Sec. V B 3.

2. IR spectrometer operating temperature

The spectrometer operating temperature did not in and of itself specifically violate a mission requirement, but did increase the system noise level and therefore degrade potential scientific return. The dark signal of the IR spectrometer was dominated not by detector dark current, but by IR emissions from the spectral imaging module (SIM) structure. Several design enhancements were made to the spectrometer design to reduce the SIM temperature and therefore the dark signal from the SIM structure. Two cold shields were added to the

FIG. 12. ITS CCD encounter LUTs for converting 14 bit DN₁₄ to 8 bit DN₈s.FIG. 13. HRI IR FPA encounter LUTs for converting 14 bit DN₁₄ to 8 bit DN₈s.

two-stage detector passive radiator to reduce the out-of-field view from the detector to the structure. The SIM cover was oversized by ~ 10 cm to increase the radiating area of the spectrometer. During cruise and approach the SIM bench ran several degrees warmer than expected due to greater instrument platform heating from direct solar radiation, which could not be well characterized during ground testing. Observation of the elevated temperature resulted in a significant effort to develop a full thermal model of the instrument platform and SIM so that the temperature of the SIM could be predicted correctly at encounter. Several encounter scenarios were modeled, and based on these results, the science team and operations teams developed a set of activities that maintained adequate SIM temperatures at encounter. The desired SIM temperature was < 130 K (as reported by the instrument temperature sensors—the absolute accuracy is ± 2 K). The SIM temperature at the time of encounter was just 137.5 K. Thus the dark signal was approximately 10 000 electrons/s over that at the desired SIM bench temperature.

The elevated background had two impacts on science return. First, and most obviously, the noise associated with the increased dark signal reduced the effective signal-to-noise ratio (SNR) of the spectrometer. Second, as described in Appendix A, the increased dark signal required an in-flight update to the IR LUTs so that the dark signal collected in the minimum IR detector integration times would be compressed to a value close to zero. This LUT change helped to maintain the effective dynamic range of the compressed IR data.

IV. CALIBRATION DATA COLLECTION

We summarize in this section the calibration data that were collected for the DI instruments both on the ground before launch and in flight. Researchers who might wish to perform their own calibrations or to verify the results presented in this paper can use the summaries provided here to help find the appropriate data for their purpose.

Calibrations and other engineering characterization of the DI instruments took place not only prior to launch in both ambient and in TV conditions but also throughout the 6 month cruise to Tempel 1. The majority of the prelaunch data used in calibrating the instruments was taken during four separate TV test programs at Ball Aerospace. These

tests and the data collected are summarized in Table V. The four tests are designated TV1–4. These data were analyzed prior to launch to determine instrument performance characteristics. The results are reported here only to the extent that they have not been superseded by in-flight calibrations.

The performance of the DI instruments was characterized at many times and in various ways throughout the post-launch period. The instruments were first turned on and checked out 2 days after launch ($L+2$ days). Calibrations using the Earth and the Moon were carried out over the next 3 weeks until the distance to these targets became too great and their direction became too close to that of the Sun. After that calibrations were carried out using other celestial targets such as stars, clusters, and planets. Radiometric standard stars and astrometric star clusters were selected for observation by the vis cameras and the IR spectrometer. Planetary nebulae with known spectral lines were also observed by the IR spectrometer. The full set of engineering and science data that could prove useful for calibrating instrument performance is listed in Table VI. Full science calibration sequences were carried out in April, May, June, and July (post-encounter). Abbreviated calibrations of selected parameters (HRI vis PSF, IR spectrometer response, and dark frames for all instruments) were obtained closer to the impact event on July 2, 3, and 9 (the last excluding the previously destroyed ITS). With only a few exceptions, the calibration data were acquired uncompressed. The extensive calibration and test program not only comprehensively characterized instrument performance but also helped the flight team learn in detail how to operate the instruments and spacecraft so as to make successful observations.

V. VIS CAMERA CALIBRATIONS

The calibration results for the three vis cameras are presented here organized by calibration type. Since the cameras all use the same detector and electronics, many of their performance characteristics are similar, and discussing each area of performance for all cameras together rather than discussing each camera separately avoids having to duplicate much of the text. The geometric calibrations of each camera in terms of its fl, which gives the pixel scale projected onto a target, the level of geometric distortion present in the images, and the relative alignments of the boresights of the various cameras are discussed. The camera's spatial resolution capabilities are presented in terms of their PSFs. In particular, the poor focus of the HRI telescope is quantified, and our success at restoring the inherent resolution of the HRI images via PSF deconvolution image processing techniques is described. The radiometric response of each camera is quantified including their response linearity, zero-exposure background levels, relationship between output DN and signal electrons, detector full-well capacities, flat-field response variations, and absolute sensitivity factors for converting DN to scene radiance or surface reflectance. In addition, we discuss quantitatively the various data degraders such as random and coherent noise in the systems, accuracy of analog-to-digital conversion, data compression effects, detector readout artifacts due to low-level cross-talk and frame trans-

fer image smear, scattered light response, bright image blooming and residual signal, and radiation noise effects.

A. Geometric calibrations

1. Focal length and geometric distortion

Images of star clusters were used to obtain fls and distortion factors for the cameras. Primary calibration data are from NGC3114, with some additional data from M11, although the latter did not provide the good spread of many detectable stars across the field of view (FOV) as did the former cluster.

The usual technique is to relate the observed coordinate of each star image on the detector (x_p, y_p in millimeters from the optical center derived from sample/line coordinates at 0.021 mm/pixel) to an “undistorted” position as

$$x_p = (\text{sample-oc}_x)/\text{pixel},$$

$$y_p = (\text{line-oc}_y)/\text{pixel},$$

$$R = \text{SQRT}(x_p^2 + y_p^2) \quad (\text{radius from optical center}),$$

$$x_c = x_p(1 + kR^2),$$

$$y_c = y_p(1 + kR^2),$$

where oc_x and oc_y are the optical center coordinates in pixels, pixel is equal to 47.619 pixel/mm, x_c, y_c are the undistorted positions (mm), and k is the “distortion coefficient.”

The differences between observed and predicted positions on the detector are the residuals, and the solution is taken where the sum of squares of the residuals is minimized. A few clearly anomalous points were eliminated; typically 4–6 out of ~ 300 star coordinates in an image. The optical center was only weakly defined for all imagers, so it was assumed to be at (512, 512).

Stellar image positions are located by a centroiding algorithm, and their RA, Dec locations⁷ and image locations in pixels are recorded. The solution of predicted versus actual positions on the detector includes camera orientation, fl, and distortion (k). Goodness of solution, and fl, depended upon inclusion of a 1/3 pixel offset between top and bottom quadrant pairs designed into the detector (i.e., in the changing line direction).

The MRI and ITS nominally have the same optics and detector geometry. For the MRI, six clear filter and seven red filter images were found suitable for analysis. 265–331 stars were measured in each image, and root-mean-square residuals were between 0.107 and 0.120 pixels.

The MRI clear filter gave $\text{fl}=2101.44 \pm 0.01$ mm and $k=-7.7 \times 10^{-7}$. The red filter gives $\text{fl}=2101.62 \pm 0.05$ mm and $k=-6.8 \times 10^{-7}$. The uncertainties are standard deviations of solutions found for individual images. These are small distortions. If no distortion is assumed ($k=0$), then fl for

TABLE V. Prelaunch calibrations of the DI instruments.

Test designation	Dates	Calibrations conducted			
		HRI vis	MRI vis	IR	ITS
TV1	6/27/02–7/2/02			Alignment, focus, emission and absorption lines, radiometry with blackbody sources and tungsten lamp, flat field, linearity, gain, read noise, full well, dark level, scattered light, latent images, bad pixels	
TV2	8/15/02–9/3/02	Dark current versus detector temperature, response intensity, gain, read noise, full well, test target imaging, alignment, focus		Alignment, focus, emission and absorption lines, spatial resolution, geometry, radiometry with blackbody sources and tungsten lamp through various windows and filters, flat field, linearity, gain, read noise, full well, dark level, scattered light, latent images, bad pixels, meteorite sample reflectance spectrum, mechanical interactions	
TV3	12/24/02–1/30/03				Focus, plate scale, geometric distortion, test target imaging, dark current versus detector temperature, response linearity, gain, read noise, full well, flat field, absolute radiometric response, shutter light leak
TV4	2/23/03–3/11/03	Alignment, focus, plate scale, geometric distortion, test target imaging, dark current versus detector temperature, response linearity, gain, read noise, full well, flat field, absolute radiometric response in all filters, shutter light leak, ADC uniformity, STIM, latent images, frame-transfer smear, electrical crosstalk, in-field scattered light, inter-instrument interactions	Alignment, focus, plate scale, geometric distortion, test target imaging, dark current versus detector temperature, response linearity, gain, read noise, full well, flat field, absolute radiometric response in all filters, shutter light leak, ADC uniformity, STIM, latent image, frame-transfer smear, electrical crosstalk, in-field scattered light, inter-instrument interactions	Alignment, focus, emission and absorption lines, spatial resolution, geometry, radiometry with blackbody sources and tungsten lamp through various windows and filters, flat field, linearity, gain read noise, full well, dark level versus mode switches and variable delays, scattered light (including glint off edge of antisaturation filter), latent images, bad pixels, ADC uniformity, inter-instrument interactions	

TABLE VI. Summary of inflight tests and calibrations using the DI instruments.

Date of Acquisition	DOY(s) of downlink	Activity	Data types				Target(s)	Purpose	Comments
			HR vis	MR vis	IR	ITS			
1/12		Launch							
1/14	15, 16	Quick instrument alignment and engineering checkout	Stars, darks, STIMs	Stars, darks, STIMs	Darks	Darks, STIMs	NGC3532	Engineering validation of operability, alignment, psf	All instruments operational with good relative alignment. Noted HRIVIS to be out of focus.
1/16	17, 18	Earth/moon calibration	Darks, moon, Earth	Darks, moon Earth	Darks, moon, Earth		Earth, moon	Radiometry, focus	Incorrect lunar exposures caused some saturation, most IR data saturated due to insufficient cooling of instrument bench, larger than expected pointing errors due to insufficient settling time allocation with the slow ADCS convergence algorithm employed for launch.
1/25	26, 27	MRI stray light and moon radiometry calibration		Darks, moon			Moon	Radiometry, scattered light	MRI turned off by onboard fault protection shortly after sequence initiation due to requested exposure time in excess of 5 min
1/26	27, 28	Flyby autonav moon test	Moon	Moon			Moon	Autonav test	Clear filter only
1/28	29	MRI stray light and moon radiometry calibration (rerun)	Darks, moon, Earth	Darks, moon, Earth	Darks, moon, Earth		Earth, moon	Radiometry, focus	Good data
1/29	30	IR radiometry and scattered light calibration and HRI focus test	Moon, star	Moon, star	Moon		Moon, iCar	Radiometry, scattered light, focus	IR scattered light test revealed error in data compression LUT (data degraded), IR bench still warmer than desired resulting in elevated background levels and noise, IR flat field only partially successful due to timing mismatch in the sequence, star imagery successful.
2/4	36	IR radiometry and scattered light calibration and HRI focus test (rerun)	Moon, star	Moon, star	Moon		Moon, iCar	Radiometry, scattered light, focus	Obtained good IR flats, IR compression LUTs still incorrect, moon only 58° from the Sun resulted in heating of HRI telescope baffle and excess 5 μ m signal
2/5	37	Flyby autonav Jupiter test	Jupiter	Jupiter			Jupiter	Autonav encounter test	Clear filter only
2/15	46	HRIVIS focus diagnostics (aka a “ghost test”)	Star	Star			Canopus	Determine focal plane position	Long exposures through all filters to bring up ghost image from reflection off filter surface, also “knife-edge” images with filter aperture partially obscuring image beam, HRI bakeout heaters on.
2/18	50	Impactor checkout				Star	Background stars at cruise attitude	Test ITS image quality	Good performance verified
2/21	52, 53	HRIVIS ghost test 2	Star	Star			Canopus	Observe focus as a function of bakeout	
2/22	55	ITS quick alignment check				Stars	NGC3532	Engineering validation of alignment and psf	Successful

TABLE VI. (Continued.)

Date of Acquisition	DOY(s) of downlink	Activity	Data types				Target(s)	Purpose	Comments
			HR vis	MR vis	IR	ITS			
3/2	62	MRI stray light and geometric calibration and precise alignment determination		Stars			NGC3114, various rich fields, encounter attitude star background	Engineering validation of focal length, geometric distortion, psf, alignment	Successful, clear filter only
3/8	67	HRIVIS ghost test 3	Star	Star			Canopus	Observe focus as a function of bake-out	
3/11	70	HRIVIS ghost test 4	Star	Star			Canopus	Observe focus as a function of bake-out	
3/13	72	HRIVIS ghost test 5	Star	Star			Canopus	Observe focus as a function of bake-out	
3/15	74	HRIVIS ghost test 6	Star	Star			Canopus	Observe focus as a function of bake-out	
3/17	76	HRIVIS ghost test 7	Star	Star			Canopus	Observe focus as a function of bake-out	
3/17	76	HRIVIS ghost test 8	Star	Star			Canopus	Observe focus as a function of bake-out	
3/18	78	HRIVIS ghost test 9	Star	Star			Canopus	Observe focus as a function of bake-out	
3/21	81	HGA "jitter" test	Star				?	See effects of the HGA tracking the Earth during encounter attitude maneuver on image smear	Smear observed was minimal
3/24	83	Encounter validation	Darks	Darks	Darks		None	Validate downlink to be used on encounter day	
3/31	91	MRI opnav photocal		Darks, stars			21 standard stars	Determine MRI response to stars of various types	Clear filter only
4/5-4/8	95-99, 119,123	April science calibration	Stars, clusters, darks, STIMs	Stars, clusters, darks, STIMs	Stars, clusters, darks	Stars, clusters, darks, STIMs	Achernar, iCar, 16CygA, Vega, Canopus, bet Hyi, HD60753, Sirius, NGC8543, M11, NGC3114, IC2391	Radiometry, focal length, geometric distortion, psf	The clusters and planetary nebula were not detectable with the IR, the ITS images were mostly mispointed due to a sequencing error (got a few of M11 and some over exposed Vega images useful for crosstalk characterization), good stellar radiometry data in all filters (except for ITS), good IR and VIS linearity data, good geometric data for HRIVIS and MRIVIS
4/18	109	HRIVIS stray light, HRI/MRI alignment	Darks, stars	Stars			NGC3114 and encounter attitude star background	Engineering validation; ADCS attitude knowledge characterization	Clear filter only; ADCS test failed due to setup errors

TABLE VI. (Continued.)

Date of Acquisition	DOY(s) of downlink	Activity	Data types				Target(s)	Purpose	Comments
			HR vis	MR vis	IR	ITS			
4/23	113, 114	HRIVIS stray light, HRI/MRI alignment (rerun)	Darks, stars	Stars			NGC3114 and encounter attitude star background	Engineering validation; ADCS attitude knowledge characterization	Clear filter only
4/27–4/26	118	HRIVIS/IR stray light	Star, darks, STIMs	Star, STIMs	Star, darks		Canopus	Stray light calibration and ADC performance	Successful
4/30	120, 121	ITS autonav Saturn test				Saturn	Saturn	Autonav ITS encounter test	Successful
5/8	126, 127	ITS stray light and alignment				Stars	Encounter attitude star backgrounds	Navigation test	Successful: no apparent stray light in ITS at encounter attitude
5/10–5/13	130–136	May science calibration	Stars, clusters, darks, STIMs	Stars, clusters, darks, STIMs	Stars, clusters, darks	Stars, clusters, darks, STIMs	Achernar, iCar, 16CygA, Vega, Canopus, bet Hyi, HD60753, Sirius, 47Tuc, NGC3114, NGC7027	Radiometry, focal length, geometric distortion, psf, VIS crosstalk, IR alignment	Obtained first good ITS calibration, good crosstalk data, measured IR alignment, obtained IR geometric distortion data over half the slit with 47Tuc, missed IR coverage of NGC7027 and 16CygA due to pointing errors (did cover 16CygB), Incomplete IR coverage of Sirius due to use of wrong scan rate, good radiometry and linearity data, good VIS geometric data
5/21–5/24	145	ITS/ADCS attitude estimation				Stars	NGC3114	Navigation test	Successful
5/28	148	LUT check outs	Darks, STIMs	Darks, STIMs	Darks	Darks, STIMs	None	Verify updated LUTs	New IR LUTs were still incorrect: VIS LUTs good
6/3	153	ITS/ADCS attitude estimation				Stars	NGC3114	Rerun of DOY 145 test to verify repeatability	Successful
6/3	154	MRI/HRI engineering images	Stars	Stars			Cruise attitude star background	Autonav test to characterize cosmic rays and hot pixels	No hot pixels found
6/7–6/9	158–161, 163, 169	June science calibration	Stars, clusters, darks, STIMs	Stars, clusters, darks, STIMs	Stars, clusters, darks	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, crosstalk data, covered other half of IR slit with 47Tuc, lost MRI data of 47Tuc, Canopus, and Achernar plus most of betHyi and 2 dark frames when camera fault protection turned it off. missed IR unbinned scans of NGC7027 due to erroneous zero scan rate in sequence (got binned-mode scan).
6/10	164	Autonav ITS engineering test				Stars	Cruise attitude star background	Autonav test to characterize image quality, cosmic rays and hot pixels	No hot pixels found, representative sample of cosmic rays
6/14	165	Minical	Star, darks	Stars	Star, darks		16CygA, NGC7027	HRI psf, IR response at encounter temperature	Good dithered HRI psf images, IR unbinned scan of NGC7027
6/16	167	IR LUT update checkout			Darks		None	Verify corrected IR compression LUTs	Verified to be correct.

TABLE VI. (Continued.)

Date of Acquisition	DOY(s) of downlink	Activity	Data types				Target(s)	Purpose	Comments
			HR vis	MR vis	IR	ITS			
6/19	170	Autonav HRI/MRI engineering test	Stars	Stars			Cruise attitude star background	Autonav test of image quality and hot pixels	No hot pixels found
6/25	178	Impactor checkout				Stars	Cruise attitude star background	Engineering validation	
6/27–7/3	179–183	ITS engineering checkout				Stars	Cruise attitude star background	Navigation test of image quality and hot pixels	No hot pixels found
7/2	183	Minical	Star, darks	Stars	Star, darks		16CygA, NGC7027	HRI psf, IR response at encounter temperature	Good dithered HRI psf images, IR unbinned scan of NGC7027.
7/3	184	Encounter dark frames	Darks, STIMs	Darks, STIMs	Darks	Darks	None	Final dark/STIM data before encounter	Data acquired successfully
7/6–7/8	188–190	Post-encounter calibration	Stars, clusters, darks, STIMs	Stars, clusters, darks, STIMs	Stars, clusters, darks		Achernar, iCar, 16CygA, Vega, Canopus, bet Hyi, 47Tuc, NGC3114, NGC7027	Radiometry, focal length, geometric distortion, psf, VIS crosstalk, IR alignment	Good radiometry, geometry, linearity data, covered both halves of IR slit with 47Tuc.
7/9	190	Minical	Star, darks	Stars	Star, darks		16CygA, NGC7027	HRI psf, IR response at encounter temperature	Good dithered HRI psf images, IR scans
7/13–7/14	195	Additional HRIVIS STIMs	STIMs, darks					Improve HRIVIS flat field	Successful

TABLE VII. Focal-length solutions for each HRI filter position.

HRI filter	fl (mm)	Images
Clear 1	10 499.94 ± 0.41	3
Blue 2	10 500.28 ± 0.56	7
Green 3	10 499.75 ± 0.28	6
Violet 4	10 502.01 ± 0.63	4
IR 5	10 500.52 ± 1.80	3
Clear 6	10 500.05 ± 0.15	4
Red 7	10 500.91 ± 0.17	2
NIR 8	10 500.93 ± 1.40	3
Orange 9	10 500.69 ± 0.02	3

clear=2101.55 ± 0.01 mm, and fl for red=2101.75 ± 0.02 mm. The pixel-scale difference between the two MRI filters tested is 9.5×10^{-5} , or ~ 0.1 pixel across the detector. The use of an average fl and distortion for all MRI images is fully warranted. The best value would be fl=2101.53 mm; $k = -7 \times 10^{-7}$.

The ITS images give fl=2103.82 ± 0.05 mm and $k = -7.0 \times 10^{-7}$. For $k=0$, the ITS fl=2103.93 ± 0.06 mm.

Data from the May NGC3114 images allowed fl for the HRI in all nine filters to be measured. The stellar images are doughnut shaped (see Fig. 19), and a specialized centroiding routine was implemented to obtain centers. The distortion factor is taken to be 0. Only 9–18 stars per image were reliably centroided, because the defocused PSF leaves low data number values. Table VII gives the results.

Image quality did not allow equal data for all filters, but the results are consistent for uniform parameter use, especially given the slightly out-of-focus situation. Residuals were between 0.1 and 0.2 pixels.

2. Relative boresight alignments

Knowledge of the relative misalignment between the HRI and MRI is important for the scientific interpretation of the encounter observations. One can break down the error budget of the relative misalignment into two parts: (1) the instantaneous error that is due to the measurements and data fit and (2) a buildup of misalignment with time due to thermal drifts between the two cameras. The latter was expected to contribute the bulk of the misalignment error.

The basic plan for the coalignment measurements was to take simultaneous MRI and HRI observations of a star field at a Sun-boresight orientation close to that used at encounter, with a sufficient number of stars in the FOV of each camera to allow an accurate determination of their instantaneous inertial orientations. From these measurements the rotation between the HRI and the MRI coordinate frames can be obtained, averaged over all measurements. The images were commanded to be simultaneous at the midpoint of their exposure with different exposures between HRI and MRI.

The instantaneous error budget can be further broken down into two parts: (a) the absolute determination of each camera's inertial orientation per picture (spatial error) and (b) timing errors between cameras that were commanded to take simultaneous pictures (temporal error). These timing errors

TABLE VIII. Image frame misalignments between the MRI and HRI.

	ψ (mrad) (1- σ)	χ (mrad) (1- σ)	Ω (mrad) (1- σ)
DOY 108	-0.119 660 (±0.000 017)	0.060 092 (±0.000 070)	2.155 0 (±0.031 6)
DOY 113	-0.120 672 (±0.000 017)	0.060 807 (±0.000 035)	2.248 2 (±0.030 9)
Mean of both days	-0.120 166	0.060 449	2.200 4

could cause an angular displacement of one camera relative to the other that could be interpreted as additional rotation of the two boresights. Initial observations at $L+2$ days showed that the relative timing error between MRI and HRI was in the 0–10 ms range (largest timing error observed), time scales over which the angular spacecraft motion is not detectable. A small number of frames was therefore sufficient to average down spatial errors from image to image.

On DOYs 108 and 113 (DOY=day of year 2005), nine simultaneous full frames were acquired per instrument of the center of the open galactic cluster NGC3114. The average number of cataloged stars per frame was 474 and 30 for MRI and HRI, respectively. All MRI images were taken with 4 s exposures, and all HRI images with 5 s exposures. The data reduction process was the same as that described above for the geometric calibration of the cameras. For each image we estimated three rotation angles, α (right ascension), β (declination), and φ (twist).

The typical 1- σ uncertainty in the determination of the camera's pointing per frame was 0.01 μ rad across boresight and 2.6 μ rad about boresight (twist) for the MRI and 0.05 μ rad across and 50 μ rad about boresight for the HRI. The misalignment between HRI and MRI is described by a matrix that rotates the coordinate system of the HRI to that of the MRI, where the right-handed coordinate system of each camera is such that X is along the sample direction (positive toward increasing sample numbers in the FITS images), Y is along the line direction (positive toward decreasing line numbers in the FITS images), and Z is along the boresight (positive in the viewing direction). This rotation matrix is decomposed into three rotations about the camera's coordinate system as

$$M_{\text{HRI-MRI}} = R_3(\Omega)R_1(-\chi)R_2(\psi),$$

where the offset angles ψ , χ , and Ω are rotations about the line, sample, and boresight axes, respectively. This description allows separation of the misalignment into across- and about-boresight components. The estimated offset angles and their 1- σ uncertainties are given in Table VIII for each day.

In matrix form, the mean rotation from HRI to MRI is

$$M = \begin{pmatrix} 0.999\ 997\ 571\ 9 & 0.002\ 200\ 422\ 1 & 0.000\ 120\ 032\ 6 \\ -0.002\ 200\ 414\ 8 & 0.999\ 997\ 577\ 2 & -0.000\ 060\ 713\ 7 \\ -0.000\ 120\ 165\ 9 & 0.000\ 060\ 449\ 5 & 0.999\ 999\ 991\ 0 \end{pmatrix}$$

This rotation places the MRI boresight at line 542.6 and sample 452.4 in a HRI FITS image. The misalignment across boresight, the root sum square (RSS) of ψ and χ , is $133.90\ \mu\text{rad}$ on DOY 108 and $135.13\ \mu\text{rad}$ on DOY 113 for a mean value of $134.51\ \mu\text{rad}$. Although the formal uncertainty of the instantaneous measurement of the misalignment across boresight on either of the 2 days is $<0.07\ \mu\text{rad}$, the difference between the 2 days is much larger than their respective uncertainties and equal to $1.23\ \mu\text{rad}$. We attribute this to thermal drifts in the course of 5 days, and we believe that this is a more representative figure for the error of the boresight alignment knowledge across the field of view.

B. Spatial resolution

1. MRI and ITS point spread functions

In flight, both the MRI and ITS were found to be very close to best possible focus even shortly after launch, so the PSFs can be well characterized using two-dimensional circular Gaussian fits. On January 14, only 2 days after launch, a quick analysis of one of the alignment images taken by the



FIG. 14. Central strip of the MRI alignment image of star cluster NGC 3532, taken 2 days after launch, heavily stretched to show dimmest star images, noise, and bias differences in the four quadrants. Note the triangular shape of the heavily overexposed images due to stressing of the optics by the mounts and six offset diffraction spikes due to the nonradial secondary mount spider vanes.

MRI (Fig. 14) was performed. This 6 s exposure has about 1400 star images, about 650 of which were identifiable from catalogs going down to 14th magnitude. After discriminating against cosmic rays, about 1296 star images had a peak SNR greater than or equal to 100, and circular Gaussians were fit to these images.

The FWHM is plotted as a function of distance in pixels from the center of the detector in Fig. 15. There is a clear trend of increasing FWHM with increasing distance from the center of the detector. Near the center the mean FWHM = 1.65 pixels; at greater than 400 pixels from the center, the mean FWHM increases to 1.9 pixels; and at greater than 500 pixels from the center (the corners of the detector), FWHM increases to more than 2 pixels.

The primary contributors to the widening of the PSF independent of radius are

- (1) smear, which even at the subpixel level can broaden the wings of the PSF, and
- (2) random registration of the stars on the pixel-line grid, which can broaden the PSF substantially (the difference in FWHM between a star at the center of a pixel and the same star at the corner of 4 pixels can be 0.2–0.3 pixels or more).

For purposes of understanding the focus variation across the field of view, we may disregard the FWHMs that do not follow the general trend: those that are low are likely unexcluded cosmic rays or stars that have inadequate Gaussian wings due to falling close to the center of a pixel, whereas those that are high are likely blended images or images contaminated with unexcluded cosmic rays.

The surfaces of best focus of the optics of both the MRI




FIG. 15. PSF FWHM in MRI quick alignment images for SNR larger than 100 (1296 stars) as a function of distance in pixels from the center of the detector.




FIG. 16. The PSF FWHM of the MRI after bakeout for 707 stars with peak SNR greater than 100 as a function of distance in pixels from the center of the detector.

and ITS are paraboloids of revolution opening toward the optics. Thus we see a paraboloidal defocus term as a function of distance from the optic axis, which was aligned to be nearly coincident with the center of the detector. There is also a limiting best focus defined by diffraction by the optics and signal electron diffusion within the detector, which tends to flatten the curve in the area closest to best focus. The optical design limits the effects of coma to <0.25 pixels at the corner of the field making its effect negligible. From this analysis, it can be seen that the MRI is very near best focus within about 300 pixels of the center of the detector but progressively further out of focus as the distance from the center of the detector increases. The design of both the MRI and the ITS presumed that focus would be set to provide best focus in a ring about 250 pixels from the center of the detector, to provide reasonably good focus over the entire detector. Comparison of this measurement with a model provided by the optical designer (James Baer of Ball Aerospace), indicated that the MRI was out of focus by approximately the amount and in the direction predicted by the uptake of humidity by the optics mounting structure, and therefore that baking the structure to remove the water should provide a restoration of the desired focus.

A similar image of the same star cluster was also taken using the ITS camera, with similar results, except that the ITS appeared to be closer to optimal focus, as expected if the moisture had not penetrated to as great a degree to the ITS, located in the interior of the spacecraft. Therefore both the MRI and ITS were baked for several days at an elevated temperature to drive off water, allowed to cool, and then images of a star cluster were taken again.

The results for the MRI after bakeout are shown in Fig. 16. As can be seen, the relatively flat area close to the center of the detector has expanded outward by over a hundred pixels, with parabolic defocus once again appearing at larger distances from the center, but with a smaller amount of defocus at the largest distances.

The similar plots for the ITS before and after bakeout are shown in Figs. 17 and 18. Note that there is essentially no change due to bakeout, indicating that very little moisture

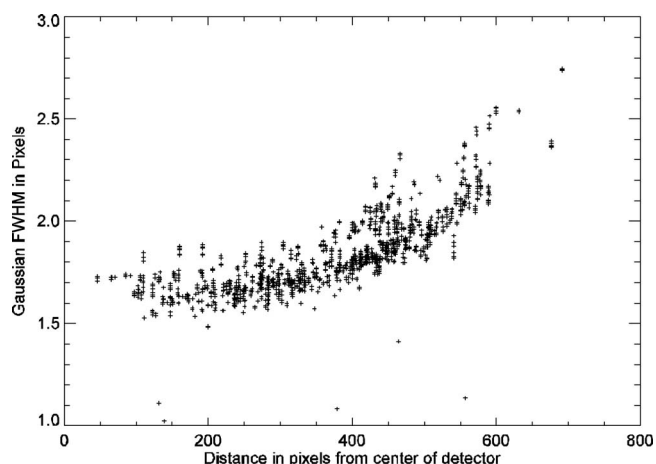


FIG. 17. The PSF FWHM of the ITS from the quick alignment images of 1638 stars with peak SNR greater than 100 as a function of distance in pixels from the center of the detector.

was taken up by the ITS structure, buried deep inside the Impactor spacecraft, which was largely inside the Flyby spacecraft.

Determination of the MRI and ITS PSFs out beyond a few pixels from their centers can be addressed to some degree using long exposures of bright stars. Some images like this were acquired in flight in the clear filter only, but analysis to determine extended PSFs has not yet been done.

2. HRI point spread function

In the first images taken after launch with the HRI, the star images were clearly out of focus, showing a ringlike structure with an outer diameter of 11.59 pixels. However, some degree of defocus was expected due to the uptake of humidity by the graphite composite telescope structure during the hours before launch. Unfortunately, bakeout reduced the degree of defocus by only a small amount, reducing the outer diameter of the ring to about 8.95 pixels and leaving the HRI best focus image location several millimeters in front of the detector. This was a surprise because the focus had changed by the equivalent of 10 pixels of image diameter during ground tests of baking out moisture. Because no focus mechanism had been included in the final design of the system, this degree of defocus continued for the remainder of the mission.

Considerable effort and observation time were dedicated to gathering data both before and after encounter to characterize the HRI PSF through all filters. The results are shown in Fig. 19. The PSF is seen to have threefold symmetry reflecting the three-point mounting of both the primary and secondary mirrors. Through all filters there is a similar amount of defocus, with PSF FWHMs ranging between 8.7 and 9.1 depending on the filter, but the longer-wavelength filters also show considerable amounts of light scattering, which tend to fill in the center of the PSF to a larger degree than is seen in the shorter-wavelength filters. The additional scattering in the long-wavelength filters is attributed to the colored glasses used to help provide out-of-band blocking having greater numbers of scattering centers than those used in the shorter-wavelength filters.




FIG. 18. The PSF FWHM of the ITS star images after bakeout for 710 stars with peak SNR greater than 100 as a function of distance in pixels from the center of the detector.

Besides providing a better representation of the actual PSF than a single critically sampled image can provide, sub-sampled PSFs can be used to restore images with subpixel resolution. Sub-pixel resolution may result in better restorations than just interpolating in solution space since it actually adds new information from the finer sampled PSF into the restoration process.

The calibrated data sets used in this PSF characterization included 45 exposures (5 exposures in each of nine filters) on 14 June 2005 and again on 9 July 2005, plus five exposures on 2 July 2005. The stars used as point sources were 16 Cyg A and B, both solar analogs, with exposure times of 2 s for the clear filter images, 8 s for green, and 10 s for all other filters. The subpixel PSFs were derived using 20 or 22 nearly randomly positioned images, obtained by using a slow drift of the spacecraft pointing over a range of several pixels both vertically and horizontally. The technique used is termed “drizzle” and was developed for the first Hubble Deep Field North and exists as a set of tasks in the STSDAS package *analysis.drither*.⁸ It is used routinely as an integral part of the Hubble Space Telescope Advanced Camera for surveys calibration pipeline.

The pre- and post-encounter data sets were independently processed to provide two measurements of PSFs in each filter. Comparison of the two sets of measurements supports the following conclusions:

- Using 16 Cyg A and B combined PSFs appears to be reasonable.
- Significant changes occurred between pre-encounter and post-encounter PSFs (systematic residuals larger than allowed by noise). This is especially true on the clear filter in position 6 (see Table II). There is also evidence of a small focus shift (ringed residuals).
- The post-encounter set is closer in time to encounter, while the pre-encounter set has slightly better coverage in fractional pixel offsets, thus slightly better resolution (for combined A+B stars).

Generally, it appears better to use the post-encounter set of




FIG. 19. The measured PSFs through the nine filters of the HRI, as determined by a short exposure of a bright star (to minimize pointing drift during the exposure). The second column displays the logarithmic stretch of a single image of the PSF intensity to show in detail the extended low-intensity part of the PSF. The third column shows a single image of the measured PSF intensity, scaled to linearly display minimum brightness as black and maximum as white. The fourth column displays a 3×3 subpixel sampled drizzled PSF derived by combining ten nearly randomly positioned images of the stars 16 Cyg A and B taken on 14 June 2005 using the STSDAS drizzle function (see text). It appears that the PSF for the Clear-6 filter is contaminated by significant drift during at least one of the exposures, presumably the very first one, which was the reference image.

PSF determinations, which varies only slightly from those illustrated in Fig. 19.

3. HRI deconvolution

Image restoration algorithms were investigated for correcting the blurred HRI PSF to improve the spatial resolution of the images via deconvolution. Under any condition, the PSF of an optical system is finite sized. In a linear and shift-invariant imaging system, the recorded image is the convolution between the “real image” and the system PSF, plus additive noise. The inverse problem is usually to divide the observed image by the system PSF in Fourier space. However, this problem is usually ill conditioned because of the loss of information due to the finite frequency distribution of the PSF and the existence of noise having a very wide frequency distribution. A number of inverse algorithms have been developed, such as maximum entropy, least-square techniques, the Richardson–Lucy algorithm, and Pixon-based algorithms, and have been widely used in many fields.^{9,10}

The algorithm adopted in our pipeline calibration process to restore the degraded HRI images is the so-called constrained least-square method,⁹ which is a linear inverse-filter type algorithm, much faster than iterative algorithms such as Richardson–Lucy, but less effective in terms of suppressing noise and artifacts. The objective in this constrained least-square solution is to minimize the norm between the solution and a trial image, while keeping constant the norm of the difference between the solution convolved with the PSF and the actual image; this difference is presumed to be the noise. The trial image is taken as the original image with all negative-valued pixels set to 0. Any imperfection in the images, such as cosmic rays, bad pixels and lines, or dust specks can be problems that affect photometry and generate PSF-like or linear artifacts; therefore, the images have to be pretreated in order for deconvolution to generate good quality images. The finite pixel size in both the images and the PSFs, as well as the nonperfect sampling of the PSFs, are the origin of ringing artifacts appearing as overshooting and undershooting at the sharp edges of deconvolved images, but the strength of these artifacts is usually less than 5% of the local brightness contrast. Total photometry is conserved, but negative pixel values higher than the noise level could exist at undershooting places. An example showing the improved resolution in convolved images is shown in Fig. 20. Readers are referred to a companion paper discussing the various methods used for deconvolution of DI images for more details.¹¹

C. Radiometric calibration

1. Linearity

The response linearity of the CCD detectors to varying light input in the visible cameras was measured during ground calibrations. Response was shown to be very linear over the full dynamic range (noise level to full well) of the detectors, as was expected from high-quality CCD devices. Figure 21 shows the results of fitting a straight line to the




FIG. 20. An example of the deconvolution step in the calibration pipeline applied to a HRI image through the clear filter in wheel position 6. Left panel shows the original image; right panel shows the deconvolved image. Small-scale features that are blurred in the left panel can be distinguished in the right panel. Also note that noise is amplified by deconvolution. The horizontal dark line at the center of the deconvolved image is a deconvolution artifact caused by a very slight dark line (1%–2% level) in the original image due to flat-field residual at the edge between the upper and lower pairs of CCD quadrants.

HRI vis response versus integration time averaged over broad areas of flat-field images taken of a constant-radiance source during TV2. The fit minimizes percentage errors in response using a least-squares criterion. Points with signal levels above 15 000 DN were eliminated from the fit because they reflect some fraction of the pixels hitting their full-well limits. Fits are shown for each of the four quadrants of the CCD with their DN levels offset progressively by 1000 DN for visibility. Differences in gain of up to 4% exist between quadrants. Figure 22 shows the percent residuals to linear fits of HRI vis TV4 data in each quadrant taken through the clear filter. These residuals are typical of linear fits for other cases. Residuals are <1% for all signal levels between 15 and 15 000 DN above the bias. Residuals tend to be negative at low and high-signal levels and positive at intermediate signal levels. A check of a representative set of individual pixels showed that their linearity functions were indistinguishable from those of the average over broad areas of pixels.

The same analysis of the MRI and ITS cameras gave similar results with only slightly higher residuals in a few MRI cases (<2%). These results justify the use of a simple linear radiometric calibration function in the vis pipeline processing.

Response linearity for the HRI vis and MRI was checked in flight using lunar images. Instrument responses were analyzed for four regions on the Moon in clear-filter images using 15 different exposure levels in each camera. Corrections for readout smear were applied (see Sec. V C 4). Linear fits to the corrected response versus integration time yielded residuals of <2% for HRI vis and <1% for MRI, confirming that the inflight linearity remained excellent for these cameras. This result was also verified later using internal stimulus lamp images of various exposure times for both cameras on the flyby spacecraft as well as for the ITS.

2. Gain/full well

The performance of the CCD cameras in terms of their system gain (ie., conversion factor from signal electrons to




FIG. 21. HRI vis response linearity. Results for the various quadrants are offset by 1000 DN for visibility.

output DN) and their full-well capacity (i.e., the maximum signal for which the detector response remains linear and charge is retained strictly within the pixel packet in which it was originally created) was determined using the photon transfer technique. This technique involves taking a number of exposures of a source with a relatively uniform radiance field and measuring the random noise level versus the signal level over the entire dynamic range of the detector. Such measurements were made during ground calibrations using images of integrating sphere apertures as well as in flight using the internal stimulus lamps.

The photon transfer technique employs the relationship between noise and signal given by

$$\text{noise}_e = \text{SQRT}(\text{signal}_e + \text{read noise}_e^2),$$

where all quantities are in units of electrons. Here the signal_e term represents the square of the shot noise in the signal due to photon statistics, and the read noise_e term is the inherent noise in reading out the CCD independent of any signal. In units of DN, this relationship becomes

$$\text{noise}_{\text{DN}} = \text{SQRT}(\text{signal}_{\text{DN}}/g + \text{read noise}_{\text{DN}}^2),$$

where g is the system gain in electrons/DN, and any quantization error is assumed negligible.

The random noise is measured in frames with relatively uniform signal across the array by subtracting two sequential




FIG. 22. Residuals to a linear fit of the HRI vis response.

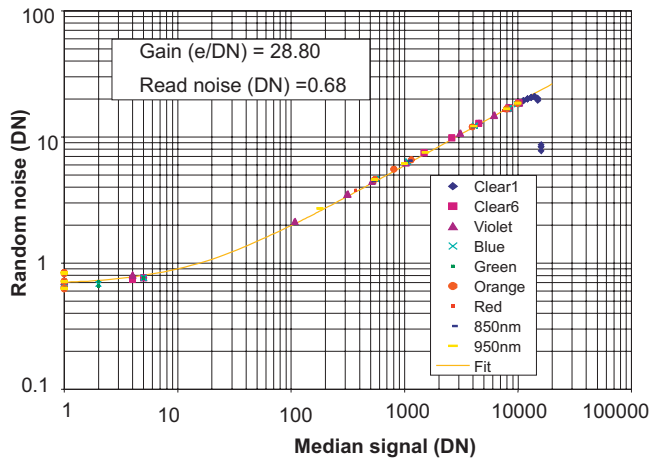


FIG. 23. (Color) Photon transfer data for HRI vis.

identical frames and calculating the standard deviation/ $\sqrt{2}$ in the difference image over a number of blemish-free subareas in the frame. The signal level in each subarea is determined by taking its dark-subtracted median value.

Photon transfer curves were determined using prelaunch data for each CCD quadrant in each filter position for each camera. Figure 23 shows an example of the results for Quad C of the HRI vis. A fit of the equation above was found by solving for read noise and g and is plotted in the figure.

As expected, no systematic differences were observed as a function of filter position in a given camera. Measurable differences were observed for the different quadrants of the different cameras, however. The prelaunch results were confirmed in flight using stimulus lamp images. Table IX lists the best-fit values for gain and read noise found in each case. The inflight data are not quite as accurate as the prelaunch data because fewer frames could be taken; we adopt the prelaunch values as the best estimate of the inflight performance.

The photon transfer curves also provide a very sensitive measure of the detector full-well level. The signal level above which the random noise value starts to fall below the value predicted by the noise equation above indicates the signal level at which charge is no longer being kept totally isolated from pixel to pixel. This signal level can be taken as a conservative level at which some pixels have reached their

full-well limit. This limit is $\sim 12\,000$ DN ($\sim 350\,000$ electrons) for all cameras in all quadrants. Most pixels will continue to respond linearly to increased photon integration beyond this level, as can be seen in the previous discussion of linearity, where a less conservative full-well limit of $\sim 15\,000$ DN ($\sim 440\,000$ electrons) can be defined with linearity residuals remaining $<1\%$ for the average response over an extended area of the array.

The only direct use of the gain and read noise values in the calibration processing pipeline is in calculating the SNR in each pixel for the radiometric quality map of an inflight image (Sec. VII C). The absolute radiometric conversion factors only depend on the product of ($\text{gain} \times \text{CCD quantum efficiency} \times \text{pixel area}$), so knowing the gain independently is not necessary. The variation in ($\text{gain} \times \text{quantum efficiency} \times \text{pixel area}$) from quadrant to quadrant is calibrated out using the normalized flat-field files. Therefore, we use averaged quadrant-independent values for gain and read noise in calculating SNR. The full-well limit is used to flag pixels whose radiometric response is unreliable because their raw DN level exceeds either the conservative full-well limit determined from the photon transfer curves or the less conservative limit defined by loss of response linearity over an extended area.

3. Zero-exposure level

The bias, sometime also called offset, is due to an intentional electrical offset applied in the output signal chain to ensure that DN levels do not fall below zero. The bias adds a constant signal to each image that must be subtracted to get the real signal from a given source. The bias varies with the CCD temperature and the signal-chain electronics temperature.

The HRI, MRI and ITS CCDs all have four quadrants, with serial overclocked (SOC) pixels that can be used to calculate the bias in each quadrant.³ This is the most robust method since bias varies with temperature (primarily that of the CCD) and time. For a fixed CCD temperature, variations in bias of <10 DN were observed in flight; an increase in bias of about 25 DN was observed when the HRI CCD temperature increased to ~ -40 °C from its normal operating temperature of ~ -103 °C. For a given frame, the bias from the SOC value is within <1 DN of the applicable bias in the

TABLE IX. Best-fit gain factor and read noise for each quadrant of each vis camera.

Camera	Test	Gain (e/DN)				Read noise (DN)			
		Quad A	Quad B	Quad C	Quad D	Quad A	Quad B	Quad C	Quad D
HRI vis	TV4	29.1	29.0	28.8	29.9	0.9	0.9	0.7	0.9
	April Cal	27.1	28.2	27.5	27.4	0.6	0.8	0.7	0.6
	May Cal	27.2	28.1	26.8	27.1	0.9	0.7	0.6	0.7
	June Cal	27.0	28.0	27.2	27.2	0.6	0.7	0.7	0.6
MRI	TV4	27.9	28.4	28.7	28.5	1.3	1.0	1.0	1.0
	April Cal	26.5	27.4	27.4	26.9	0.8	0.9	0.6	1.0
	May Cal	27.9	27.4	26.7	26.8	2.1	1.0	0.7	1.0
	June Cal	25.7	27.9	27.2	28	0.9	1.3	0.7	1.3
ITS	TV4	30.3	30.8	30.2	30.5	1.0	1.4	1.2	1.1
	June Cal	30.3	32.0	29.4	30.2	1.3	1.1	1.0	1.5

TABLE X. Bias (DN) for modes 7 and 8, where there are no serial overlocked pixels, and values of the parameters C (nA/cm²) and E_0 (eV) used for the dark current model (see equation in text).

		Quadrant A	Quadrant B	Quadrant C	Quadrant D
HRI vis	Bias	376	367	374	361
	Dark (C ; E_0)		0.58; 1.215		
MRI	Bias	356	358	354	359
	Dark (C , E_0)		0.58; 1.201		
ITS	Bias	393	428	384	404
	Dark (C , E_0)	0.48; 1.213	0.53; 1.226	0.49; 1.213	0.36; 1.209

active array. Therefore, the SOC value is used for bias subtraction in our processing pipeline. The actual bias in the active array has been observed to vary by up to 1 DN on time scales shorter than a line readout time.

Unfortunately, the above method cannot be used for CCD operating modes 7 and 8 (64×64 subframe modes) since they do not produce SOCs. As a consequence, for in-flight data in modes 7 and 8, the bias in each quadrant is taken to be the median value of all the SOC pixels in each quadrant for all the other modes, derived from the May in-flight science calibration (Table X), and this value is used for the pipeline bias subtraction.

The dark current level also depends on the CCD temperature. Even without any incident photons, thermal agitation within the CCD induces a small current (signal) that is added to the real signal when the CCD is read. The dark current increases rapidly with temperature and must be subtracted from the measured signal.

For obvious time and storage reasons, we could not take dark frames during encounter for all CCD and electronics temperatures and exposure times. For in-flight data, we rely on a dark current model, based on the physical characteristics of silicon, and widely used in the past. The model was successfully applied to ground and in-flight science calibration data

$$\text{dark current (DN/s)} = Cp_x^2 2.55 \times 10^7 T^{1.5} \times \exp(-E_g q / 2kT),$$

where $E_g = E_0 - (7.021 \times 10^{-4} T^2) / (1108 + T)$ is the silicon band gap (eV), C is current density (nA/cm²), p_x is pixel size (mm), T is CCD temperature (K), q is electron charge (Coulomb), and k is Boltzmann's constant (J/K). The two unknown parameters are E_0 and C , and we used ground and in-flight science calibrations to determine them (Table X). Values of C ranging from 0.36 to 0.58 nA/cm² indicate good quality CCDs for all three cameras. Dark current values derived from the equation are subtracted along with the bias in our pipeline processing.

Figure 24 illustrates the dark current as a function of CCD temperature for HRI. We obtain similar results (but slightly different values) for MRI and ITS. HRI operates at about -103 °C, where the dark current is negligible, $\sim 3 \times 10^{-5}$ DN/s. MRI operates at about -96 °C, where the dark current is also usually negligible, $\sim 2 \times 10^{-4}$ DN/s. ITS operates at about -30 °C, where the dark current is not negligible, ranging from 7 to 10 DN/s depending on the quadrant. In practice, the dark current correction model is only

required for ITS. A more detailed analysis indicates that the dark current in ITS is line and sample dependent, but we do not correct for this effect, so that our model is good within ± 1.5 DN at short exposure time (100 ms) and ± 4 DN at long exposure time (2000 ms). Concerning ITS bias, we studied its variations during the last 24 h before impact for the different modes. We notice some short time variations of ± 1 DN for uncompressed data and up to ± 3 DN for compressed data for a given mode. However, within those error bars, the bias values agree with Table X, so that the May calibration still gives a good approximation for the mode 7 and 8 biases. Overall, the bias and dark subtraction add an uncertainty of less than 2 DN for HRI and MRI and up to 7 DN for ITS.

4. Frame transfer smear

The DI CCDs are split-frame, frame-transfer devices. Image integration time is controlled by the time between successive frame transfers out of the active area of the array into the masked storage regions. While the cameras do include "shutters," these are more properly called light blockers. They do not control the signal integration time. They serve primarily to keep bright scenes from continuing to produce signal in the active portion of the array while image storage and readout are taking place, which could cause charge bleeding into the storage regions before frame readout is complete.

In this type of CCD, charge integration takes place not only during the commanded integration period when the CCD parallel clocks are held fixed and the image is collected




FIG. 24. Dark current as a function of CCD temperature for quadrant C of the HRI instrument. The model is the solid line and the data are the squares.




FIG. 25. 100 ms exposure, dark subtracted, Gaussian stretched image showing frame transfer smear.

but also while the charge packets are being shifted across the array during the last frame flush prior to the integration period and during the frame transfer period after integration is complete. The extraneous charge that is collected in a frame during the last frame flush and the frame transfer we call frame transfer smear. For both the top and bottom halves of the split-frame CCD, all charge packets in a column within a frame pass under the scene being imaged on that column and collect charge as they do so. The line shift rate during the last flush prior to stopping the parallel clocks for integration is the same as that used for the frame transfer after integration. So given no pointing instabilities during a frame time, each pixel in the top and bottom halves of a column will accumulate the same amount of extraneous frame transfer signal. The time required to shift a row across the unmasked region of the array is 5.46 ms; therefore, frame transfer smear can contribute >1% of the total signal for exposure times of 500 ms or less. Figure 25 shows an example of the effect of frame transfer smear in a 100 ms mode 1 image. The DN values in the bright areas range from 100 to 400 DN. The frame transfer smear signals (vertical streaks) are at the 1–4 DN level.

For proper radiometric calibration of an image, this excess charge needs to be subtracted out. Two methods for determining the amount of frame transfer smear signal in each half-column of an image have been developed. The first involves using the parallel overlocked (POC) pixel values to determine the smear signal. The POC pixels accumulate frame transfer smear charge during the frame transfer process only; their charge packets have not even been created by the time image integration ends. So they should provide a fairly accurate measure of the frame transfer smear signal in each half-column. Note the frame transfer smear signal extending into the POCs (top and bottom eight rows) of Fig. 25. Four rows are summed in the CCD serial register to create each POC value, so the POC values need to be divided by four prior to subtracting them from the image columns.




FIG. 26. Version of Fig. 25 corrected for frame transfer smear using parallel overlocked pixel value.

This correction technique can only be used in camera modes 1–6; in modes 7 and 8 no POCs are created. The number of POC rows with reliable values varies with mode; in mode 1 only the first and last five rows have good values, in modes 2–4 the first and last four rows are good, and in modes 5–6 only the first and last two rows are good.

The second frame-transfer smear correction algorithm is used in modes 7 and 8. It involves estimating the frame transfer smear signal as a fraction of the average dark-subtracted signal in each column of the actual image area, where the fraction depends on the ratio of the frame transfer time to the total signal integration time. This approach can be subject to significant errors if (a) there are saturated signals in the image area, or (b) there are significant differences between the scene brightness in the top and bottom few rows of the returned subframe and that in the rows outside the subframe that are not returned.

Figure 26 shows the results of correcting the frame transfer smear seen in Fig. 25 using the POC values. The residual frame transfer smear correction errors are about 1 DN.

5. Absolute sensitivity

The goal of radiometric calibration is to be able to accurately measure the radiance of a scene at each point in its image. The pertinent equation for the single-pixel response to a spatially resolved source radiance assuming a linear instrument response function is

$$\text{DN} = tg\Omega a \int S_{\lambda} T_{\lambda} F_{\lambda} N_{\lambda} P_{\lambda} d\lambda + \text{DN}_0(t),$$

where DN is the linear instrument response in DN, t is the integration time (s), g is the gain setting in DN/signal electron, Ω is the solid angle of optics equals area of entrance pupil/effective fl^2 (sr), a is the area of pixel (m^2), λ is the wavelength of light (μm), S_{λ} is the detector quantum efficiency (signal electrons / photon) at wavelength λ , T_{λ} is the

transmission of the optics at wavelength λ , F_λ is the transmission of the filter used at wavelength λ , N_λ is the radiance of the scene ($\text{W m}^{-2} \text{sr}^{-1} \mu\text{m}^{-1}$), P_λ is equal to $\lambda/1.98648 \times 10^{-19}$ photons $\text{W}^{-1} \text{s}^{-1}$, $d\lambda$ is the wavelength interval (μm), and $\text{DN}_0(t)$ is the zero-exposure DN output for integration time t .

Making radiometric measurements with the DI cameras requires inverting this equation to solve for the scene radiance, N . In the case where we do not know *a priori* the spectral shape of the radiance curve, we assume that the radiance is constant over a given filter bandpass f and solve for

$$N_f = \frac{\text{DN} - \text{DN}_0(t)}{t g \Omega a \int S_\lambda T_\lambda F_\lambda P_\lambda d\lambda}.$$

For solar-illuminated scenes, we can also solve for the reflectance of the scene, assuming it is constant over the filter bandpass. Reflectance $R(\lambda)$, also known as I/F , is related to radiance by

$$R(\lambda) = \frac{I(\lambda)}{F(\lambda)} = \frac{\pi N(\lambda)}{H(\lambda)}$$

where $H(\lambda) = \pi F(\lambda)$ is the solar irradiance incident on the reflecting surface at wavelength λ , and $I(\lambda) = N(\lambda)$. Thus for filter f ,

$$R_f = \frac{\pi[\text{DN} - \text{DN}_0(t)]}{t g \Omega a \int S_\lambda T_\lambda F_\lambda H_\lambda P_\lambda d\lambda}.$$

While an initial derivation of the vis camera absolute sensitivities (i.e., $1/[g\Omega a \int S_\lambda T_\lambda F_\lambda P_\lambda d\lambda]$ for each camera and each filter) was done using ground-based radiance sources during TV calibrations, the most accurate determinations were done with in-flight calibrations using photometric standard stars. The stars used for this purpose are listed in Table VI. Absolute radiometric response calibrations were carried out four times in flight—in April, May, June, and July, 2005. Note that when deriving the spectral response calibration using point sources as opposed to spatially resolved sources, the response equation is more usefully written as

$$[\text{DN} - \text{DN}_0(t)]_{\text{total}} = t g A \int S_\lambda T_\lambda F_\lambda I_\lambda P_\lambda d\lambda,$$

where $[\text{DN} - \text{DN}_0(t)]_{\text{total}}$ is the total dark-subtracted signal from the star over all pixels, A is the area of optics entrance pupil (m^2), and I_λ is the irradiance of the starlight at the camera entrance aperture ($\text{W m}^{-2} \mu\text{m}^{-1}$).

The objective of our in-flight stellar calibration is then to determine the $gAS_\lambda T_\lambda F_\lambda$ that causes the above equation to be true for any star's I_λ . From that, we can derive the absolute camera sensitivity factors required to convert from dark-subtracted DN/ms to radiance or reflectance for spatially resolved sources. $S_\lambda T_\lambda F_\lambda$ is assumed to be the same across the entire CCD array; any pixel-to-pixel or quadrant-to-quadrant differences in response are corrected by a separate flat-field correction as discussed in Sec. V C 6.

The absolute spectral irradiances from 200 to 1200 nm of the standard stars observed were obtained from a combination of data and model spectra.^{12–18} Morel and Magnenat¹² provided U , B , V , R , I , and J absolute photometry of most of

our standards; supplementary photometric measurements were obtained from Kharchenko,¹⁶ Monet *et al.*,¹⁷ and Neckel *et al.*¹⁸ Burnashev¹⁴ and Kharitonov *et al.*,¹⁵ published absolute irradiance spectra but only shortward of 750 nm. Pickles'¹³ stellar irradiance model spectra for the appropriate spectral type and luminosity, which span the entire 115–2500 nm range at a resolution of 0.5 nm, were scaled to fit the data points of Morel and Magnenat and the spectra of Burnashev or Kharitonov *et al.* for each star to provide interpolations with higher spectral resolution. These scaled irradiance spectra were then used as our absolute references. Figure 27 shows the absolute irradiance spectra adopted for three of our standard stars. The uncertainties in our reference spectra are estimated to be $<5\%$ for a given star at a given wavelength.

Using these standard sources, the DI vis camera in-flight absolute spectral sensitivities were determined in a best-fit sense. Very good repeatability was seen throughout the four in-flight calibrations spaced over a 3-month period bracketing the encounter event. However, the in-flight sensitivities showed significant differences from the prelaunch values, which had been derived from camera responses to an integrating sphere flat-field light source, whose spectral radiance had been absolutely calibrated, viewed through the TV chamber window. Part of the reason for these differences might have been the need to extrapolate the integrating sphere radiance and the chamber window transmission below 350 nm, since their measurements cut off there. In addition, the chamber window transmission was estimated by adjusting the measured transmission of another similar, but thinner, chamber window rather than measured directly. The in-flight data suggest that the absolute calibration of the integrating sphere source radiance and/or the estimate of the chamber window transmission was systematically high.

Figure 28 plots the ratio of the HRI and MRI camera responses in terms of total DN/ms through each filter for each in-flight calibration ratioed to those of the May calibration; the ratios are plotted versus the effective wavelength of each filter. The ratios are all very near 1; the ratios for the different calibration dates are offset by 0.5 units in the plot for visibility. The variations are seen to be $<10\%$ with no systematic trends. The variation in ITS response between May and June (the only two successful ITS calibrations) was $<2\%$.

The scatter in the measurements is larger than the formal errors that would result just from random noise in the images. One probable cause of the increased scatter is uncompensated cosmic ray signal either within the star images themselves or in the surrounding background area used to determine the offset level to subtract. Another probable cause is inaccuracies in removing the frame transfer smear signal. The data suggest that any individual star calibration measurement is uncertain to about the 5% level for the MRI and to about the 2% level for the HRI (the defocused HRI means that much higher total star signals can be achieved without saturating the peak pixel, which results in much higher SNRs than are possible with the MRI). The greater degree of scatter in the April ratios is due to using a mix of short and long exposure times and no repeat frames for this first calibration

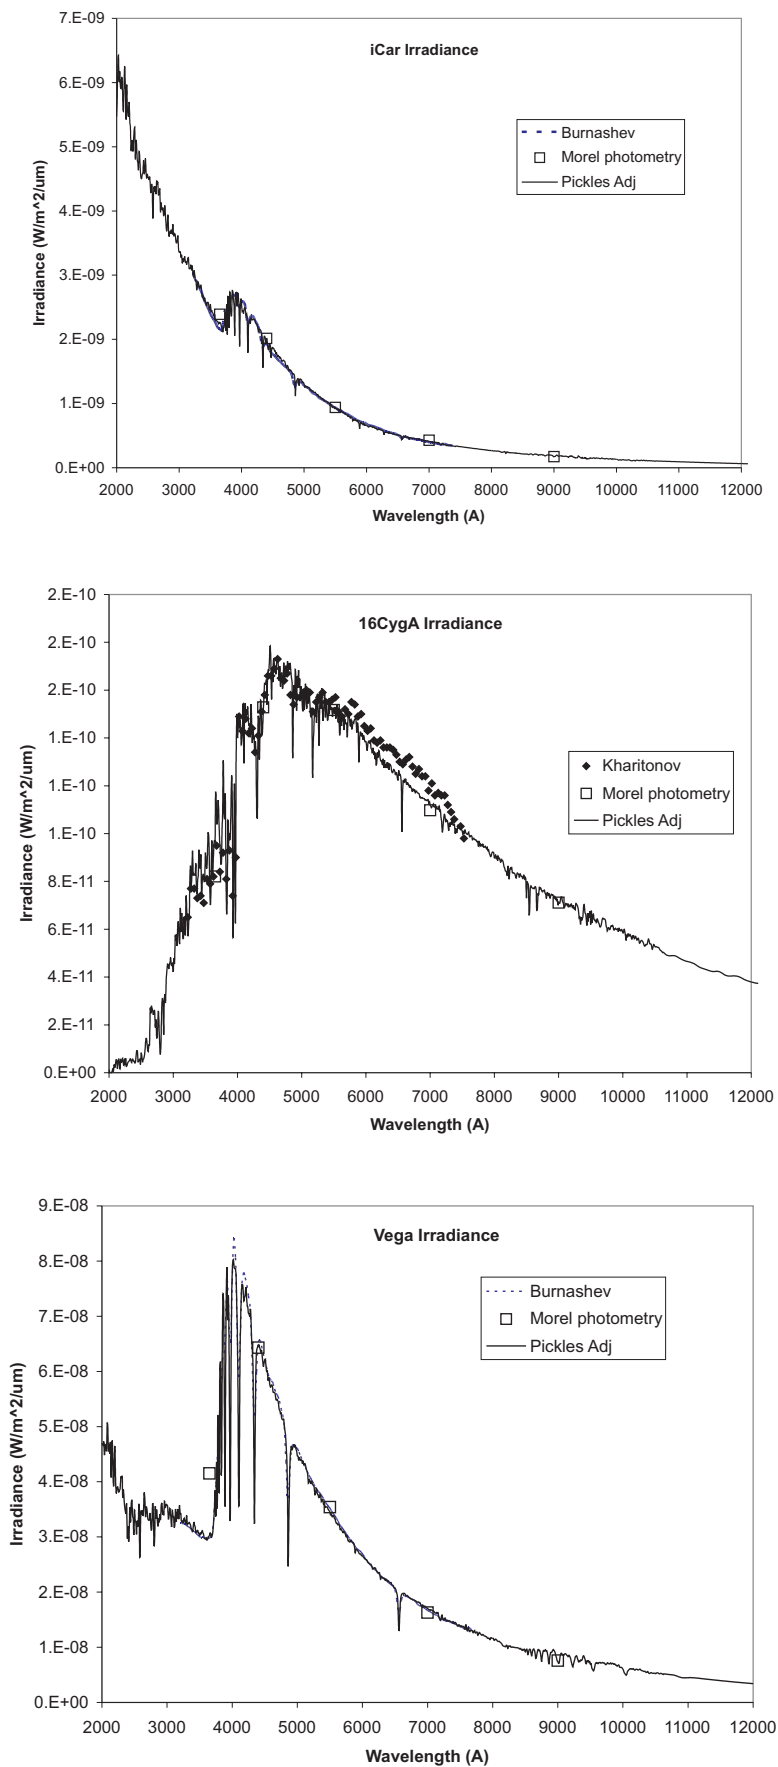


FIG. 27. (Color) Absolute spectral irradiance of selected photometric standard stars.




FIG. 28. (Color) Camera response during inflight star calibrations ratioed to those obtained in May 2005 (HRI, top; MRI, bottom). When error bars cannot be seen, they are smaller than the data symbol.

and to severely underexposing the HRI images because the exposure times were set assuming an in-focus system.

Figure 29 plots the ratio of the observed in-flight response to the various photometric standards averaged over all four calibrations to that predicted based on the prelaunch calibration; as in Fig. 28, ratios are plotted versus the effective wavelength of each filter. For all three cameras, the in-flight response was about 25% higher than predicted in the clear and midwavelength filters, with response at the UV end being elevated the most and response near $1 \mu\text{m}$ being close to the predicted level.

The in-flight response rates were used to adjust the spectral response math models for each camera that were developed prior to launch. The CCD quantum efficiency (QE) estimates were adjusted first because they were the least well-known aspect of the model. For MRI, additional adjustments were made to the peak transmissions of the three

narrow-band UV filters. These adjustments were needed to reconcile the star responses in the clear filter for all spectral types with those through the UV filters. We determined that in order to reduce the discrepancy in the model predictions in the 387-nm MRI filter for β Hya compared to those for the other stars, we require either that the filter bandpass be shifted 2 nm longward from its prelaunch value or that the deep absorption band in the β Hya spectrum around 385 nm be shifted 2 nm shortward. What combination of these adjustments is actually correct cannot be determined. In general, calibration of the MRI response in the UV filters is problematic. Many of the stellar sources have narrow absorptions in their spectra below 400 nm, and high-spectral-resolution irradiance measurements of these sources are not available (we are relying on Pickles' model spectra for the high-resolution features). In addition, the clear-filter MRI responses, particularly to blue stars, includes added uncertainty




FIG. 29. Ratios of actual inflight stellar response rates to those predicted using the prelaunch camera models.

because nothing in the system cuts off UV response, and the component-level calibrations of detector QE were not made below 350 nm nor were clear-filter transmission data taken below 300 nm. Deriving a model that reconciles the clear and UV-filter MRI predictions with the observed responses is not very deterministic. Therefore, absolute radiometric calibration accuracy in the MRI UV filters is no better than 20%. The HRI escapes similar problems because its filter bandpasses are all broad, and the beamsplitter in the HRI does not allow light shortward of 300 nm to reach the detector.

The other terms in the camera response model equation were fixed at their nominal values of $g=1$ DN/30 electrons, $a=4.41 \times 10^{-10}$ m², and $\Omega=6.41 \times 10^{-4}$ sr for the HRI and 2.57×10^{-3} sr for the MRI and ITS. T_λ was taken to be the

value determined from prelaunch component-level calibrations. A few of the filters have red leaks longward of 1 μ m, the worst being the HRI 350 nm filter, with less severe leaks for the HRI 550, 650, and 850 nm filters. Moreover, while the long-wavelength leak in the MRI 309 nm filter is no worse than 10^{-5} , the system response rate is so low in the main bandpass at this short wavelength that the out-of-band response can be non-negligible in this filter. While these leaks are not significant in most cases, they need to be accounted for when doing the calibrations.

With adjustments to the detector QEs and the MRI UV filters in the camera models, the calculated response rates could be made to match all the observations to within 20%. Excluding the MRI UV filters, the matches were good to




FIG. 30. Ratios of actual inflight stellar response rates to those predicted using the adjusted camera models.

TABLE XI. Radiometric conversion constants and effective wavelengths for each filter in each vis camera.

HRI spectral response model results				
Filter	Gray radiance conversion constant ($\text{W m}^{-2} \mu\text{m}^{-1} \text{sr}^{-1}$)/DN/ms	Effective wavelength (nm)	Solar-illumination reflectance conversion constant I/F at 1 AU/(DN/ms)	Effective wavelength (nm)
Clear-1	0.117 4	653.02	2.508×10^{-4}	624.52
Clear-6	0.117 5	653.00	2.511×10^{-4}	624.50
350-4	2.867	375.39	8.737×10^{-3}	375.15
450-2	0.962 2	452.48	1.611×10^{-3}	454.19
550-3	0.594 9	550.43	1.010×10^{-3}	549.88
650-9	0.586 3	648.44	1.154×10^{-3}	646.63
750-7	0.664 6	746.18	1.634×10^{-3}	744.40
850-8	1.014	844.94	3.105×10^{-3}	843.24
950-5	1.931	951.82	7.335×10^{-3}	948.71
MRI spectral response model results				
Clear-1	0.034 07	626.07	7.505×10^{-5}	610.28
Clear-6	0.034 11	626.08	7.513×10^{-5}	610.29
309-9	30.02	309.48	1.026×10^{-1}	309.63
345-8	13.42	345.41	4.084×10^{-2}	345.38
387-7	10.36	388.80	2.617×10^{-2}	388.97
514-2	1.759	515.31	3.002×10^{-3}	515.20
526-3	3.676	525.95	6.203×10^{-3}	525.92
750-4	0.212 5	746.00	5.223×10^{-4}	744.23
950-5	0.587 9	957.35	2.258×10^{-3}	953.31
ITS spectral response model results				
None	0.035 51	626.07	7.822×10^{-5}	610.28

better than 10% with a few exceptions. Most of the residual scatter is attributed to uncertainties in the absolute spectral irradiances of the various sources used and to uncorrected cosmic ray and readout smear signal. Some systematic discrepancies were observed; for example, Canopus appears to be redder than indicated by the Ref. 12 spectrophotometric data by 15%–20% at 950 nm compared to 350 nm. This trend is seen in both cameras across all calibrations. However, by averaging over all the stellar sources, we estimate that the absolute calibration is good to 5% for the HRI for all filters except for 950-nm, which is only good to 10%, and to 10% for MRI except for the UV filters, which are only good to 20%. Figure 30 plots the ratio of the observed stellar responses to the revised modeled responses and shows good agreement. The three data points for each filter are from each of the three calibration sessions in May, June, and July. The error bars represent the standard deviation in the measurements of each star within each calibration session (typically eight repeat frames were taken of each star in each filter during each calibration).

The revised camera response models were used to calculate the camera predicted response rates to spatially resolved sources with spectral radiances equal to $1 \text{ W m}^{-2} \text{sr}^{-1} \mu\text{m}^{-1}$ at all wavelengths and equal to that of a surface with a reflectivity of 1.0 at all wavelengths illuminated by the Sun at a distance of 1 AU. The solar irradiance at 1 AU used is the ASTM E490 model.¹⁹ Figure 31 plots the modeled response rate per pixel for each camera/filter combination when viewing an extended source of radiance

$1 \text{ W m}^{-2} \text{sr}^{-1} \mu\text{m}^{-1}$ at all wavelengths. The resulting radiometric conversion constants and effective wavelengths for each filter in each camera are given in Table XI.

The relationship between the total integrated DN/s for a star and its absolute visual magnitude, V , have also been derived for HRI and MRI in the clear 1 filter and for ITS. These relationships²⁰ are

HRI:

$$V = m + 12.003 - 0.146(V - R),$$

$$V = m + 12.004 - 0.127(B - V).$$

MRI:

$$V = m + 17.274 - 0.459(V - R) + 1.022(V - R)^2,$$

$$V = m + 17.304 - 0.381(B - V) + 0.374(B - V)^2.$$

ITS:

$$V = m + 17.485 - 0.303(V - R),$$

$$V = m + 17.417 - 0.728(B - V) + 0.637(B - V)^2,$$

where m is the instrumental magnitude $= -2.5 \log_{10}$ (integrated DN/s), $V - R$ is the absolute magnitude difference in the Cousins²¹ V and R bands, and $B - V$ is the absolute magnitude difference in the Cousins B and V bands.

The residual scatter is about 0.08 magnitudes for HRI, 0.03 magnitudes for MRI, and 0.04 magnitudes for the ITS using $B - V$ (but 0.121 magnitudes using $V - R$).






FIG. 31. (Color) Modeled camera response rates to an extended gray source of radiance of $1 \text{ W m}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$: (a) HRI vis, (b) MRI, and (c) ITS.

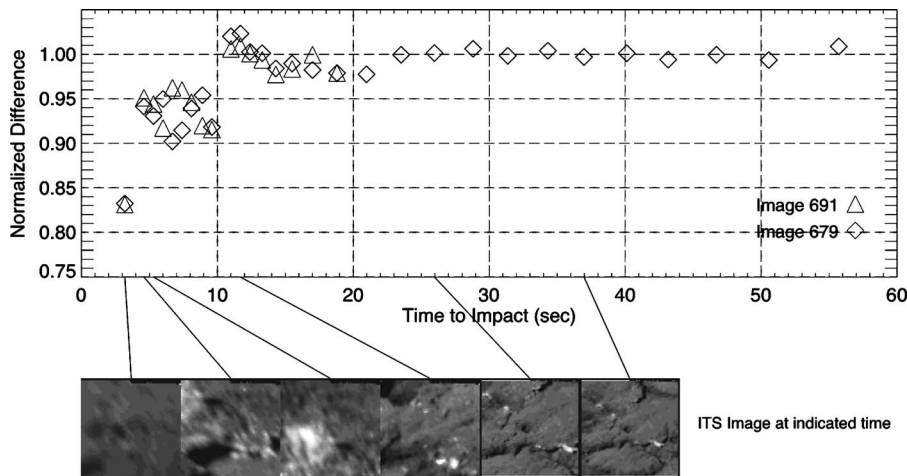


FIG. 32. Dropoff in ITS signal rate just prior to impact.

While it was obviously not possible to continue to monitor the ITS calibration all the way in to impact, we have examined the ITS nucleus images during its late approach phase to see if its radiometric response seemed to change. The observed ITS response rate starts to decrease about 10 s before impact and continues to decrease roughly linearly with range to the nucleus (see Fig. 32). The total reduction by the time of the last image (about 4 s before impact) is about 15%. The images used to produce Fig. 32 are small subframes of the full CCD; the nucleus is much larger than these subframes at this time, so the apparent brightness reduction seen cannot be due to an increased fraction of scattered light being lost outside of the image field. Another possibility is that this was due to degradation of the ITS primary mirror due to dust impacts. However, using empirical dust damage models used for HST,²² the amount of damage expected was two orders of magnitude less than could account for a 15% reduction in throughput. As the impactor is deflected due to dust impacts during approach, the ITS images parts of the nucleus at phase angles different than the average approach phase angle. The greatest deviations from the average approach intensity coincide with the two largest impactor deflections indicating that the change in intensity may simply be a measure of the local phase function on the nucleus.

6. Individual pixel response

Every CCD has a few bad pixels that need to be identified to fulfill the scientific objectives. For our analysis, we looked for 6 different types of bad pixels. The results are illustrated in Fig. 33.

- “Hot” pixels that reach saturation rapidly or that always saturate (marked as a on Fig. 33). This group includes all pixels that have a full well under 10 000 DN and all pixels that are always saturated with a signal higher than 16 000 DN (above full well) independent of the integration time. The first and last rows (numbers 8 and 1015) of the active CCD are hot pixels (i.e., excluding the overclocked pixels) for all cameras in every mode.
- “Cold” pixels that always have a low signal (marked as b on Fig. 33). This group includes all pixels that have a signal >5 times lower than the average signal of their

quadrant. There is a group of 33 cold pixels in the lower left corner of the HRI CCD that are returned in mode 1. There are no cold pixels for MRI or ITS.

- Pixels that vary significantly from one image to the other. This group includes all pixels that vary randomly and significantly from one image to the other one by more than 1%. There are no pixels of that type for any of the cameras.
- Overclocked pixels that behave strangely (marked as c on Fig. 33). This group includes overlocked pixels that have a signal too high or too low compared to their neighbors. This effect is seen in the overlocked pixels located in the rows that correspond to the rows of hot pixels for all cameras in every mode.
- Pixels with nonlinear response residuals above the 5% level rms (marked as d on Fig. 33). Pixels from row 1014, column 8–17, in HRI mode 1 have a nonlinear response. The response is linear for all other HRI pixels and all pixels in MRI and ITS.
- Pixels that have higher bias level (marked as e on Fig. 33). This group includes all pixels that have a bias at least 5% higher than the average bias of the quadrant. The first and last columns (Nos. 8 and 1015) of the active CCD have higher bias (i.e., excluding the overlocked pixels) for all cameras in every mode.

Recovering useful measurements from bad pixels is de-




FIG. 33. Bad pixels for HRI, MRI, and ITS in mode 1. The figure is not to scale for illustration purposes. See the text for the definition of marks a–e.

sirable but not always possible. For all cameras, the first and last rows and columns of the active area of the CCD are always bad. Since they comprise large areas (an entire row and/or column) on the edge of the detector, it is not possible to fill in reasonable DN values at these locations using the median value of the neighbor pixels, for example. Useful DN values are also unrecoverable for the group of cold pixels in the lower left corner of the HRI, since they span a large area. Concerning the nonlinear pixels in HRI, a more detailed analysis indicates that those pixels are linear with an accuracy of 12%–18%, which we consider usable. So, we use those pixels with caution in our scientific analysis with no correction.

The flat-field correction is an early step in the data reduction. Each pixel of the CCD detector has a somewhat different response to a uniform source. At small scale, the response is nonuniform because of the different characteristics of each pixel (e.g., QE). At larger scale, the response is nonuniform because mirrors and filters may have small dust particles on their surfaces that decrease the number of photons received by the detector in specific areas. Moreover, the optics themselves, even if excellent, induce small variations in the system sensitivity across the detector. We also have quadrant-to-quadrant variations since each quadrant has its own signal-chain electronics. To correct for all these effects, we divide all raw images by an image of a uniform source: the flat field. Since flat fielding corrects for optical and electronic effects, we must use a different flat-field correction for each filter and each mode of the instruments.

Flat fields were taken on the ground, during TV tests using a 65 in. integrating sphere, in every mode and every filter for each camera. The ability to correct for high-spatial-frequency response variations using those flat fields was then successfully validated in flight. To generate the flat field, we normalized to one the average of several flat fields with good SNR by dividing by the mean value of all pixels. In more detail, we used the following algorithm.

- Select a set of flat-field frames with good SNR, same mode and same filter.
- Subtract the *bias*, pixel by pixel, for each frame (Sec. V C 3).
- Subtract the *dark current*, quadrant by quadrant, for each frame (Sec. V C 3).
- Normalize each frame to 1, dividing by the mean value of all good pixels of the active CCD.
- Create an image *sdev*, which is the standard deviation $\sqrt{\langle x^2 \rangle - \langle x \rangle^2}$ of the value at each pixel across the n normalized frames. The image *sdev* is used to track and remove flat-field frames with cosmic rays.
- Create an image *flat field* which is the average of the n normalized frames (excluding those with cosmic rays).
- Set overclocked pixels of flat field to 1 (to leave them unchanged).
- Write in FITS format the resulting flat field.

We repeated the above algorithm for all modes and all filters of HRI, MRI, and ITS. The number of frames averaged together to generate each flat field varied from 2 to 16, depending on the instrument, mode, and filter. A typical

number is 8 frames. Figure 34 shows the flat fields for HRI, MRI, and ITS, in mode 1, using filter 1 (no filter for ITS). Figure 35 gives a profile for each of them along row 250 as an example.

From Fig. 34, we can identify large-scale nonuniform responses. The quadrant-to-quadrant correction is the most significant. Then, the two central rows of each CCD have a lower value, about 20% lower than the average flat field. This is a well-known effect on the DI CCD due to the 1/3 pixel gap between the top and bottom quadrant pairs. Next, the cold pixels in the lower left corner of the HRI CCD are clearly visible, as well as a ring, due to a dust particle on filter 1. Similar patterns are visible in other filters for HRI and MRI.

From Fig. 35, we can quantify the small-scale nonuniform responses. The pixel-to-pixel variation is $\sim 3\%$ rms for all cameras. A more detailed analysis indicates that this level of pixel-to-pixel variation is representative for all modes and all filters of all cameras, except for mode 9, where the pixel-to-pixel variation is $\sim 1\%$ (not shown here). We note that we did not use mode 9 during the entire mission.

We also performed an analysis to compare full-frame flat fields (mode 1) with subframe flat fields (modes 2, 3, 4, 5, 6, 7, and 8). The result is that the two different flat-field files are the same within errors of 2% for HRI and MRI and 4% for ITS. This gives us confidence in the stability and reliability of our CCD detectors when they operate in subframe modes.

The final test to validate and quantify the quality of our flat fields was to apply them to real data. For that purpose, we took a series of stimulator images during the in-flight science calibration phases. The stimulator is a light, located between the filter wheel and the CCD that provides diffuse illumination of the CCD. It cannot be used to test large-scale nonuniform response of the CCD, but it is extremely useful for observing the pixel-to-pixel variations. Figure 36 shows the profile of a stimulator image before and after flat-field correction for HRI (top), MRI (middle), and ITS (bottom). The improvement due to the flat field is significant, as the pixel-to-pixel variation is reduced in all cases. A more quantitative analysis indicates that before flat-field correction, the pixel-to-pixel variation is $\sim 3\%$, while after flat-field correction, it is down to $\sim 0.5\%$. The flat fields we generated from the TV tests on the ground apply very well to correcting inflight data.

For HRI vis, in order to recover the expected spatial resolution using deconvolution, we had to improve the pixel-to-pixel correction beyond the accuracy achieved from the ground-based calibration. For that purpose, we took a hundred HRI internal stimulator images in flight at the end of the post-impact lookback sequence. We combined these hundred stimulator images to derive improved pixel-to-pixel relative responses, using a resistant mean to reject cosmic rays and a high-pass filter to remove low-spatial-frequency intensity variations. Then, the previously generated flat fields were corrected with these new pixel-to-pixel relative-response variations. Figure 37 illustrates the results with the profile of a coma image corrected with the new and old flat-field files. To first order they look similar, but a close look indicates that




FIG. 34. HRI (top), MRI (middle), and ITS (bottom) flat field for mode 1 and filter 1 (no filter for ITS). Each flat field is a 1024×1024 pixel frame covering all four quadrants. Data span the range from 0.8 (black) to 1.2 (white) in this display with the mean response normalized to 1.0. The upper left quadrant corresponds to quadrant A in the text, the upper right quadrant to B, the lower left quadrant to C, and the lower right quadrant to D.






FIG. 35. Profile of HRI (top), MRI (middle), and ITS (bottom) flat field corresponding to Fig. 34 along row 250.

the new correction removes some of the high-frequency variations, like on pixels 23 and 53, for example. A more quantitative analysis indicates that the bottom part of the figure (corrected using the old flat field) has a standard variation of ~ 14 DN, while the upper part (corrected using the new flat field) has a standard variation of ~ 13 DN. This is only a slight improvement, but enough to affect the deconvolution process, which is very sensitive to high-frequency variations. Therefore, we have adopted the revised flat-field corrections for use in the calibration pipeline.

7. Noise

a. Random. The random noise in CCD images consists of the detector read noise, signal shot noise due to photon statistics, and analog-to-digital conversion (ADC) quantization error. Measurements of the random noise levels in the DI instruments were made as part of the photon transfer analysis discussed in Sec. V C 2. Figure 23 shows the relationship between the random noise level and the signal level illustrating that we are shot noise dominated over most of the detec-

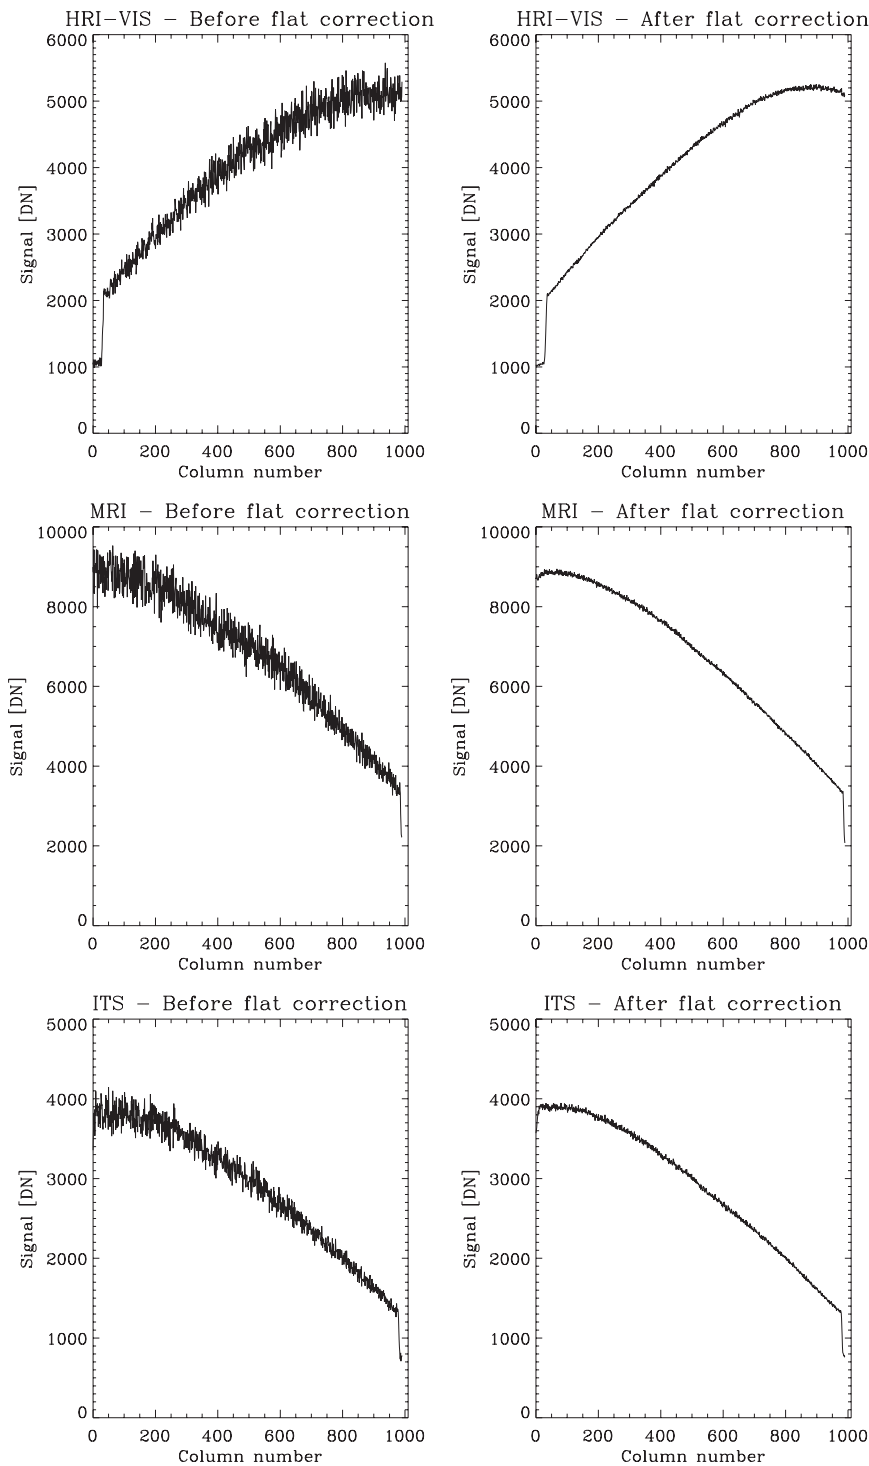


FIG. 36. Profile of a stimulator image before (left) and after (right) flat-field correction for HRI (top), MRI (middle), and ITS (bottom). The flat-field correction reduces the pixel-to-pixel variations from $\sim 3\%$ to $\sim 0.5\%$ rms.

tor's dynamic range. Table IX lists the read noise levels measured for each quad of each camera. Assuming a perfect ADC, quantization error would be at only the 0.3 DN level and would be insignificant. Actual ADC performance is discussed in Sec. V C 7 c.

b. *Coherent*. The DI CCDs and associated readout electronics were largely free of coherent noise, which we define as obvious and patterned electrical interference from external sources. Most images show a small (on the order of 1 DN) variation in bias level that varies every few rows, which results in some mild horizontal striping when images are stretched to enhance low-signal levels. The result is a reduc-

tion in effective SNR that must be accounted for in SNR calculations. Its effect is included in the readout noise determined above. The effect is most strongly seen in MRI narrow-band filtered images, for which signal levels are typically low.

Superimposed on top of many images is an apparent interference signal, sometimes of a few DN (< 10) amplitude, peak to peak. The signal is largely sinusoidal or sometimes half-sinusoidal (signal is only increased), and is shown in Fig. 38. In the time domain, the frequency of the sinusoid is ~ 8 kHz. This sinusoidal noise is not truly coherent in that it does not occur at the same time, and therefore the same

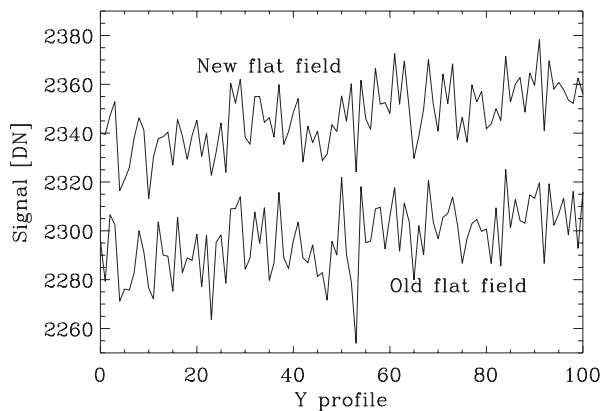


FIG. 37. Comparison of the profile of a coma image with the new flat field including the stimulator pixel-to-pixel relative-response variations and the TV4 flat field.

position, within the image. In fact, many images show no sign of this sinusoidal noise. When it does occur, it is common to all four quadrants and occurs symmetrically within the CCD, though each quadrant will have a different amplitude of the noise. In the few cases where the noise has been examined in some detail, the pattern is stable in frequency and may even exhibit a long-term pattern (over many readout rows) that could be used to create an improved subtraction technique. With the relatively low amplitude and unpredictable occurrence within an image, there were no specific calibration steps derived to subtract this noise signal. A routine that detects rows in which the noise occurs was developed, but is limited to images with primarily dark backgrounds and illuminated scenes covering fewer than a few 10 s of pixels.

A 100 ms HRI CCD dark exposure with the CCD at room temperature ($\sim 22^\circ\text{C}$) is shown in Fig. 39. There are two distinct features in the image that are common to all three CCDs. First is a series of diagonal stripes, which are due to metallization layers in the CCD. The bright bands are repeated every 8 pixels in an individual row and shift horizontally by one pixel from row to row. The stripes are continuous across left-to-right boundaries, but are discontinuous over the top-to-bottom boundary. The second is a large increase in the dark signal within a horizontal area near the center of the image. The first is observed over all CCD tem-




FIG. 39. CCD room-temperature dark current pattern. The center 256×256 pixels of a room-temperature HRI image are shown.

peratures, while the second seems to diminish when the CCD is operated below -30°C . These patterns become undetectable at normal HRI and MRI operating temperatures.

c. *ADC*. Analysis of the accuracy of the digital encoding of the CCD analog signals was performed using HRI images taken with the internal stimulator (STIM) turned on. The STIM results in images with a smoothly varying left-right signal gradient that spans roughly a factor of 2 in brightness. An ideal ADC would output data from such an image that has a smoothly varying histogram with approximately equal numbers of DN's within the neighborhood of each signal level in the image. Figure 40 shows a representative histogram from one of the prelaunch STIM images. The histogram indicates that the CCD ADC performance is less than ideal, with encoding bin sizes that vary from level to level by as much as a factor of 2. The 14 bit encoding is actually only reliable to 13 bits at certain DN levels, particularly those separated by an interval of 16 DN (and to a lesser extent, those separated by 8 DN). The ADC performance appears to be repeatable in a statistical sense in that certain DN levels




FIG. 38. An example of a recurring low-level noise sometimes seen in all three CCD imagers. A segment from one quadrant of a HRI CCD image is shown below, and two DN profiles across the image are plotted above. The duration of the noise corresponds to about two times the readout time for one row. Profile (a) is a row with the noise, and profile (b) is a row without the noise. Profile (b) is offset by +5 DN for clarity.




FIG. 40. Histograms of HRI CCD STIM images showing uneven ADC bit weighting: (a) full histogram, (b) expanded view.

are always under-represented while others are always over-represented. Therefore, an improved measure of the true analog signal for a given pixel could be obtained by mapping each DN level to the center of its actual ADC encoding bin (fractional DN). Such a correction would improve the image SNR, at least for lower signal levels (<100 DN) where encoding error is a significant noise source. The data exist to determine such corrections; however, we have not yet derived correction files for pipeline use. MRI and ITS ADC performance has not yet been analyzed; however, we expect their performance to be comparable to that of the HRI vis ADC.

d. *Compression.* Application of the data compression LUTs to convert from 14 bit encoding to 8 bit causes an increase in encoding error and therefore an increase in the noise level of any pixel's measurement. The LUTs for the vis cameras were designed with a square-root structure, based on the observation that the detector performance is shot-noise limited over most of the dynamic range of each detector. Because the bias varies from quadrant to quadrant for each detector and also with the temperatures of the detector and the associated electronics, the LUTs each start out with a linear section for minimal signal levels covering the range of bias levels expected during operations. The second section of each LUT is then a square-root function. As shown in Figs. 10–12, the various LUTs span different sections within the

TABLE XII. Comparison of encoding step size to signal shot noise for VIS data compressed using various LUTs.

LUT		Uncompressed DN range	Compressed DN step/shot noise
HRI/MRI	ITS		
3	0	0–2000	1.72
2	1	0–6000	3.23
1	2	0–full well	5.2
0	3	0–3000	3.3
		3000–full well	12.6

system dynamic range with the increases in encoding error being less when a smaller range is spanned and larger when the entire range up to full well is covered. One LUT has a dual-slope character with smaller encoding steps at the low end of the dynamic range and larger steps at the high end. Table XII summarizes the approximate range of uncompressed DN's above the bias spanned by each LUT and the ratio of the compressed DN encoding step to the shot-noise value that applies for each LUT.

e. *Cross-talk.* The DI cameras exhibit electronic cross-talk between the various quadrants of the CCDs. The cross-talk is an electronic effect generated within the readout electronics. Images with significant and structured signal show mirrored, low-intensity images in other quadrants. The magnitude of this effect is 0.15% or lower depending on the combination of quadrants used. This effect is negligible for faint sources (e.g., coma), with the mirrored signal being less than 1 DN for a primary signal lower than 650 DN, but becomes non-negligible for bright sources (e.g., nucleus) producing mirrored signals larger than 10 DN for a primary signal higher than 6500 DN. In order to quantify this effect, we designed a specific sequence that was executed during the in-flight science calibration phases. The sequence is an observation of Canopus (a bright star) in each quadrant, which we exposed normally and then saturated about 100 times above full scale (16 383 DN) in order to have a good cross-talk signal in the other quadrants. From these observations, we were able to derive the fraction of signal (called the cross-talk ratio) that goes from one quadrant to the other for every combination of quadrants. The results are given in Table XIII for HRI, MRI, and ITS. The cross-talk ratio varies from 0.03% to 0.15%. Using these ratios, a correction for cross-talk can be made by subtracting the correct percentage of the mirrored signal in each quadrant from every other quadrant.

Figure 41 illustrates the results before and after the cross-talk correction. Before the correction, we see Canopus (saturated) on the lower right quadrant and an electronic ghost in the other three quadrants. After the correction, the electronic ghosts disappear from these three quadrants. A more detailed analysis indicates that the ghosts do not completely disappear but remain at a level lower than 2 DN, a reduction of at least a factor of 3. We consider this result acceptable.

8. Scattered light

Measurements were made of the vis cameras' stray light responses, i.e., the response at a given pixel location due to a

TABLE XIII. Cross-talk ratios for HRI, MRI, and ITS. The ratio is the fraction of the primary signal that is produced in another quadrant.

Quadrant combination	HRIVIS crosstalk ratio(10^{-4})	MRI crosstalk ratio(10^{-4})	ITS crosstalk ratio (10^{-4})
B over A	4.0	5.0	6.0
C over A	3.5	3.0	5.0
D over A	3.5	3.0	5.0
A over B	3.5	6.0	6.5
C over B	7.0	8.0	14.0
D over B	4.0	5.0	3.0
A over C	3.0	3.5	5.5
B over C	8.0	9.0	15.0
D over C	5.0	3.0	8.0
A over D	3.3	4.0	6.0
B over D	4.0	4.0	3.0
C over D	6.0	3.5	10.0

source not imaged directly on the pixel (including sources outside the FOV). Such calibrations are impractical on the ground for a well-baffled telescope, so they were conducted in flight. It is useful in theory to consider the “stray light” due to sources in the FOV simply as extended characteristics of the PSF. However, some contributions to the in-field stray light, such as reflections off the filter substrate surfaces, depend on the position of the source within the FOV, whereas the PSF from the telescope has very little variation over the FOV. We did not attempt to decouple and characterize the far wings of the PSF from scattered light due to other spatially dependent causes.

Two calibrations were performed using the moon for MRI and Canopus for HRI. Reference images of both objects were taken along with long exposure images of adjacent fields with the object illuminating the inside of the baffle (not the CCD). Figures 42 and 43 show the stray light pattern for HRI. A quantitative measure of the HRI stray light response along a diffraction spike is plotted in Fig. 44. Figures 45 and 46 show the stray light for MRI. A quantitative measure of the MRI stray light response along a diffraction spike is plotted in Fig. 47.

The in-field stray light was measured using lunar images taken shortly after launch. An initial assessment of the stray light showed it to be so low that correction techniques were not necessary. More detailed analysis of the stray light proved this to be true for all filters except the IR (950 nm). The method used for measuring the stray light was to create a lunar mask image slightly larger than the actual lunar image and use it for measuring the total lunar signal inside the mask and for measuring the total signal falling outside the mask. The ratio is a measure of the integrated in-field scattered light. In Fig. 48 the results for the HRI are shown. In Fig. 49 the similar results from MRI are shown.

Note that MRI vis shows about six times more scattered light than HRI vis. This is expected because the focal ratio of HRI is $f/35$ compared to $f/17.5$ for MRI and because the obscuration of MRI (31%) is $2.3\times$ higher than that of HRI (13%). Errors in this measurement can result from the zero reference not being exactly zero and from the imperfection in

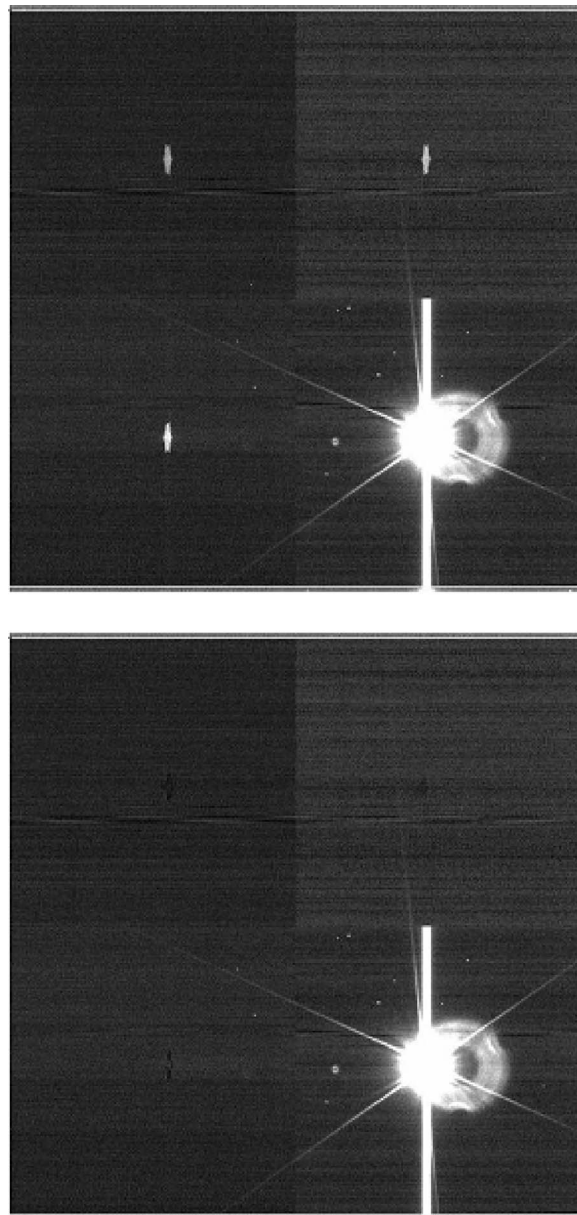


FIG. 41. Saturated image of Canopus before (top) cross-talk correction and after (bottom) cross-talk correction. The electronic ghosts nearly disappear.

the lunar mask. In Fig. 50, a scattered light column trace for MRI is shown, and a zero reference line is also shown. The zero reference error is negligible. The inherent PSF of the instrument contributes only a small amount to the light outside the lunar mask. As a test of possible measurement errors due to an error in the size of the lunar mask, the size of the mask was progressively increased, and the rate of decrease in the calculated scattered light observed. Changing the size of the mask by 5 pixels decreased the scattered light by about 10%.

All the filters except the IR filters show very similar performance leading to the conclusion that the enhanced scatter in the IR filters is due to either the filters themselves or to the CCD. Light longward of 900 nm will not be completely absorbed in the CCD and will reflect off the CCD gate structure. The magnitude of this effect is currently unknown for this CCD design.




FIG. 42. The stray light pattern for HRI. Canopus is shown scaled as a reference. There are five images surrounding the boresight image. Each image has an exposure time of 800 s. The Canopus image saturates with an exposure time of about 20 ms. The secondary mounting structure (the spider) is the major contributor to the observed stray light.

The in-field scattered light is very low in all filters except the IR—3.6% of the incident flux spread over the focal plane for MRI and 0.6% for HRI for the example lunar images that were analyzed. For an extended object such as the Moon it shows up most noticeably in the wings of the edge-spread function. Even for the IR filter, the amount of scattered light should normally be small enough to be ignored. A full characterization of the scattered light pattern in the images as a function of source solid angle and its location relative to the FOV is beyond the scope of this paper.

9. Charge bleeding/residual

The CCD for DI was chosen not to include blooming control circuitry (primarily to maintain a large full well, $\sim 400\,000$ electrons). Therefore images that overexpose a scene will exhibit effects of charge bleeding as shown in Fig. 41. Theoretically, even if charge fills the well in a pixel and begins to bleed down its column, the total amount of charge collected should be conserved. In ground tests of the CCDs with a 16 bit digitizing system, the CCDs showed that even when saturated, such that charge is bleeding up and down CCD columns, the total charge collected is still conserved, with the total amount of charge collected matching that expected to within around 10%. Due to the desire to maintain a reasonable single DN step size in the CCD system, the flight CCD system is set up to reach the maximum amplitude of the ADC before the CCD is fully saturated and starts to




FIG. 43. The same HRI data as in Fig. 42 converted to polar coordinates with 1° increments.

bleed. This situation is shown graphically in Fig. 51. In this case, images that show bleeding charge provide a lower limit of the actual charge produced in that frame. Also shown in the figure is the detail of an image of Canopus purposely overexposed by around $10\times$ to examine electrical cross-talk. A secondary effect of saturating a region of the CCD is that the CCD readout electronics undershoot when the CCD read-




FIG. 44. One of the bright columns from the HRI polar plot in Fig. 43 (at 94°) normalized with respect to the total Canopus signal showing the worst-case HRI scattered light response to a point source outside the FOV.




FIG. 45. The stray light for MRI. The moon is shown scaled as reference. There are five images surrounding the boresight image. Each image has an exposure time of 1000 s. The moon saturates with about 15 ms exposure time. The MRI telescope is similar to the HRI, so similar stray light patterns can be seen.

out changes from a saturated column to a dark column. This also adds to uncertainty in determining the total photometric signal from a saturated scene.

Ground tests of charge residual, i.e., remaining charge in saturated regions in subsequent images, show that even from several times oversaturated images, residual charge was below the limit of detectability.

10. Radiation noise

A cosmic ray (CR) passing through the vis or IR detector can cause generation of a cluster of signal charge that can lead to incorrect interpretation of an image if it is not correctly distinguished from scene-induced charge. The characteristics of CR charge clusters in images made during the flight of DI were studied for both typical and peak (during a flare) solar activity periods.

CR events are most easily detected in dark images, so these were used to derive statistics for CR events. Our studies were based on analysis of raw images. Only pixels with digital numbers $DN > 370$ for MRI and $DN > 390$ for HRI were considered candidates for being CR events (i.e., signals more than ~ 15 DN above the bias level). On the top panel of Fig. 52, the observed number of CR events per cm^2/s is plotted versus the number of pixels affected by the event for dark images taken during both typical solar activity and during a flare. On the bottom panel of the figure, the mean value of DN (including the bias level) within the CR signature is




FIG. 46. The same MRI data as in Figure 45 converted to polar coordinates with one degree increments.

plotted versus the number of pixels in the event. This figure is based on analysis of several MRI images with exposure time $9 \leq t \leq 30$ s, where t is the commanded exposure time. For $t \leq 0.2$ s the number of detected CRs can vary by a factor of several in different images with the same exposure time. We note that the CR integration time is not the same as the image integration time. It also includes the light blocker




FIG. 47. One of the bright columns from the MRI polar plot in Fig. 46 (at 190°) normalized with respect to the total lunar signal.




FIG. 48. HRI vis in-field scattered light shows the comparison between the clear filter and the IR filter. Total stray light is expressed as a fraction of the total signal within the masked area. Each image is 512×512 pixels.

close time and, on average, $\frac{1}{2}$ the readout time. For mode 1 images, the CR integration period is about 0.82 s longer, on average, than the total image exposure time.

For most HRI and MRI visual images made during low solar activity at $t > 4$ s, the number of CR events per s/cm² of CCD was about 2–4 (typically ~ 3), and generally there were no events consisting of more than $2t$ pixels. Most CR events in an image consist of not more than 4 pixels. The largest CR signatures have a linear form in contrast to the more circular form for star images. At high solar activity, the CR event rate can increase by a factor of 5 compared to that at low activity, and long CRs due to grazing paths of the cosmic ray through the detector membrane can exceed $8.5t$. At $t=30$ s the maximum number of pixels in one long streak exceeded 200, while no CR events consisted of more than 45 pixels at $t=30$ s for images outside the period of solar flares. The ratio of the number of CR events consisting of n pixels obtained at high solar activity to that at low solar activity was greater for greater n . For example, this ratio was greater at high solar activity than that in out-of-peak activity by a factor of about 1.5, 2, 3, 3.5, and 7 for rays consisting of 1, 2, 3, 4, and 5 pixels, respectively. This suggests that events caused by energetic particles from the Sun tend to produce larger signatures than do interstellar CRs.

Based on comparison of CR signatures on dark and sky images, we can make two main conclusions that can be used for recognition of CRs.

- (1) Even for out of peak of solar activity, most CR events are smaller than the camera PSF. At $t \geq 4$ s, almost all ($>80\%$) 1–4 pixel charge clusters in typical sky im-

ages are CRs (excluding images of dense clusters of stars).

- (2) Large CR events have a linear form in contrast to the more circular form of star images. We calculated the ratio $k_p = n_{\text{pix}} / (d_x^2 + d_y^2)$ for different clusters, where d_x and d_y are the maximum differences of coordinates x and y in a charge cluster each increased by 1 pixel, and n_{pix} is the total number of pixels in the cluster. At $n_{\text{pix}} > 30$ we found that $k_p < 0.17$ for all CRs (dark images) and $k_p > 0.17$ for all stars. However, when charge clusters consist of ≤ 10 pixels, it is difficult to distinguish between CRs and stars based on k_p .

The most reliable way to recognize CR events in astronomical images is to compare different images of the same region of the sky, but it is not always possible to do this. We analyzed the performance of four programs (IMGCLEAN, CR-FIND, DI_CRREJ, and RMCR) written by E. Deutsch, R. White, D. Lindler, and S. Ipatov, respectively,²³ that seek to recognize CRs in a single image. The first three programs run well in many cases (e.g., for typical calibrated sky images, for analysis of which they were created), but usually they do not work well with raw images. Some of the programs have problems with long (oblique entry) rays, and they delete pixels near the edge of a comet nucleus or from its coma. CR-FIND and DI_CRREJ only identify pixels corresponding to CRs, but do not replace these pixels. The RMCR program was written to work on both raw and calibrated images and to replace detected CRs with values of their neighboring pixels. Below we briefly discuss the performance of these programs, giving main attention to those images for which the above programs do not work well.

In the RMCR program, which was adapted expressly for DI, only those pixels for which raw DN or calibrated radiances are greater than some limit lim are considered to be possible CRs. This limit can be an input parameter (e.g., for raw images) or it can be calculated as $\text{lim} = \text{limit } 0 \times k \text{ lim}$ (e.g., $k \text{ lim} = 3$), where $\text{limit } 0$ is the median value of all pixels on an image. For calibrated images, one may not know lim in advance, so it is better to use the latter calculation of lim unless the area of interest is black sky. This program deletes all “long” CRs—clusters with $k_p < 0.17$. Charge clusters consisting of not more than $n \text{ limit}$ pixels are considered as CRs. Depending on a considered image and problem, the input parameter $n \text{ limit}$ can be chosen to take any value (e.g.,




FIG. 49. MRI vis in-field stray light. Each image is 300×300 pixels.

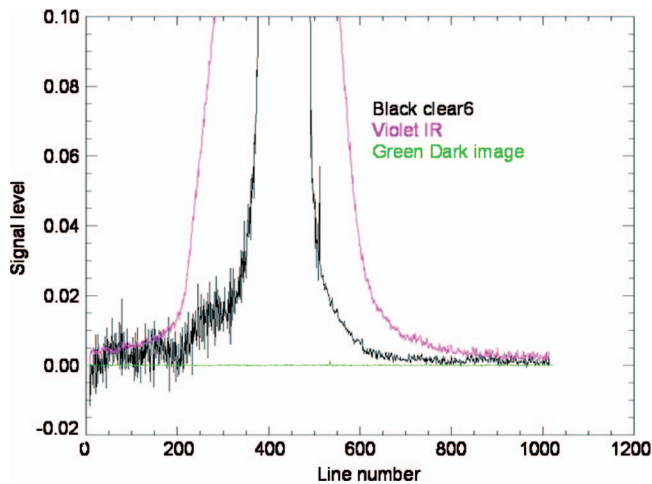


FIG. 50. (Color) MRI in-field scattered light response down column 650. Signal level is in arbitrary units with all filters normalized to the same signal amplitude.

10). In one variation of the program, the objects (exclusive of long objects) that are closer than a specified number of pixels to a defined rectangle that includes the comet and its coma are not considered as CRs. For small values of n limit, it may be useful to run RMCR for calibrated images two times—first with a greater lim, and then with a smaller lim. The program runs slowly when there are a lot of pixels in all objects (e.g., a comet occupies a considerable part of an image) because the program analyzes the entire image at once, rather than by small portions of the image at a time as do the other programs.

The effectiveness in recognizing cosmic rays with the different programs varied for different kinds of images. Programs IMGCLEAN, CRFIND, and DI_CRREJ often do not work normally with raw images. Sometimes they had errors during their processing and didn't run to completion; sometimes they deleted a lot of arbitrary pixels of background.

Figures 53–55 illustrate clusters recognized as CRs by the different programs on three images. The images presented on Figs. 53 and 54 consist of 256×256 pixels ($t=20$ s and $t=5$ s, respectively) and include Temple 1 and its coma. That on Fig. 55 consists of 181×181 pixels ($t=60$ s) and images a bright star with HRI. The expected number of CRs is about 15, 4, and 20, respectively. On all three images, CRFIND and DI_CRREJ defined too many pixels

as CRs that are not CRs. IMGCLEAN does not recognize the whole long CR on Fig. 53, but it deletes fewer pixels from the coma than RMCR (which does recognize the whole long CR). So IMGCLEAN seems to be the best technique to use for Fig. 53. Pixels erroneously classified as CRs by DI_CRREJ and IMGCLEAN at the edges of an image are actually due to bright pixels at the boundary between the imaging array and the overclocked pixels. RMCR excludes these pixels from its search for CRs. In the case of Fig. 54, the choice of the best program depends on the purpose for which the frame is being analyzed: DI_CRREJ did not delete pixels near the edge of the comet nucleus, but it considered about 80 small clusters (not well seen on the figure) as CRs instead of expected number of ~ 4 ; IMGCLEAN worked normally far from the nucleus, but it deleted many pixels around the nucleus; it is possible to find parameters for suitable operation of RMCR, but they are not those that are usually used; CRFIND did not run to normal completion, so it gave the worst results (but for some other images it can give the best results). In the case of Fig. 54, IMGCLEAN replaced the brightness of some pixels recognized as CRs near the nucleus by a wrong brightness, and after such replacement the image incorrectly looked as if some material was ejected from the comet. In such cases it may be better not to attempt CR removal or to try to find input parameters for which RMCR gives suitable results. On Fig. 55, IMGCLEAN recognized as CRs pixels that belong to the star, but did not recognize as CRs the objects near the star, which must be deleted for proper correction. So for Fig. 55, RMCR gave better results.

For a calibrated dark image at a peak of solar activity, IMGCLEAN removed most of the CRs, but a few tens of them were left on an image. Even several long rays were left. IMGCLEAN did not recognize well long CRs on some other images. For a calibrated image of conglomerations of stars, IMGCLEAN left most of stars, but sometimes it deleted many small stars. It classifies as CRs a lot of pixels near the comet which are not CRs. It is a fast program and is easy to use.

For the calibrated images considered (dark images and images of clusters of stars) with maximum radiance of $\sim 0.0001 \text{ W m}^{-2} \mu\text{m}^{-1} \text{ sr}^{-1}$, DI_CRREJ and CRFIND did not work normally if we used the same default parameters for which these programs worked normally with calibrated images with maximum radiance of $\sim 1 \text{ W m}^{-2} \mu\text{m}^{-1} \text{ sr}^{-1}$. For the small radiance case, they deleted a lot of pixels of back-

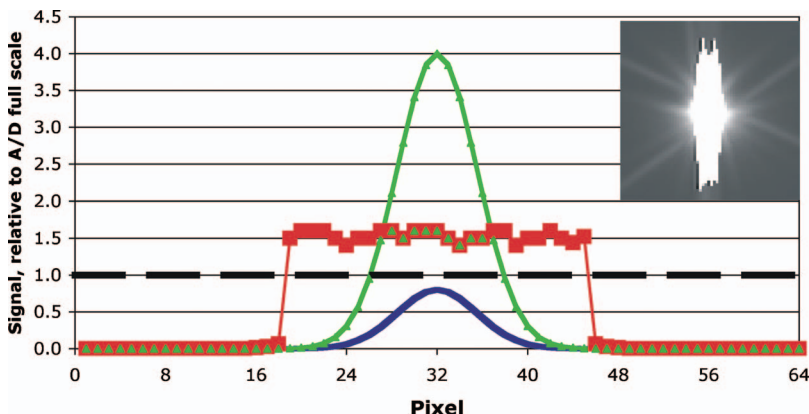


FIG. 51. (Color) CCD saturation and charge conservation. The plot shows an idealized cross section of signal of an unsaturated star (blue line without symbols) and a saturated star (green line with triangles) in the case where the physical limitation of the CCD full well is above the range of the ADC. The signal level is shown relative to the maximum value digitized by the ADC. The resulting signal level in the CCD is also shown (red line with squares). Any signal above the dashed line is digitized as 16 383 counts. The figure in the upper right of the plot is a HRI image of Canopus saturated about $10\times$, showing the charge bleeding up and down columns.

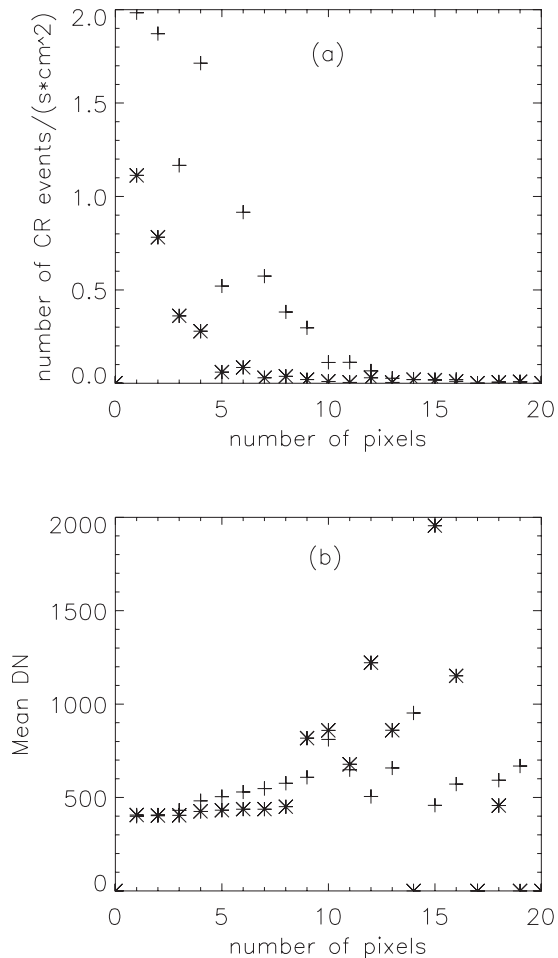


FIG. 52. The number of cosmic ray events per s/cm^2 (top) and mean values of DN (bottom) vs the number of pixels affected by an event at a peak of solar activity (crosses) and for out-of peak activity (stars) on dark MRI images with $9 \leq t \leq 30$ s.

ground. We have not found parameter settings that work well at small radiances for these programs.

IMGCLEAN is a more reliable program if one needs to remove CRs automatically from a large number of images, but depending on a specific image and a specific problem, other programs can work better (e.g., sometimes CRFIND is the best when there is a large image of a comet). All of the CR detection programs discussed here will be made available in the DI PDS archive.

CR detection has proved more difficult in ITS images than in either MRI or HRI images. The problem may be because the values of the background DN in an ITS raw dark image vary more than those for MRI and HRI images (by a factor of 2 or more) due to the higher operating temperature and increased dark current of the ITS CCD. For raw ITS images only RMCR works. The number of charge clusters classified as CRs on dark ITS images per second per quadrant was about the same as that for MRI and HRI if we consider only pixels with $DN > 50$ above the median value for a quadrant. The difference between the median values for different quadrants of ITS raw dark images can exceed 40. If we consider pixels with $DN > 20$ above the median value of DN for a quadrant, then the number of clusters recognized by RMCR as CRs is greater by a factor of 5–10 than that for MRI

and HRI. The fraction of clusters detected as CRs and common for a pair of images (presumably, therefore, not true CRs) is about 5% of the clusters detected as CRs on one image when a threshold of 50 DN above the background is used. None of the programs considered worked well with calibrated dark ITS images using their default parameter settings. The number of charge clusters consisting of ≤ 4 pixels deleted by IMGCLEAN and RMCR was greater by a factor of several than even the expected number of CRs at the peak of solar activity, so most of the deleted clusters were not real CRs. CRFIND and DI_CCREJ designated even more pixels as CR. These excess CR detections appear to be due to the inadequacy of simply using a quadrant-mean dark current subtraction technique for the ITS. The excess CR detections tend to occur at the same pixel locations in all frames. This problem can be corrected by implementing a pixel-by-pixel dark current subtraction technique for ITS, and this work is in progress.

VI. IR SPECTROMETER CALIBRATIONS

This section presents the calibration results for the IR spectrometer instrument. The geometric calibration of the IR in terms of its fl, which gives the pixel scale projected onto a target, the precise orientation of the spectrometer slit and its spectral dispersion axis, and the relative alignment of its boresight to those of the various cameras are discussed. The IR's spatial resolution capability is presented in terms of its PSF. The radiometric response of the IR is quantified including its response linearity, zero-exposure background level as a function of time and instrument temperature, relationship between output DN and signal electrons, detector full-well capacity, flat-field response variations, and absolute sensitivity factors for converting DN to scene radiance or surface reflectance. We specify the mapping between wavelength and pixel location in the detector array and discuss the spectral resolution of the spectrometer. In addition, we discuss quantitatively the various data degraders such as bad (uncalibratable) pixels, random and coherent noise in the system, accuracy of ADC, data compression effects, scattered light response, bright image blooming and residual signal, and radiation noise effects. A detailed layout of the SIM bench is shown for reference in Fig. 56.

A. Geometric calibrations

1. Focal length

The fl of the overall IR system was measured by two methods—on-the-ground target tests and in-flight scan tests. The ground tests confirmed that the scale of IR images of a test target, taken nearly simultaneously with visible images, was 5.00 ± 0.01 times that of the HRI vis for binned IR pixels (2.5 times for unbinned physical pixels). Combined with the measured instantaneous FOV (IFOV) of the HRI vis images in flight, this results in an effective fl of 2.10 ± 0.04 m for the IR system, and an IFOV of $10 \mu rad$ for binned pixels.

The in-flight method of determining the IR fl is an analysis of images taken during scans of the star cluster 47 Tuc from the May and June instrument calibrations. This method is less reliable than the ground test because it is subject to




FIG. 53. Calibrated MRI image 6 002 420 (256×256 pixels). From left to right: initial image, pixels recognized as CRs by DL_{CRREJ} , $CRFIND$, $IMGCLEAN$, and $RMCR$ (at $k \text{ lim}=3$ and $n \text{ limit}=10$).




FIG. 54. Central part of calibrated MRI image 9 000 907 (1024×1024 pixels). From left to right: initial image, pixels recognized as CRs by DL_{CRREJ} , $CRFIND$, $IMGCLEAN$, and $RMCR$ (at $k \text{ lim}=15$ and $n \text{ limit}=10$).




FIG. 55. Calibrated HRI image with a large star (at maximum radiance $\sim 0.5 \text{ W m}^{-2} \mu\text{m}^{-1} \text{ sr}^{-1}$). From left to right: initial image, pixels recognized as CRs by DL_{CRREJ} , $CRFIND$, $IMGCLEAN$, and $RMCR$ [at $\text{lim}=k \text{ lim} \times \text{limit } 0=0.0005 \text{ W m}^{-2} \mu\text{m}^{-1} \text{ sr}^{-1}$ and $n \text{ limit}=30$].

errors in the scan rate and direction. The spacecraft was commanded to scan the IR slit across the star cluster in what was determined to be the across-slit direction. Any component of the scan in the along slit direction could add a bias (for a drift in motion) or noise (for random jitter) to the measurement. Therefore the results from this test are used only as a sanity check to make sure the ground tests did not suffer from a significant error in the analysis. The results from the two tests match to well within the uncertainty of the in-flight measurement.

2. Relative boresight alignments

We used two methods, ground calibration and in flight, similar to the focal-length methods, to determine the relative alignment of the HRI vis boresight and the center of the IR detector. Unlike the IR fl measurements, the in-flight measurement is more representative of the relative boresights than the ground-based measurement and is used as the encounter configuration. The ground calibration method is described to document the small shift seen between the two calibrations.

The ground calibration method involved near-simultaneous IR and HRI vis imaging of a resolution target (called the Air Force Test Target) at the focus of the ground test collimator. Using cross-correlation techniques between an along-slit IR profile and an along-slit profile created from the vis image, the along-slit position was determined. The relative cross-slit position was found by varying the expected

IR slit position in the visible image. The modeled slit position that produced the greatest cross-correlation coefficient was deemed the cross-slit position. The positions determined in ground calibrations at 1 G are shown in Fig. 57.

In flight, the relative boresights were determined from analysis of a scan of the star Vega during the May and June calibrations. Both IR and vis images contain a time stamp with a well-known offset to the center time of the exposure for that image (as described in Sec. III). The central position of the star in the HRI vis images versus time is well modeled using a linear fit, and shows that the scan rate during the maneuver was constant. Thus, the position of the star in the vis image could be interpolated during the scan. When the time that the star signal peaked in the IR image set was found, the horizontal (across slit) and vertical (along slit) positions are determined from the interpolation, and with simple offset corrections the relative HRI vis boresight to the center of the IR detector is determined, as shown in Fig. 58. The alignment relative to the MRI is derived from the MRI-to-HRI boresight relationship. The in-flight alignment of the IR spectrometer boresight relative to that of the HRI vis shifted by -12 HRI vis lines and -5 HRI vis samples relative to the ground-based calibration.

3. Slit alignment

To investigate the in-flight alignment of the IR slit with respect to the MRI field of view, we observed the globular

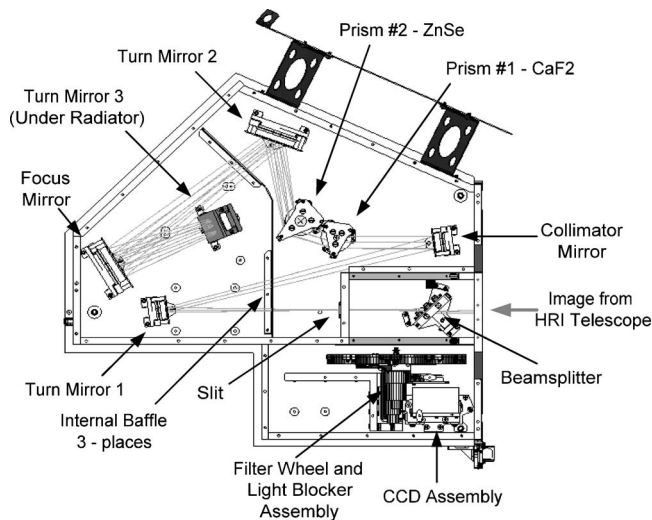


FIG. 56. Detailed layout of the SIM (Ref. 3, used with permission).

cluster 47 Tuc with a 100-image IR scan. The stars seen in the IR FOV during this scan are shown in Fig. 59.

A subimage containing the IR FOV has been extracted from the full-frame MRI image of the same star field taken just before the scan and is shown in Fig. 60.

From the IR image of 47 Tuc, 31 bright stars were selected, and a Gaussian plus a constant background were independently fit in the along-slit direction and in the scan direction. The same stars were identified in the MRI image, and a Gaussian plus a constant background were also independently fit in the two pixel directions. The differences in




FIG. 57. Relative boresight alignment of the HRI vis and IR spectrometer as measured prelaunch (boresight offsets in figure are not drawn to scale; only the central ~100 pixel subarea of HRI and the central 10% of IR spectrometer slit are shown).




FIG. 58. Relative locations of the HRI-IR slit with respect to the MRI and HRI fields of view measured in flight (boresight offsets in the figure are not drawn to scale; only the central ~250 pixel subarea of HRI, the central ~120 pixel subarea of MRI, and the central 25% of IR spectrometer slit are shown).

location between the IR and MRI images in the across-slit direction are graphed in Fig. 61 as a function of along-slit position. Note that there is no noticeable trend over the range of over 250 IR physical pixels, indicating that the IR slit is aligned with the MRI pixel direction to within a fraction of a MRI pixel over this distance.

In addition to the alignment of the IR spectrometer spatial slit to the FOV of its own and the HRI and MRI vis detectors, there is also the question of the alignment of the spectral dispersion axis with the IR detector array. By plotting the central line number of a dispersed star spectrum against column number, the dispersion axis was found to be nearly aligned with the fast-readout axis of the detector array, but with a slight variation with column number or wavelength, as illustrated in Figs. 62 and 63, amounting to almost exactly one physical pixel change in position from one edge of the detector to the other. The cause of this slight (~1 mrad) misalignment is presumably due to a small rotation of the prism axes with respect to the detector raster orientation. This slight misalignment was also seen on the ground.

B. Spatial resolution

The analysis of the in-flight spatial resolution of the IR spectrometer used the across-slit scan of the calibration star Beta Hyi taken during the May 2005 science calibration. After pipeline processing, a data dark was derived by performing a pixel-by-pixel resistant mean over the 50 spectral images (2.5-σ threshold), and this data dark was subtracted from all 50 spectral images. Analyzing this data set, the maximum signal from Beta Hyi was found in scan image 29 of 50, and the peak signal in the center column of that image



FIG. 59. IR FOV during scan of 47 Tuc. The FOV is 100 scans wide and 512 physical pixels tall. To have matching spatial scale in both directions, the image has been stretched by a factor of 2 in the scan (horizontal) direction.

was found in row 417. Further analysis of dark-subtracted image 29 consisted of fitting a Gaussian plus a quadratic polynomial to each column of the image, from column 6–1019 (the light-sensitive columns). The sigma of the Gaussian was converted to equivalent FWHM through application of the scaling factor of 2.354. The result of this analysis is plotted in Fig. 64. The spatial resolution of the IR spectrometer as a function of wavelength is plotted in Fig. 65.

Figure 65 indicates that the IR spectrometer spatial resolution is diffraction limited at wavelengths longward of about $3\ \mu\text{m}$. Shortward of that, other optical performance




FIG. 60. An extracted subimage of the MRI FOV containing the FOV of the IR spectrometer during its scan of 47 Tuc.

characteristics limit the spatial resolution. The most likely cause of less-than-diffraction-limited performance at short wavelengths is the defocus in the HRI telescope. If the IR spectrometer defocus were the same as that of the HRI vis, the PSF spot would have been 3.5 physical pixels. It appears that the IR spectrometer optics shift its focus location so it is somewhat better than that of the HRI vis camera.

Figure 66 plots the IR spectrometer PSF along-slit profile at three wavelengths where the performance is diffraction limited. The first Airy ring is noticeable at 5%–10% of the peak signal level, and the width of the Airy ring is seen to increase linearly with wavelength as expected.

C. Radiometric calibration

1. IR linearity

The response linearity of the IR FPA was analyzed using the measured dark signal during the cruise to Tempel 1. The dark signal is a combination of the detector dark current and the background thermal IR radiation emitted by components on the SIM bench. IR spectrometer linearity was first measured during ground tests using a blackbody flat-field source; these tests showed that the response nonlinearity is independent of wavelength to less than 1% uncertainty, which allows it to be measured reliably using only SIM bench radiation. Results from the TV4 ground tests indicated that one quadrant of the IR array had a signal-chain diode failure that required repair before flight, and new ground calibration data could not be collected after the repair. Therefore, only the flight data were used to determine the response linearity for the IR array.

Data were collected during three different calibration events during the cruise to Tempel 1, in April, May, and June, as well as in July soon after impact. The temperature for the FPA for each calibration event remained under 85 K,




FIG. 61. Difference in across-slit IR and MRI star positions as a function of along-slit position.

and the standard deviation of the temperature during a single event was less than 0.025 K (Table XIV). The optical bench temperatures remained constant during each calibration event as well; however, the temperature shift between calibrations is significant. Each calibration event involved collecting data at ten different integration times in each of four imaging modes. The integration times were not the same for all four modes in order to collect data at as many different signal levels as possible. The first two calibration events, in April and May, collected three frames for each integration time, and the June and July events collected up to five frames for each integration time.

With only ten data points for each integration mode, pixel-by-pixel linearity equations would be suspect and unreliable due to insufficient SNRs. Therefore, the linearity equations were calculated as averages for each quadrant by combining the data from all pixels in a quadrant in each of the four imaging modes. The different imaging modes analyzed included one unbinned and three binned modes (the

pixels in the subframe binned modes were binned with a one row offset relative to the full-frame binned pixels). Reference rows and bad pixels were ignored, and the resistant mean over the entire quadrant was calculated by removing outliers beyond 2.5 standard deviations. The data from each calibration event for Quadrant A can be seen in Fig. 67. The optical bench temperature decreased for each successive calibration event, so the earliest/warmest (April) event has the highest slope, while the latest/coolest (July) event has the smallest slope. The response is close to linear with integration time, but measurable departures from linearity are observed. The response rates are seen to be slightly different across the different modes within a given calibration event.

Although the data were collected at different integration times for each operational mode, a slice at a single integration time shows that the subframe modes consistently demonstrate higher response than either full-frame mode. The two full-frame modes, binned and unbinned, show the same




FIG. 62. The inflight variation of the centers of Gaussian fits to the image of the calibration star Beta Hyi as a function of IR spectrometer pixel (column) number. The total variation from one edge of the IR detector to the other is almost exactly one physical pixel.




FIG. 63. The inflight variation of the centers of Gaussian fits to the image of the calibration star Beta Hyi as a function of IR spectrometer wavelength. The total center variation in row numbers from minimum to maximum observed wavelength is almost exactly one physical pixel, and the variation is very close to linear with respect to wavelength.




FIG. 64. The IR spectrometer spatial resolution (FWHM of the Gaussian fit) as a function of pixel (column) number for the 10 May 2005 observation of calibration star Beta Hyl. Note that the spatial resolution is approximately constant at a value of 1.6–1.7 pixels FWHM for pixels 6 through about 700 and then rises approximately linearly to nearly 3.0 pixels for pixel numbers around 1000 due to optics diffraction.

response during all four collection events. The two subframe modes use only 50% or 25% of the total pixels available, and the out-of-frame pixels are not reset in these imaging modes. It is likely that in the subframe modes the out-of-frame pixels have saturated, and the built-up charge from these pixels induces an increased response rate in the subframe pixels that is proportional to the number of out-of-frame pixels. In order to remove this effect, data in the range of 3000–6000 DN were used to calculate signal-level ratios of the full-frame to subframe modes at the same integration time. Using this DN range allowed all four imaging modes to be used by avoiding integration times when a significant fraction of pixels begin to saturate, and this range exhibited consistent linear behavior in all four calibration events. This range includes four to six data points for each dataset. The shortest integration times could not be included since data could not

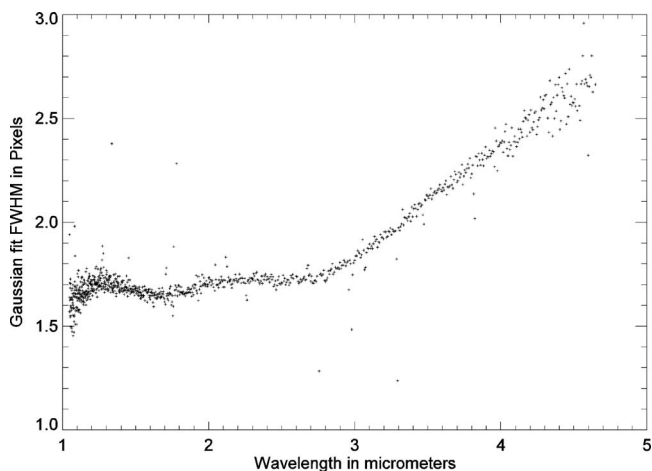


FIG. 65. The IR spectrometer spatial resolution (FWHM of the Gaussian fit in physical pixels) as a function of wavelength for the 10 May 2005 observation of calibration star Beta Hyl. Note that the spatial resolution is dominated by the defocused PSF of the HRI telescope for wavelengths shorter than about $2.8 \mu\text{m}$ and by the diffraction limit of the telescope for wavelengths longer than about $2.8 \mu\text{m}$. One physical pixel is about $8.5 \mu\text{rad}$.




FIG. 66. Along-slit traces of a star image in the IR spectrometer for 3.5, 4.0, and $4.5 \mu\text{m}$ (top to bottom). Also plotted are Gaussian fits; each data set is normalized by the maximum value of the Gaussian fit and then respectively displaced by 0.5, 0.25, or 0.0 intensity units upwards. The data points are joined by lines, and the Gaussian fit values are dotted.

be collected below approximately 3 s for the unbinned mode. Also, the subframes taken at the shortest integration times were not affected by out-of-frame saturation since these frames were collected immediately after resetting the array in the prior mode, and the out-of-frame pixels were not yet saturated. The ratios were calculated to be approximately 0.95 for BINSF1 and 0.92 for BINSF2. The difference between these ratios is consistent with BINSF2 having 50% more out-of-frame pixels than BINSF1. Once these factors have been applied to the data, the linearity curves from each month agree well for all the modes. Figure 68 plots the signal generation rate versus the mean signal level for each calibration event.

In order to properly calculate the linearity equations, certain data points should be omitted from the fits. It is obvious in the May data that a significant number of pixels have saturated for the longest integration time. However, there is not enough granularity in the integration times to determine exactly when an individual pixel saturates. With a limited number of data sets, including the saturated data would significantly impact any curve fitting to be performed. The two shortest integration times were also removed from the subframe datasets since they do not seem to be affected by the out-of-frame saturated pixels in the same way as subsequent

TABLE XIV. IR spectrometer focal plane and optical bench temperatures from the cruise calibration events.

Month	FPA temperature	FPA temperature δ	Bench temperature	Bench temperature δ
April	84.97	0.02	139.52	0.06
May	84.28	0.01	137.80	0.01
June	84.07	0.01	137.05	0.02
July	84.37	0.02	136.93	0.07

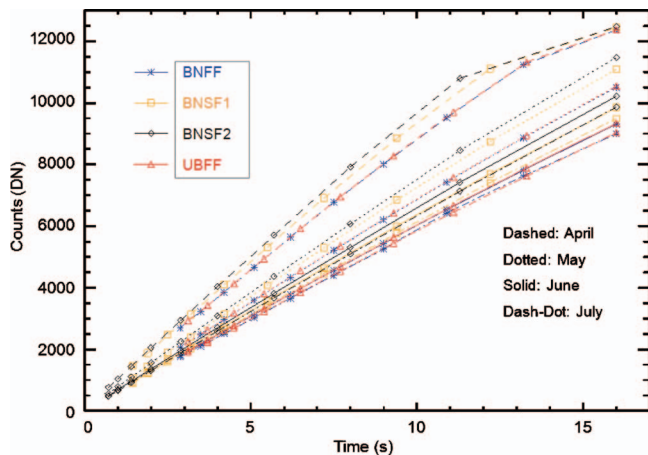


FIG. 67. (Color) The number of IR spectrometer DN measured in the left quadrant A using a resistant mean vs integration time for each imaging mode is shown. Imaging modes are identified by color, and the month of the event is defined by linestyle.

frames taken at longer integration times. Finally, the first frame for each integration time was not included because of the dark-level instability discussed in Sec. VI C 3 c.

Although the temperature of the optical bench was relatively constant over the entire cruise, the IR detector is sensitive to small changes from month to month. In order to check for changes in response linearity throughout the mission, each data set was normalized to 1 DN/ms at 5000 DN. Figure 69 shows that the response linearity remained constant over all four calibration events to within the accuracy of the measurements. The degree of nonlinearity is $<5\%$ over the dynamic range of the detector.

A series of polynomials was used to find the best fit, and the appropriate degree of polynomial was selected using a chi-square goodness-of-fit statistic. We selected a third-order polynomial since it provided a better fit than a quadratic equation by almost 20%, while a fourth-order polynomial provided less than 0.2% improvement. Higher-order polynomials saw much smaller improvements. Since each quadrant of the detector uses separate signal-chain electronics, a different nonlinearity equation was derived for each quadrant. The data and curves can be seen in Fig. 69.

The curves shown in Fig. 69 were derived from the first three calibration events. This preliminary solution was necessary in order to begin processing encounter data as soon as they arrived after impact. Fortunately the flyby spacecraft survived passing the comet, and the July post-impact calibration was accomplished. Table XV includes the linearization equations calculated with and without the July calibration data. Although there are some changes in the values of the polynomials, both quadrants show a difference of less than 0.15% between the two equations below 11 000 DN. This results in differences of less than 10 DN in the linearization correction, which is below the noise and error associated with the data (less than 0.15%). The largest deviations in the equations occur above 11 000 DN, since the available datasets do not have enough resolution to solidify the equations beyond this full-well level (see discussion of full well in Sec. VI C 2). The July difference was judged small

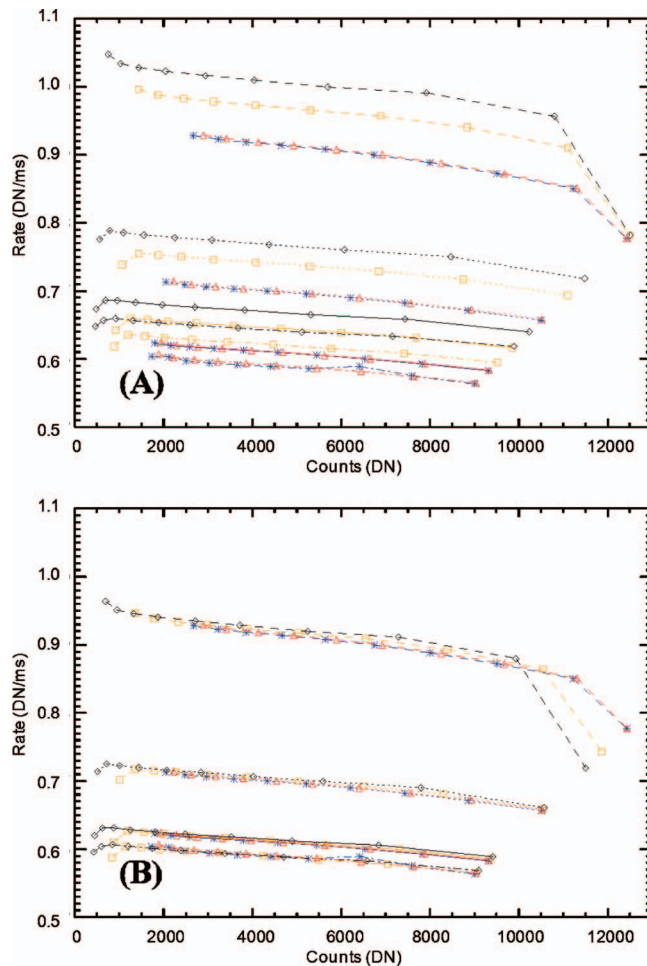


FIG. 68. (Color) The IR spectrometer response rate is calculated using the data shown in Fig. 67 for quadrant A. Plot A shows the subframe imaging modes with a higher response rate. Plot B was created by applying the correction ratios of 0.95 and 0.92 for BINSF1 and BINSF2, respectively. As in Fig. 67, imaging modes are identified by color, and the month of the event is defined by linestyle. The nonlinearity in the IR response becomes more apparent in these curves.

enough that the linearization equations derived from the first three calibration events were adopted for all the data collected by the IR spectrometer.

2. Gain/full well

The gain and full well of the IR detector can be characterized using the photon transfer technique in a way similar to that used in CCD calibrations. For the IR spectrometer, the best data set for doing this analysis is a series of dark frames acquired in flight with a range of exposure times from the minimum possible up to times that cause detector saturation. The emission from the SIM bench provides a relatively uniform illumination over the detector and allows data to be obtained at the lowest possible signal levels so the low end of the detector's dynamic range can be characterized. Using additional external flat-field radiance sources, calibration data were also acquired prelaunch. Photon transfer measurements were made several times both before launch (TV 1, 2, and 4) and in flight (April, May, June, and July 2005 calibrations).

As discussed in Sec. VI C 3 c, the IR dark frame level

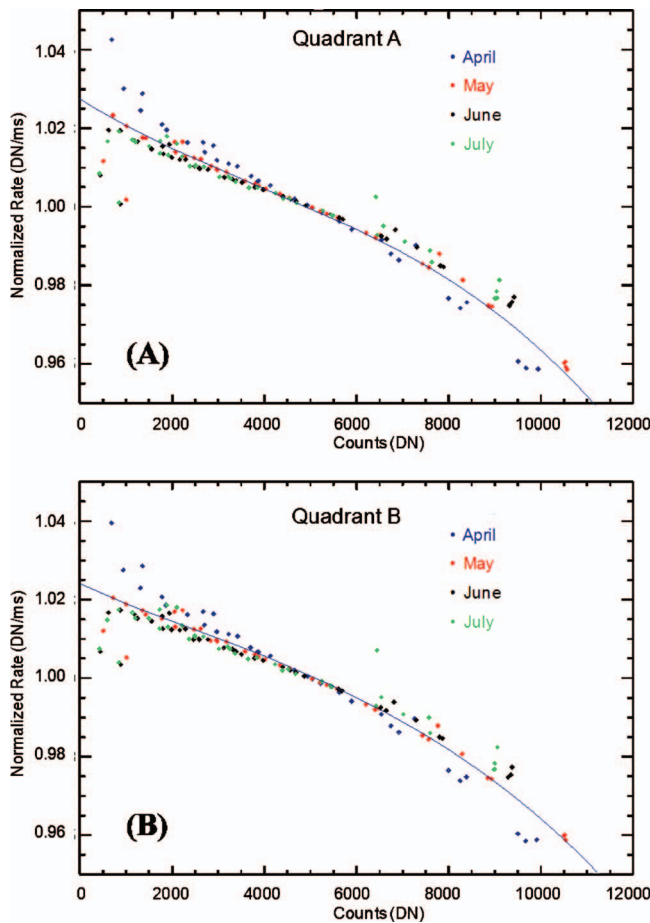


FIG. 69. (Color) The IR spectrometer nonlinearity equations are derived after applying the subframe scaling ratios, and all data are normalized to 1 DN/ms at 5000 DN. Plot A shows the data from quadrant A, and plot B shows the data from quadrant B. Data from each month are represented by a different color.

observed in a series of successive frames decreases significantly between the first and second frame and more gradually for each frame thereafter. Whatever is causing this variation also seems to cause an increase in the random noise level in the first few frames. Therefore, the photon transfer analyses were done using frames that were no less than third in a series of successive readouts to minimize bias from this effect. Median signal values and standard deviations in the frame-differenced signal were determined for 33 20 × 20 pixel standard areas (10 × 10 in the binned modes) dis-

TABLE XV. IR spectrometer linearization equations calculated for both quadrants with and without the July calibration data. X =raw DN value; Y =linearized DN value. Differences in the polynomials are less than 0.15% below 11 000 DN.

Quadrant	July included	Equation
A	No	$Y=1.02748-7.29603 \times 10^{-6} X$ $+6.18492 \times 10^{-10} X^2-5.16479 \times 10^{-14} X^3$
A	Yes	$Y=1.02846-8.50283 \times 10^{-6} X$ $+9.16036 \times 10^{-10} X^2-6.97525 \times 10^{-14} X^3$
B	No	$Y=1.02406-5.22930 \times 10^{-6} X$ $+2.54496 \times 10^{-10} X^2-3.28636 \times 10^{-14} X^3$
B	Yes	$Y=1.02537-6.74768 \times 10^{-6} X$ $+6.25803 \times 10^{-10} X^2-5.55182 \times 10^{-14} X^3$

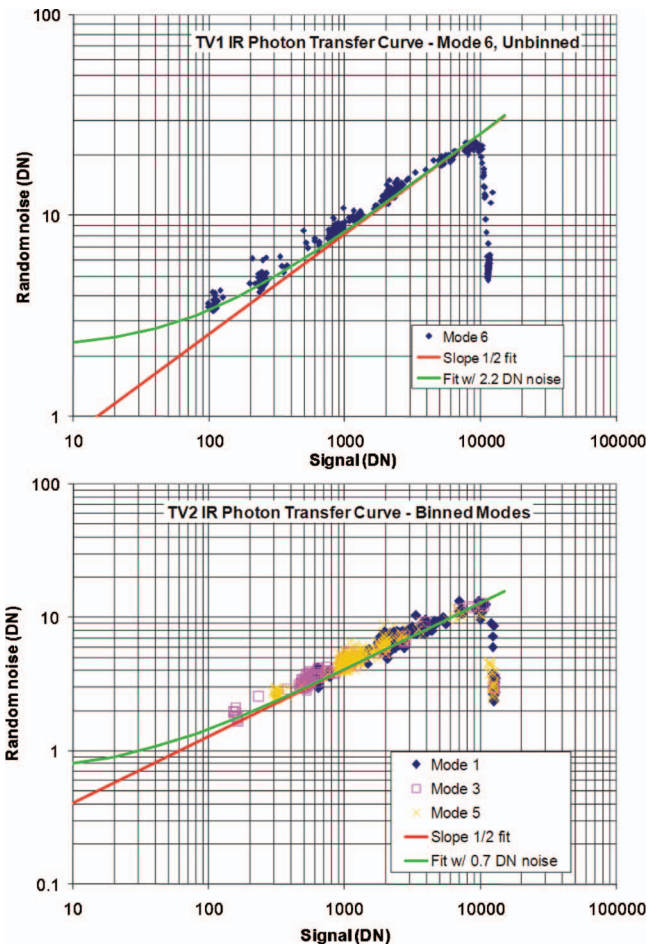


FIG. 70. (Color) IR spectrometer photon transfer data in the unbinned (top) and binned (bottom) modes.

tributed across each image. The areas were selected to avoid any obvious hot pixels. They were positioned to avoid crossing the boundaries of the antisaturation filter rows.

No obvious differences in the photon transfer data points are seen as a function of frame, subarea, illumination source, SIM bench temperature, or mode (within the binned and unbinned categories) over the few days of each calibration data set and between calibrations spanning from TV1 through the end of the mission (except for TV4 as discussed below). Figure 70 shows the photon transfer data for unbinned and binned modes from TV1 and TV2.

The TV4 photon transfer data revealed an obvious change in the Quad B results compared to its performance in TV1 and TV2. This change was subsequently traced to a failed diode in the Quad B signal chain, which was replaced prior to launch. Unfortunately no recalibration of the IR spectrometer was possible after this repair until the in-flight calibrations.

Photon transfer analysis of the in-flight IR spectrometer showed no measurable change in system gain, read noise, or full well from those measured prelaunch. For the unbinned mode, the measured gain was 16 electrons/DN. The measured random noise level begins to fall below the shot-noise fit above about 8000 DN (128 000 electrons) indicating that this is the conservative full-well limit of the detector, since some pixels are beginning to saturate beyond this point caus-

ing the random noise level to be less than it should be. The majority of pixels saturate at signals above about 11 000 DN (176 000 electrons); this can be taken as a less conservative definition of full well. At the coldest SIM bench temperature achieved in our calibrations (~ 123 K) with the shortest available unbinned integration time (1.4 s), we were not able to obtain signal levels below ~ 100 DN, so our measurement of the read noise floor is somewhat uncertain. We estimate a read noise level in the unbinned modes of ~ 2 DN (32 electrons). Knowing the read noise level for the IR spectrometer is not very important, however, because in flight we were never able to achieve signal levels less than 400 DN; we were always operating in the shot-noise dominated regime.

For the binned modes, the random noise measured in output signals from the 2×2 pixel averaged bins is a factor of 2 less than would have been measured in unbinned data at the same output DN signal level. Thus, the measured binned-mode gain is four times that in the unbinned modes [because at a given signal level, gain is proportional to (random noise) $^{-2}$]. The measured gain in binned modes is 64 electrons/DN, consistent with the value measured for the unbinned modes. The number of electrons here is the total number in the 4 pixels that were averaged to yield the output DN. Applying the unbinned gain factors to binned DN data will give the average number of electrons in each corresponding unbinned pixel. A conservative full well limit of ~ 8000 DN and a less conservative value of 12 000 DN are again measured. The read noise floor is ~ 1 DN (equivalent to ~ 32 electrons of noise when initially encoding each pixel value prior to binning, which is again consistent with the value determined from the unbinned photon transfer data). This level of read noise is insignificant at typical signal levels.

3. Zero-exposure level

a. *IR background: Introduction.* The analysis of the HRI-IR spectrometer background level focused in two areas: the level as a function of temperature and the level as a function of time. The vast majority of the signal comes from the background “glow” of the instrument (e.g., the surrounding structure and the optics), while a minor component comes from the thermal dark current of the FPA itself. In both cases, the signal is temperature dependent, with higher temperatures producing a higher signal. The background is also time dependent, because signal is constantly integrating on the FPA, and is only cleared during a reset or readout of the image. Unfortunately, the reset/readout process does not entirely clear the accumulated signal, so the background level in one frame is weakly dependent on the time that has elapsed between the start of the readout of the previous frame and the start of the readout of the current frame, as well as on the signal level that was present in the previous frame. Other time dependencies are also present and will be discussed more fully below.

The majority of the background analysis was done using data from the ground-based TV tests, because they were designed to cover a wide range of different temperatures and exposure times. Additional work was done using the in-flight

data as well, but even though these tests covered a range of temperatures, they were severely limited in coverage of temporal phenomena. While working to understand the data, a number of problems were discovered relating to the behavior of the background signal. Although much work was done in an attempt to resolve these issues, the available calibration data were inadequate to fully characterize some of them. (Most of the problems were discovered during the analysis of the ground-based TV test data, and the in-flight tests were too limited in quantity and scope to provide enough information to resolve them.) In the discussion below, these issues are addressed, along with the impact that they have on the data. For some of the circumstances, particular solutions were devised to minimize the problems, and in these cases, the reduction procedures are described.

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 across the array, the signal increases proportionally (up to the saturation point), and a subframe 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 antisaturation 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 changes across the entire FOV. (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). A resistant mean algorithm was applied to the columns to remove any pixels that differ by more than 3σ , and then the average of the remaining pixels was used to provide a single value that would represent the background level for the frame.

Typically, the HRI-IR observations consist of sequences of frames, all with the same observing parameters. Because the temperature does not change significantly during a sequence, the background level from one frame per sequence was sufficient for use in the analysis of temperature dependence. However, because the reset command does not fully clear the integrated signal, the first frame of a sequence has a measurably higher background level than those of subsequent frames (this is one of the problems discovered in the testing and will be discussed more fully below), so a later frame is more representative of the background level for the sequence as a whole. In order to remain consistent, the eighth frame in each sequence was typically used in the temperature analyses of ground-based data discussed here.

b. *IR background: Temperature dependence.* Most of the work involving temperature dependence was done using unbinned, full-frame mode (mode 4) data, because these include the entire array and comprise the most extensive set of measurements available in the ground-based tests. Although fewer measurements are available for the other modes, they appear to follow the same dependencies. Ultimately, a combination of ground-based and in-flight data was used to derive the final temperature results, and data from modes 1–4 were

combined to compute a single set of coefficients in the equations for temperature dependence.

The background signal was known to be dominated by the glow of the optics and surrounding structure in the instrument, with a secondary contribution from the dark current of the FPA. To determine the temperature dependence of the instrument glow, the averaged value from temperature sensors on the two prisms in the light path was used to represent the temperature of the SIM bench. Another sensor recorded the temperature of the FPA.

To determine the background signal's dependence on these temperatures, data from the second and fourth TV tests (TV2 and TV4) were used. The primary results come from the TV2 measurements, where background levels were obtained with nominal prism temperatures of 130, 136, and 142 K, and the corresponding FPA temperatures were 83, 85, and 88 K. In TV4, all of the data were obtained at only a single operating temperature. The prism temperature was ~ 139 K, and although this was in the same range as in TV2, the FPA at 107 K was significantly warmer than in TV2 and in flight. As will be discussed below, this fortuitously turned out to be useful in extracting the background contribution from the FPA dark current. All relevant data obtained when the instrument was at its operating temperature were incorporated in the analysis.

Because the wavelengths observed by the FPA are at the short-wavelength end of the blackbody function for the SIM bench operating temperature, the Rayleigh–Jeans approximation can be used to represent the background from both the instrument glow and the dark current. Thus, the temperature dependence of the background is expected to follow the Arrhenius functions for the two components:

$$B_{\text{tot}} = c_0 + a_0 \exp(a_1/T_{\text{prism}}) + b_0 \exp(b_1/T_{\text{FPA}}).$$

This relation was assumed throughout the analysis and there was no indication that the dependencies did not behave in this manner.

From the analysis of the TV2 results, it was found that, near the nominal operating temperatures, the contribution from the instrument glow completely dominates the background signal. By themselves, the TV2 data can be well fit using only the Arrhenius function relating to the prism temperature. This indicated that the FPA dark current provided only a relatively small contribution (<100 DN for nominal operating temperatures), which could not be constrained by the TV2 data.

Fortunately, TV4 provided good measurements of the dark current signal. At 139 K, the prism temperature was within the range covered during TV2, so the contribution from the instrument glow could be constrained. In addition, the FPA temperature was high enough that the dark current contribution was comparable to that from the instrument glow. So, using both the TV2 data and the TV4 data, it was possible to solve for all of the parameters to produce a fit to the observations. For mode 4 at its minimum integration time, the data were fit by the function

$$B_{\text{tot}} = -14.7 + 1.203 \times 10^{12} \exp(-2764.8/T_{\text{prism}}) + 1.89 \times 10^{11} \exp(-2002.0/T_{\text{FPA}}).$$

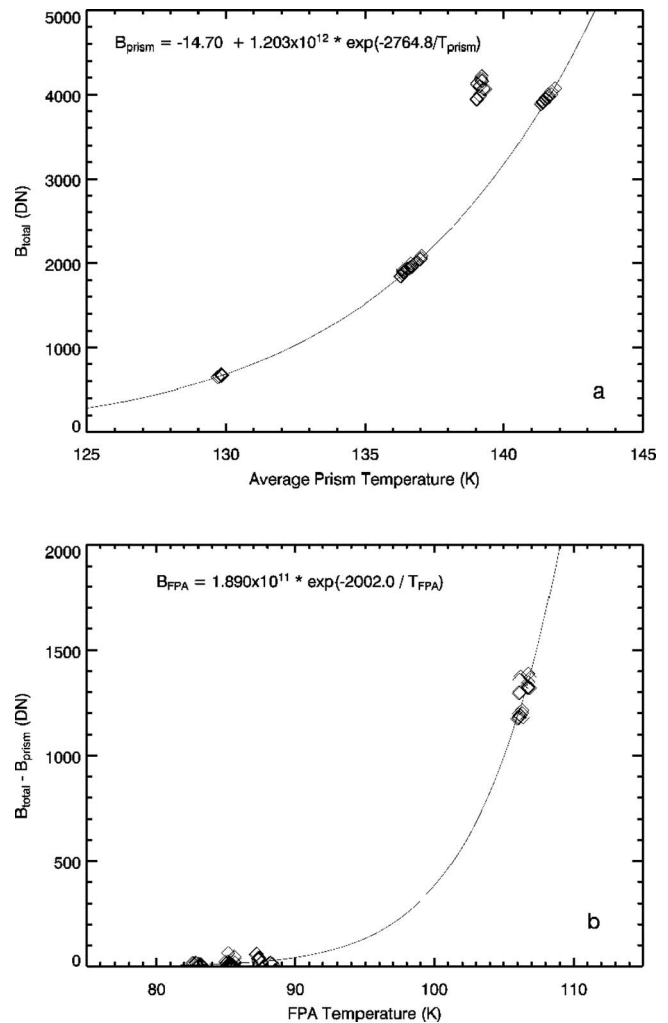


FIG. 71. IR spectrometer background level as a function of prism temperature for mode 4 ground-based data at zero integration delay. (a) Background as a function of the prism temperature. Symbols indicate the measured values of the total background signal, while the line indicates the computed values for the best fitting contribution from the instrument glow. (b) Residual background as a function of the FPA temperature (Ref. 3, used with permission). The symbols indicate the residual signal after the computed contribution from the instrument glow has been removed. The line shows the best fit for the dark current contribution.

Figure 71(a) shows the measured background level as a function of prism temperature, with the solid line denoting the component computed for the instrument glow. The excess signal above this curve is that from the FPA dark current. The small excesses at 130, 136, and 142 K reflect the small contribution of the FPA near its nominal operating temperature, while at the 139 K measurement, the excess is much higher due to the warmer detector in the TV4 measurements. Figure 71(b) shows the excess background component after the computed instrument-glow contribution was removed, with a curve to show the computed contribution from a best fit to the Arrhenius relation for the FPA. It is clear that when the FPA is near its nominal temperature around 84 K, the dark current signal is almost negligible compared to the instrument glow.

Additional measurements were obtained during the cruise phase of the mission, as the prism temperature dropped from 146 to 137 K. These data were used to derive




FIG. 72. IR spectrometer background rate as a function of the prism temperature for the in-flight data. Symbols indicate the measured values of the total signal, while the line indicates the computed values for the best fitting modeled contribution from the instrument glow.

the final prism temperature coefficients. The FPA temperature remained nearly constant during these observations, so an in-flight solution for the FPA temperature dependence was not possible. Thus the dependence found in the ground-based tests was adopted for the FPA contribution, with the understanding that, at the given temperatures, its contribution was nearly negligible, and so any error should be entirely negligible.

In order to incorporate the measurements from different modes, the images were first corrected for nonlinear response and then scaled by their exposure time, to convert them to linearized DN per millisecond. In this format, the background rate is consistent from one mode to another, allowing measurements in different modes and with different integration times to be combined into a data set from which a single set of coefficients can be derived. When the in-flight data from all modes are used, the best fit to the data produces the function

$$B/t = 0.0583 + 3.006 \times 10^{10} \exp(-3384.4/T_{\text{prism}}) + 6.563 \times 10^7 \exp(-2002.0/T_{\text{FPA}}),$$

which represents the final parameters for scaling the background level to account for the effects of temperature. The prism temperature dependence for the linearized in-flight data is shown in Fig. 72, and the residuals, after the computed contributions for both the instrument glow and FPA dark current have been removed, are shown in Fig. 73.

The plot of the residuals in Fig. 73 reveals another notable item in the background analysis. Although the amount of scatter within each of the temperature groups is fairly constant at around 3% (3σ), there are small systematic shifts between the groups (points at 137 K are systematically low, those at 138 K are systematically high, etc.). This suggests that the FPA dark current and the instrument glow are not the only contributors to the background level. Because the groups were obtained at different times and under different spacecraft orientations, it is possible, if not likely, that some other component of the instrument is changing, and causing




FIG. 73. (Color) Fractional residuals in the background level of inflight IR spectrometer data after the computed FPA dark current and instrument-glow contributions are removed. Residuals for different modes are shown in different colors.

the slight offsets. Fourteen different temperature sensors are located on components of the HRI instrument to record information about the environment, and an analysis was done to determine if any of them might be associated with the unknown contributor. Unfortunately, none showed any consistent correlation to the background residuals, suggesting that any additional background component comes from a source that did not have a sensor, and remains unidentified.

An examination was done of the amount of hysteresis in the temperature dependence (i.e., background level dependence on the past history of temperature). The variations in TV2 were less than 10 DN, suggesting that any hysteresis has an effect at a level of less than 1%, even for rapid temperature variations. The in-flight temperatures remained very stable (in the 2 weeks before encounter, the main temperature change was a gradual drop of about 2 K; around encounter it remained nearly constant), so any hysteresis in the flight measurements is expected to be negligible.

c. IR background: Temporal dependence. The IR background exhibits two types of temporal dependence. The first is the behavior of the background level as a function of the location where the frame falls in a sequence of images (read-out slot), and the second is the behavior of the background as a function of the time since the chip was previously read out (the inter-sequence gap). Some amount of interconnection exists between the two. A sequence is defined here as a series of images, all with the same observing parameters, that are obtained in rapid succession using a single command with a fixed, mode-dependent timing pattern between the images. For some of the following analyses, it was necessary to normalize the observations to remove the effects of temperature changes that occurred during the testing sessions by computing the expected background level for the given temperature and dividing by this value.

Figure 74 shows the background level in each frame of a sequence of 32 images from TV2, plotted as a function of time since the first frame was read out. It is clear that the first frame in the sequence has a background level that is significantly higher (in this case, about 4%) than the second frame,




FIG. 74. Background level in a TV2 sequence of 32 IR spectrometer images. The level in the first frame of the sequence is much higher than those in the following images, though subsequent frames continue to drop with time throughout the sequence.

and subsequent frames continue to fall throughout the sequence, even after 32 images. The rate of change is more extreme in the earlier images than in the later ones, where it settles into a very shallow, linear decrease. This figure illustrates the first type of temporal dependence, where the background level in each frame is dependent on the readout slot in the sequence. It should also be noted that, although this “overshoot” is always present in the first frame of a sequence, the amount of overshoot is not constant, ranging from 2% to 7% higher than in the rest of the images in the sequence. Furthermore, in the example shown here, the subsequent falloff is well behaved, dropping continuously with time. There are other cases in which the falloff oscillates throughout the subsequent frames, and others in which a gradual increase has been seen. These will be discussed later.

Qualitatively, the overshoot of the first frame can be explained by the basic operation of the detector. Because there is no cold shutter, the signal from both the dark current and the instrument glow continuously integrate, increasing the background level until it is reset or it reaches the saturation point. To eliminate this buildup prior to an exposure, a “reset” command clears the chip by reading it out. Unfortunately, not all of the charge that has accumulated is removed when the reset command is executed, so the background level comes from the signal that has built up since the reset combined with the leftover signal that was not cleared by the reset command. Subsequent readouts occur rapidly enough that the residual signal does not have time to build up to a significant level. Additional analysis of the first frame problem in TV2 data revealed that the overshoot level was also dependent on the amount of time that had passed since the previous readout of the chip (the intersequence gap) and on the mode that was used in the previous observation (i.e., changing modes has an effect on the overshoot level).

A series of tests was done during TV4 to further investigate these issues and attempt to quantify them. However, the data that would be needed to explore all of the different background issues would have required many sequences to





FIG. 75. Normalized IR spectrometer background level as a function of the intersequence gap for the first image in a sequence in mode 1. (a) The full range of observations showing a constant overshoot level (dashed line) for longer gaps. (b) Blowup of the first 33 s, showing the linear rise with intersequence gap time. The diamonds are measurements obtained for mode change from 2 to 1, and the squares are for no mode change (1-1). Linear fits to the rising portion in each of the two cases, indicated by the dashed lines, have the same slope but are offset in time.

be run for a variety of intersequence gaps, modes, and mode changes. Unfortunately, there was only a limited time available for all of the calibration measurements needed from TV4, and only a subset of the background tests could be obtained. Unlike the TV2 observations, which primarily focused on the full-frame unbinned mode (mode 4), the TV4 tests focused on the full-frame, binned mode (mode 1) with a secondary emphasis on the binned, subframe modes (modes 2 and 3) because these were the modes that would primarily be used during the encounter sequences.

Using the TV4 observations, it was confirmed that the overshoot level was strongly dependent on the intersequence gap. For small gaps (<15 s for mode 1) the overshoot level showed a linear increase with the gap time, while for longer times, it flattened out at a constant level (see Fig. 75). This behavior is the result of the constant accumulation of background charge and then the partial clearing when the reset command is executed. The linear sloped segment represents the residual for times when the signal is increasing after a




FIG. 76. Falloff in the IR spectrometer background level after the first frame in mode 1. The dotted lines and diamonds show the falloff of the background when the first-frame overshoot reached the plateau, while the solid lines and squares depict the sequences in which the first frame did *not* reach the plateau level.

recent readout. After about 15 s, however, the accumulating background signal saturates and no more charge is collected. When a reset is sent after this saturation, the uncleared residuals have a constant value, producing the plateau in the overshoot value that is observed in Fig. 75 for any time longer than 15 s.

Figure 75 also shows that the overshoot level is dependent on mode changes that occur before the start of the sequence. Measurements in which the mode changed from 2 to 1 before the sequence show a higher overshoot for a given intersequence gap time than those in which the mode did not change. However, the rising slopes and the plateau levels are the same for both cases, indicating that the background level accumulates at a constant rate and saturates at the same point, regardless of the previous readout mode. The offset in the rising portions may be due to timing differences between the two modes (e.g., mode 2 is a subframe of mode 1, and takes less time to read out, introducing an offset in the time when a pixel starts integrating), though there are not enough data to evaluate this conjecture. The data for modes 2 and 3, though more limited, seem to show the same general behavior as for mode 1, with a linear increase up to a flat plateau. The rise times for each mode are different, and a limited number of mode change observations means that the offsets for each case cannot be accurately determined or compared to those from mode 1. The plateau levels are also different for each mode (~ 1.05 for mode 1, ~ 1.08 for mode 2, and ~ 1.10 for mode 3), though within each mode they are the same, regardless of the mode of the previous readout.

While the typical behavior of the background level in a mode 1 sequence is exemplified in Fig. 74, there are cases that do not follow this pattern. Figure 76 shows a plot of multiple sequences, with the vertical axis denoting the difference in the background between two consecutive frames. In general, these sequences exhibit the sharp drop from the first to the second frame, and then a gradual monotonic drop through the rest of the images. (Note, however, that even these typical cases show exceptions where the dotted lines

jump upward.) Sequences in which the background never reached the plateau show a distinctly different behavior, as shown by the solid lines. The sharp initial drop is still present, but rather than a monotonic decrease with subsequent frames, the following frames tend to oscillate around zero. This indicates a fundamental difference in the behavior of the background level for frames that have reached the saturation point and those that have not. To complicate matters even further, the subframe modes (2 and 3) do not even exhibit the relatively well-behaved monotonic falloff when the first frame plateau is reached, and there is no consistent pattern with time for the behavior in these modes. This difference in the subframe modes is likely caused by the portions of the chip that are not reset or read out. Because these regions are never cleared, they remain saturated throughout the sequence, and it is believed that electronic interference from the unread regions is affecting the background levels in the pixels that are read out.

Because the first-frame overshoot effect is not predictable enough to model accurately, it was decided that corrections for it would not be automatically applied in the calibration pipeline. Possible manual adjustments in the background level of an initial frame are discussed in Sec. VI C 3 d.

d. *IR background: Removal procedures.* The ultimate goal in the background analysis was to understand how to remove it from individual images. As was stated earlier, as long as the dark current contribution remains negligibly small, the background increases proportionally across the field of view, so a master background frame can be created and scaled up or down to match the level in the image of interest. Tests showed that under normal circumstances, subtracting this scaled background from an image worked well for removing the background, if the amount of scaling could be accurately determined. To this end, a master background frame was created for each mode, using 32 background frames from the in-flight calibration data. In each case, the 32 frames were averaged together using a resistant mean routine to remove points that varied by more than 3σ . The master frames were normalized to remove the dependencies on SIM bench temperature and exposure time.

Initially, the master frame was simply scaled automatically using the relevant exposure time and temperatures to create a background subtraction frame for each observation. Because of the uncertainties in the time and temperature dependencies, this solution was known to be good to only a few percent. Early results using this technique to reduce the data at encounter indicated that it produced poor results in the observations of interest. Even in observations that produced a high signal, the resulting spectra are very sensitive to the background subtraction, and in images with low SNR, an accurate background removal is critical for interpretation of the data. It became clear very early that background uncertainties of a few percent were much too large to be useful for the data analysis that was being done (e.g., a 2% error in a 2000 DN background level would have a severe effect on “good” signals of 50–100 DN, and would completely overwhelm low SNR measurements that register only 5–10 DN). Figure 77 shows an example of the sensitivity to background removal errors, with the comet’s continuum showing dra-

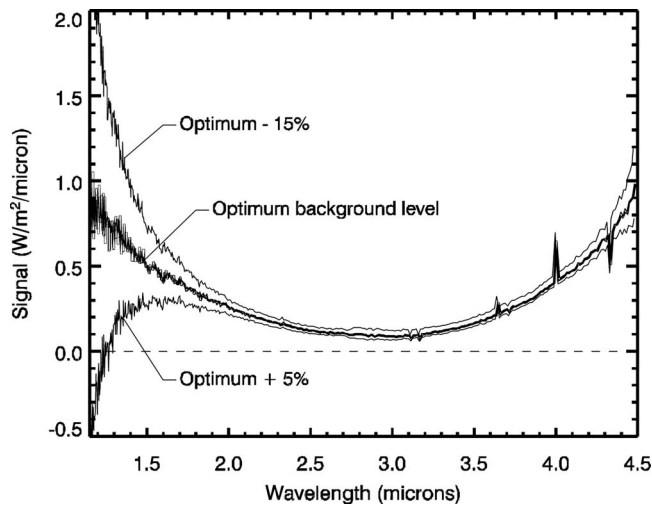


FIG. 77. Example comet spectrum showing the effects of improper IR spectrometer background removal. The heavy line shows the result for the “optimum” background level found as described in the text. The upper thin line shows the result for a 15% decrease in the background level, and the lower thin line shows the result for a 5% increase. Around $3\ \mu\text{m}$ the effect is $\sim 25\%$, while at the shortest wavelengths, it overwhelms the signal.

matic differences as a result of 5%–15% changes in the level of background removal. Because of the accuracy that was needed, it was determined that generating a representative background frame strictly from recorded temperatures and integration times would not be possible. Instead, a new strategy was needed that incorporated information from the images themselves to help constrain the background level. This need introduced a new level of involvement, where manual input is required to produce the best result.

Examination of the cometary data showed that at wavelengths near $1\ \mu\text{m}$ the scene-induced signal is usually weak, even when a strong signal is present at other wavelengths. This factor, combined with the fact that most of the cometary sequences used spatial scans that have at least one frame that is situated far from the nucleus, means that each sequence contains a region that is essentially just background, with no comet signal. Since the entire frame scales proportionally, this region can be used to scale the master background frame to represent the whole image. Furthermore, since all images in a sequence are 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 scaling of the one frame represents that for all the frames in the sequence. Using this information, a routine was developed to scan through all of the frames in a sequence, extracting either a default sampling of pixels or a sampling that can be designated manually, looking for the frame with the lowest count level and outputting information that can be used to determine whether or not it is actually at the background level. If a region is found to be purely background signal, then the scaling factor between it and the corresponding region in the master background frame is used to compute an appropriate background frame for the entire sequence.

It may be possible to apply a second-order correction for the monotonic falloff (or sometimes the increase) as a function of time since the beginning of each sequence. Many of

the sequences have clean background regions at both the start and end of the sequence, and these frames can be used to determine the changing background rates throughout the sequence. This refinement provides an even more accurate measure of the background level.

Although this strategy works for the majority of the cometary images, there are still a few problem cases. Not all of the sequences contain a “clean” background region. Around the time of the encounter with Tempel 1, the coma was bright enough that it produced a significant signal at all wavelengths throughout the scan, leaving no region that can be used for scaling. Fortunately, around the time of encounter, the temperatures remained constant, suggesting that the background would remain unchanged as well. For mode 1, which was used most often, the observations from before and after encounter support this idea, because sequences where the temperature is constant and clean background regions are available have scaling factors that remain nearly constant. Thus, the frames throughout this period can all use the same background scaling factors. There are fewer data available to fully assess the behavior for modes 2 and 3, but they do not appear to maintain a constant scaling factor. This may again be a result of not reading out the entire detector. For these modes, other information must be used to adjust the background level that is removed. This information can include the results obtained from mode 1, where the background levels are better determined.

Finally, there are two major issues for which there is no simple background removal technique. The first of these is the problem of the overshoot in the first frame. Because the magnitude of the overshoot is dependent on many factors and is difficult to model, there is no systematic way to compute what the level should be, so the images themselves must be used. If there is a clean background region in the frame, this can be used to scale the master frame as was done for the other images. If the frame has coma signal throughout, then other information must be utilized to determine the dark level. This is straightforward for the first frame of a sequence, where subsequent images can provide the necessary information. More problematic, however, are single frame images that are not part of a sequence. These still have the first-frame overshoot, but no subsequent images that can be used to indicate what the final result should look like. In these cases, neighboring scans should be used to constrain the background as well as possible, though it is likely that the uncertainties on these frames will be higher than for others. The residual error in the background level determination for such frames is estimated to be in the 1%–7% range depending on how similar the neighboring scans are to the one being calibrated; $\sim 2\%$ would be typical.

The second issue that presents major challenges to the background removal involves some of the sequences of the nucleus, during which a high signal (potentially near the saturation point) was observed. The temporal dependence of the next sequence will then have multiple components: pixels that were near saturation in the previous sequence, which are more likely to exhibit a higher overshoot and monotonic falloff, and the coma pixels with low signal, that will have a lower overshoot and an oscillating falloff. The net result of

this difference is to create structure in the background that is not present when the master background frame is scaled to match the image. Although the analyses have detected evidence that this scenario may be present, the extent of the problem is not currently known, and no procedures have been developed to deal with it. It is likely that a manual, user-intensive analysis, using a combination of the previous frame and the master background image, will need to be done on a case-to-case basis to produce a correction for this effect. The residual uncertainty in the background level determination in these cases is estimated to range from 0% to 7% and might vary across the frame depending on how much saturation occurred in the previous frame.

When the HRI-IR data are archived with the PDS,²⁴ different volumes will include updated procedures that reflect the state of the reductions up to that point. The first volume, frozen on 24 December 2005, incorporates the pipeline processing for most of the data, with background scaling coming strictly from the FPA and prism temperatures. The exception to this is the sequences around impact (9 000 024–9 010 003), where the manual scaling has been applied to improve the background removal. In future volumes, all sequences will include the manually determined scaling factors, as well as a correction for the monotonic falloff in sequences where it can be determined.

4. Wavelength map

The registration of spectra on the detector of the IR spectrometer was extensively characterized through ground tests and then verified in flight.

The first spectral registration data set was obtained during TV1, during which spectra of argon and krypton spectral lamps were taken at spectrometer temperatures of 135, 128, and 141 K bracketing the nominal operating temperature. This data set confirmed the dispersion of the spectra as a function of temperature as calculated from the Code V design model and indicated a departure from the desired alignment by +33 physical detector pixels in the spectral direction and –5 physical detector pixels in the spatial direction. In all cases, departures from desired alignment are measured positive in the direction of increasing row or column number of the image and negative in the direction of decreasing row or column number in the image. To improve the registration of the desired spectra on the detector, slight adjustments were made in the focal plane position after TV1.

The second spectral registration data set was obtained during TV2, during which spectra of argon and krypton spectral lamps were again taken, as well as spectra of methane (CH₄) and carbon monoxide (CO) absorption cells. These tests confirmed that the spectral registration adjustments that had been carried out between the two TV tests had been successful. The measurements of spectral registration showed that the registration was only +1 physical detector pixel from the desired location in the spectral direction and +2 physical pixels from the desired location in the spatial direction.

The third spectral registration data set was obtained during TV4, during which spectra of argon and krypton spectral lamps were once again taken, as well as spectra of methane




FIG. 78. Wavelength vs unbinned column number for the IR spectrometer for a bench temperature of 140 K.

and carbon monoxide absorption cells. These tests showed that the spectral registration had been slightly altered by one or more of thermal cycles, vibration tests, and/or the removal and reinstallation of the IR detector assembly to add stray light baffles near the detector. As a result, the spectral registration was measured to be +5 physical pixels from the desired location in the spectral direction, and +3 physical pixels from the desired location in the spatial direction. However, this amount of nonoptimal registration was within acceptable limits, so no effort was made to improve the alignment of the system before launch.

The best data set obtained prelaunch was the one obtained during TV4. That data set had the best SNR in the spectral lines and included many spectral lines that covered the entire long-slit spatial range of the detector. As a result, it was possible to construct a detailed temperature-dependent model of the spectrometer spectral dispersion and geometry and thus to provide detailed wavelength and dispersion maps for every pixel of the detector for any given temperature. That model was checked against the TV2 data set to confirm that it also reproduced those results when the appropriate spectral and spatial displacements were included. Figure 78 shows the wavelength imaged on each column for the central row of the IR detector array for a bench temperature of 140 K. A given wavelength shifts toward the right (higher column number) with decreasing bench temperature by an amount of about 0.7 pixels/K. The variation in column location of a given wavelength for rows away from the center, the spectral “smile,” is defined by the following equation:

$$\Delta\text{column} = 102.298 - 102.298\sqrt{1.0 + [(\text{row} - 261.688)/740.984]^2},$$

where the row numbers start with 0 at the bottom of the frame and increase upward, and a negative change in column means a shift toward the left (lower column number). The model is integrated into the calibration pipeline to provide pixel-by-pixel spectral registration and spectral dispersion information for every spectral image.

After launch, the spectral registration was confirmed to be unchanged from that determined during TV4 by observations of planetary nebula NGC 7027. The observed lines

TABLE XVI. Width of spectral lines observed in the IR spectrometer.

Pixel location	Wavelength (μm)	FWHM (pixels)	Source
33	1.067	2.11	Argon
264	1.296	1.93	Argon
469	1.694	1.81	Argon
625	2.313	1.87	Argon
782	3.32	2.65 ^a	CH ₄
987	4.66	4.8 ^a	CO

^aNote that the sharpest features in the CH₄ and CO bands are sufficiently wide that these are upper limits of the resolution of the spectrometer at these longer wavelengths and not measurements of the true resolution of the spectrometer.

from the planetary nebula were compared to published IR line lists for this nebula [ISO SWS line list for NGC 7027^{25,26}] and found to be indistinguishable from the predicted positions, that is, nominally 0.52 ± 0.5 physical detector pixels away from predicted, using the same spectral registration offsets from optimal that had been determined during TV4.

5. Spectral resolution

The IR spectrometer of the DI mission was designed to measure and characterize spectra of extended sources such as the coma of a comet, or the resolved surface of the nucleus of a comet, or the debris from an impact event. The calibration of the spectral resolution of the instrument on the ground was therefore also done with extended sources, including a low-pressure argon lamp and CH₄ and CO absorption cells, which were arranged to fill as uniformly as possible as much of the slit as possible, both in width and in length. Even the relatively low dispersion of 350–700 generally resolves the molecular bands and surface absorptions we wish to characterize, so simply showing that the effective spectrometer resolution for resolved sources matches the design resolution is adequate for our purposes. As shown in the IR spectrometer spatial resolution section VI B, even when we look at point sources such as calibration stars in flight, we see a soft focus image that is comparable in size to the slit width, due to the defocus of the high resolution telescope, so we cannot characterize the spectrometer spectral resolution in response to true point sources. We report only the effective spectral resolution of the spectrometer to spatially extended sources, and show that it meets the design specification.

The detailed analysis of the IR spectrometer spectral resolution used four selected, unblended, isolated argon spectral lines fairly evenly spaced across the lower-wavelength two-thirds of the detector, from spectral images taken during TV4, supplemented by the longer-wavelength CH₄ and CO absorption bands imaged during TV2 and TV4, for characterization of the long-wavelength end of the detector. The emission lines were fit with a Gaussian together with a quadratic polynomial for the background, and this same method was adapted for use with sharp features in the CH₄ and CO absorption bands. Using this method the FWHM spectral resolution of the spectrometer was found to be approximately 2 physical pixels over the entire detector, with minimal variation (Table XVI), indicating excellent internal reimaging of

the entrance slit of the spectrometer onto the detector. The uncertainties in these FWHM values range from ~ 0.1 pixel for the argon lines to ~ 0.3 pixels for the CH₄ line and ~ 1 pixel for the CO line. Thus the spectral resolution of the spectrometer is indeed as designed: two physical pixels (one binned pixel) over the entire detector.

6. Absolute spectral sensitivity

The purpose of calibrating the radiometric response of the IR spectrometer is to enable us to accurately convert output instrument DN values to the spectral radiance of the scene imaged onto the detector at each pixel location. Analogous to the case of the CCD cameras, the pertinent equation for a single pixel in a given instrument operating mode at a given bench temperature assuming a linear instrument response function is

$$\text{DN} = tg\Omega a S_{\lambda} T_{\lambda} F_{\lambda} N_{\lambda} P_{\lambda} \Delta\lambda + \text{DN}_0(t),$$

where the notation is the same as for the CCD case except that no integral over wavelength is taken since the wavelength interval on each pixel is very narrow, and F_{λ} is the spectral transmission of the antisaturation filter for the central third of the slit (1.0 elsewhere), $\Delta\lambda$ is the wavelength interval imaged on a pixel (μm), and $\text{DN}_0(t)$ is the zero-exposure DN output for the given integration time for the same SIM bench temperature and detector temperature as those when the data were taken.

Since we know that the IR spectrometer response function is not linear, we must first adjust for the nonlinearity in the instrument response before we apply this function; our derivation of the IR absolute spectral sensitivity constants assumes that the required linearity correction has been applied to all data as a first step. The values of λ and $\Delta\lambda$ are known at every pixel using the wavelength maps described in Sec. VI C 4.

While an initial derivation of the IR spectrometer absolute sensitivity (i.e., $1/[g\Omega a S_{\lambda} T_{\lambda} F_{\lambda} P_{\lambda} \Delta\lambda]$) was done using ground-based radiance sources during TV calibrations, the most accurate determinations were done with in-flight calibrations using photometric standard stars. The stars used for this purpose are listed in Table VI; however, the measured SNR was high enough to be useful only for Vega, Sirius, and β Hyi. Absolute radiometric response calibrations were carried out four times in flight—in April, May, June, and July 2005. Note that when deriving the spectral response calibration using point sources as opposed to spatially resolved sources, the response equation is more usefully written as

$$[\text{DN} - \text{DN}_0(t)]_{\text{total}} = tgAS_{\lambda} T_{\lambda} F_{\lambda} I_{\lambda} P_{\lambda} \Delta\lambda,$$

where $[\text{DN} - \text{DN}_0(t)]_{\text{total}}$ is the total dark-subtracted signal from the star over all pixels, A is the area of the optics entrance pupil (m^2), and I_{λ} is the irradiance of the starlight at the optics entrance pupil ($\text{W m}^{-2} \mu\text{m}^{-1}$).

The objective of our in-flight stellar calibration is then to determine the product $gAS_{\lambda} T_{\lambda} F_{\lambda}$ for each pixel that causes the above equation to be true for any star's I_{λ} . From that, we can derive the absolute camera sensitivity factors required to convert from dark-subtracted DN/ms to radiance or reflectance for spatially resolved sources. $S_{\lambda} T_{\lambda} F_{\lambda}$ is assumed to be

the same across the entire IR detector array; any pixel-to-pixel or quadrant-to-quadrant differences in response are corrected by a separate flat-field correction as discussed in Sec. VIC 7 b.

The absolute spectral irradiances from 1 to 5 μm of the standard stars observed were obtained from a combination of data and model spectra.^{12,13,27–29} Morel provided *I*, *J*, *H*, *K*, *L*, *M*, and *N* absolute photometry of most of our standards; supplementary photometric measurements were obtained from Gezari and the 2MASS and IRAS catalogs. Pickles' stellar irradiance model spectra, which extend to 2.5 μm at a resolution of 0.5 nm, were chosen for the appropriate spectral type and luminosity and then scaled to fit the photometric data points in their spectral range. Smooth curves were fit to the photometric data points beyond 2.5 μm for each star to provide interpolations with higher spectral resolution. These scaled irradiance spectra were then used as our absolute references. Figure 79 shows the absolute irradiance spectra adopted for three of our standard stars. The uncertainties in our reference spectra are estimated to be $<10\%$ for a given star at a given wavelength.

Calibration data of the standard stars were obtained by smoothly scanning the IR spectrometer slit across the star location at a rate that moved the slit by one slit width during the time required to integrate each spectrum. Scanning was required to ensure capture of all the light from the star given the spectrometer's PSF and the spacecraft pointing uncertainty. To compute the total dark-subtracted signal from a star over all pixels, the DN values for each column were summed over all rows and over all frames in the scan that contained star signal. The dark level to subtract from each frame in the scan was determined by taking the median dark level at each pixel of frames in the scan on either side of those that contained star signal (excluding the first three frames of a given scan due to the transient dark level drop-off at the beginning of a series of frames) and linearly interpolating between them in time to find the applicable dark level for each pixel in a given frame containing star signal. The resulting total signal levels in each column for each star scan are shown in Fig. 80.

Quite good repeatability is observed between calibrations for the same star—no systematic changes with time were seen during the mission except for a possible hint of a few percent loss in response after encounter. Shot noise and CRs produce a fair amount of noise in the data of Fig. 80.

The in-flight response rates were used to derive radiometric calibration curves by dividing the stellar irradiances by the observed response rates and multiplying by the correct wavelength interval at each column. To create calibrations that convert from response rate to scene radiance for a spatially resolved scene (as opposed to the total signal from stellar point sources), the curves were also divided by the pixel solid angle. The appropriate wavelength map for the row in which the star image was centered and for the applicable bench temperature was determined using the previously described wavelength map equations (Sec. VIC 4). Curves were derived for the binned mode (behind the antisaturation filter only due to SNR limitations) and the unbinned mode (outside the antisaturation filter only).

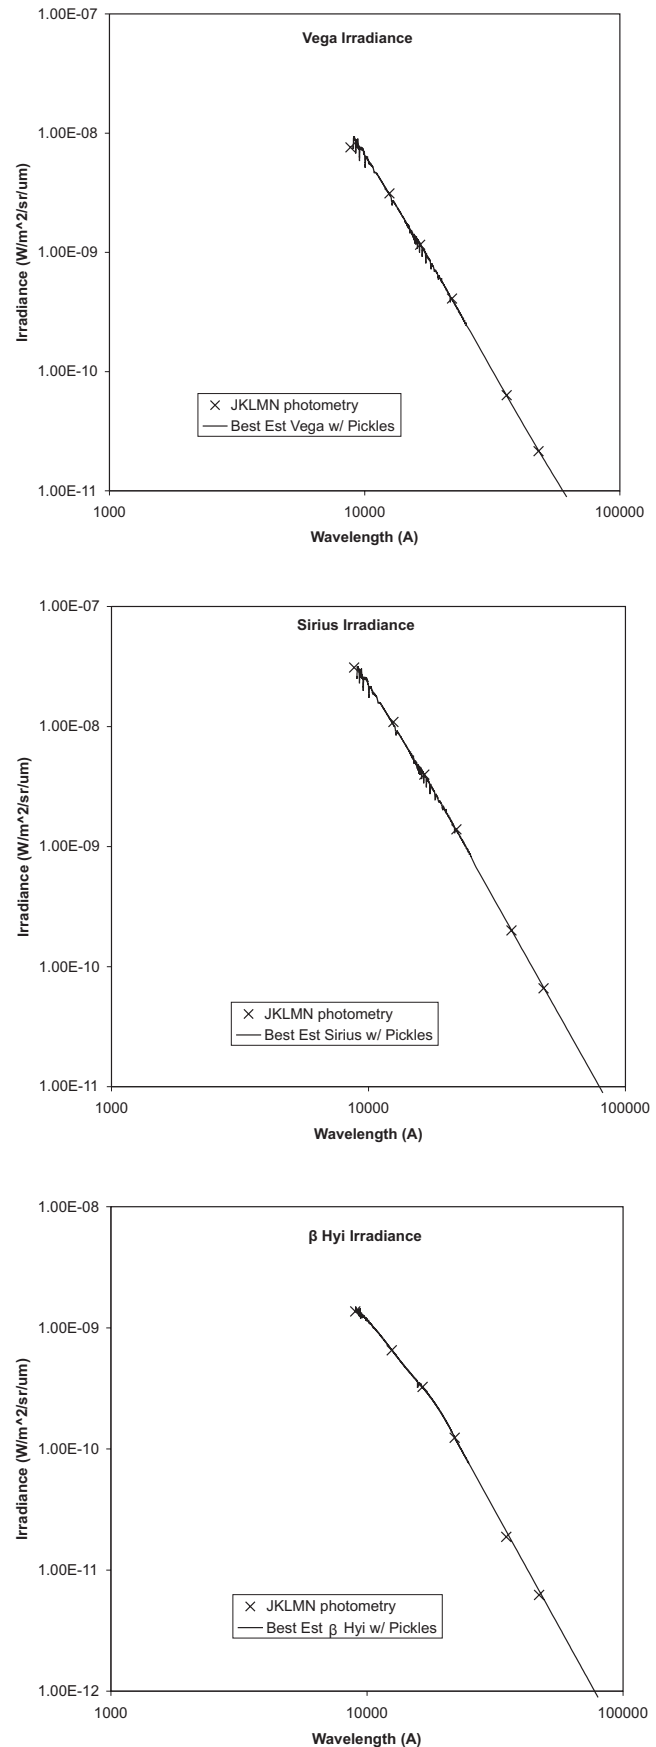


FIG. 79. Absolute spectral irradiance of selected photometric standard stars.

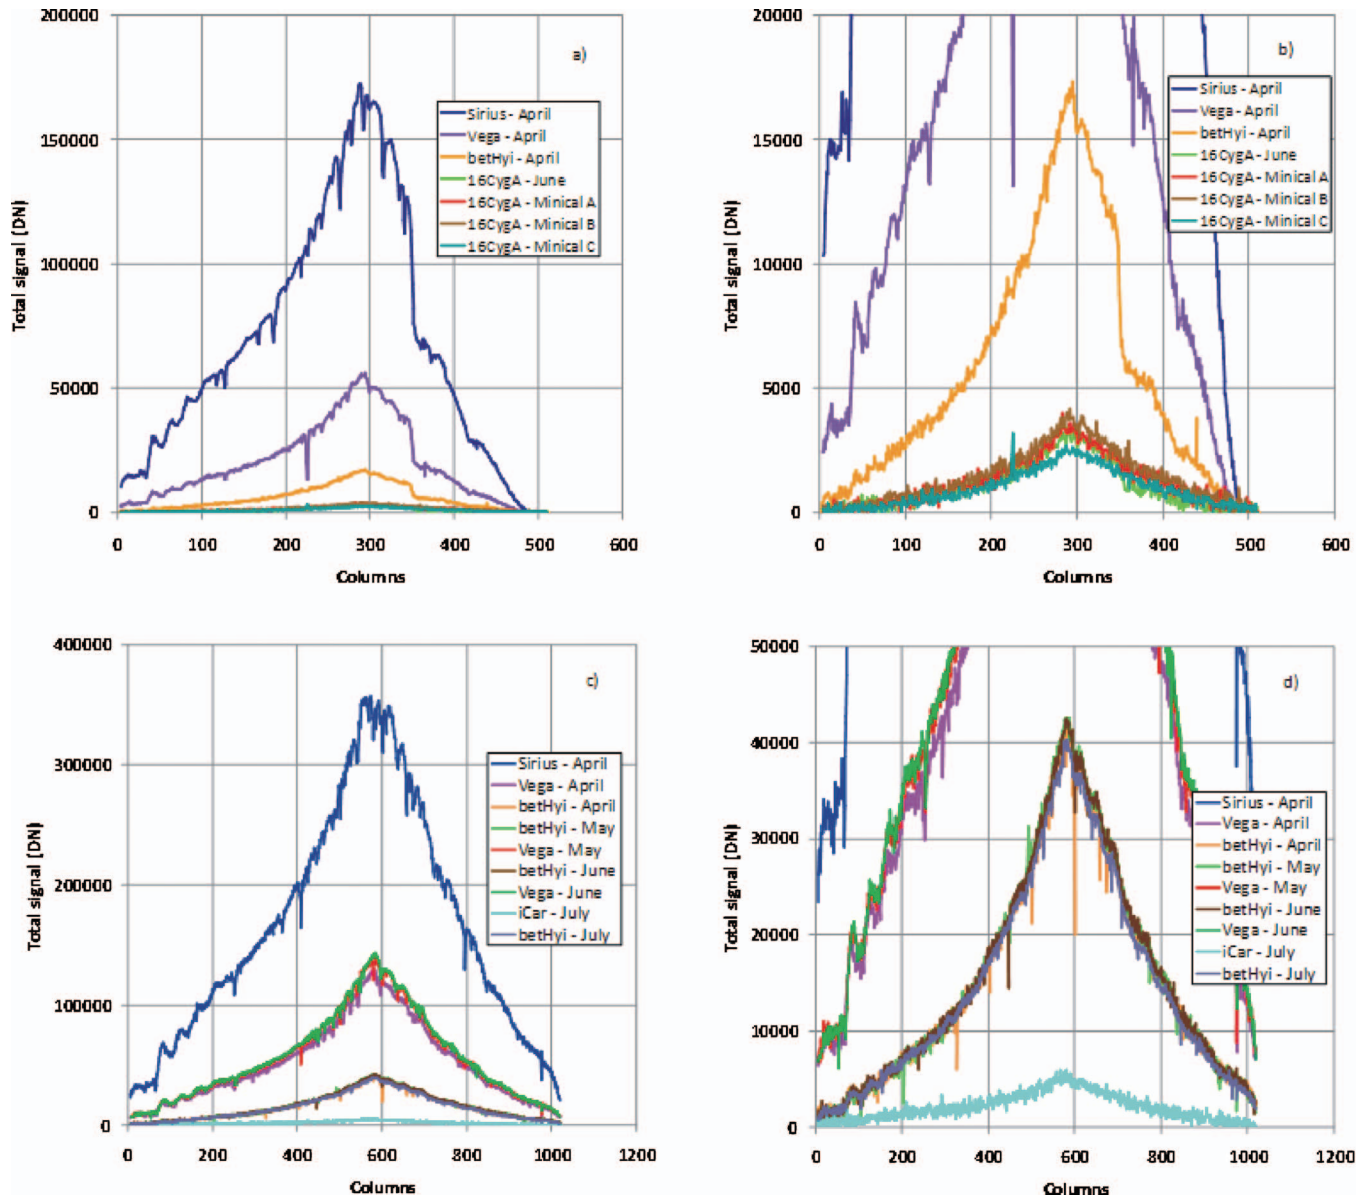


FIG. 80. (Color) IR spectrometer total response during inflight star calibrations: (a) binned mode, behind antisaturation filter except for minicals; (b) binned mode with expanded scale; (c) unbinned mode, all outside antisaturation filter; (d) unbinned with expanded scale.

Figure 81 shows the resulting calibration curves derived for each in-flight calibration. The relatively high values at the wavelength extremes reflect the loss of detector QE at these wavelengths. The humps around 2.7 and $3.7 \mu\text{m}$ in the curve behind the antisaturation filter are due to drops in transmission of the filter at those wavelengths. Obvious noise spikes were identified by their nonrepeatability in different calibrations; these have been smoothed over. Sirius' irradiance value initially appeared systematically too low by $\sim 20\%$ beyond $2.7 \mu\text{m}$; the literature irradiance was assumed to be in error and was scaled down in this region to make it match the results from other stars; the revised Sirius-based calibrations are included in Fig. 81. Sirius provides the best SNR at the longest wavelengths, so it is important to adjust its irradiance rather than to exclude it from the fits. Although we anticipated mismatches between calibrations in areas where the star signals exceeded the nominal IR detector saturation

level, we generally did not see such discrepancies. Total signal seems to be conserved until well beyond full well, perhaps as far as ADC saturation at 16 383 DN. Only the most severely saturated case, Sirius unbinned in April, showed a slightly too-high calibration curve toward the center of the range of saturated columns—data for Sirius in this range (1.4 – $2.9 \mu\text{m}$) were eliminated from subsequent data reduction steps.

Since the measured calibrations are repeatable for all stars throughout the mission, the different measurements behind and outside of the antisaturation filter were averaged to derive best-estimate curves with reduced random noise effects. These curves are shown in Fig. 82.

Using the ratio of the two curves of Fig. 82, the spectral transmission of the antisaturation filter was determined using in-flight data. Figure 83 shows the resulting curve. Also plotted is the preflight measurement of the filter trans-

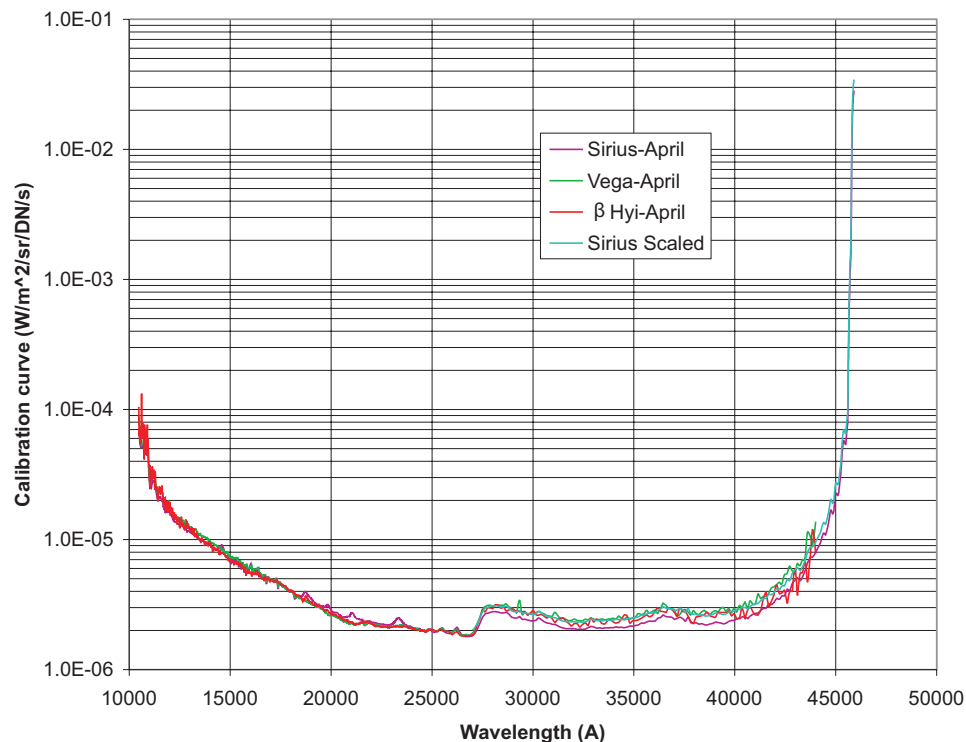
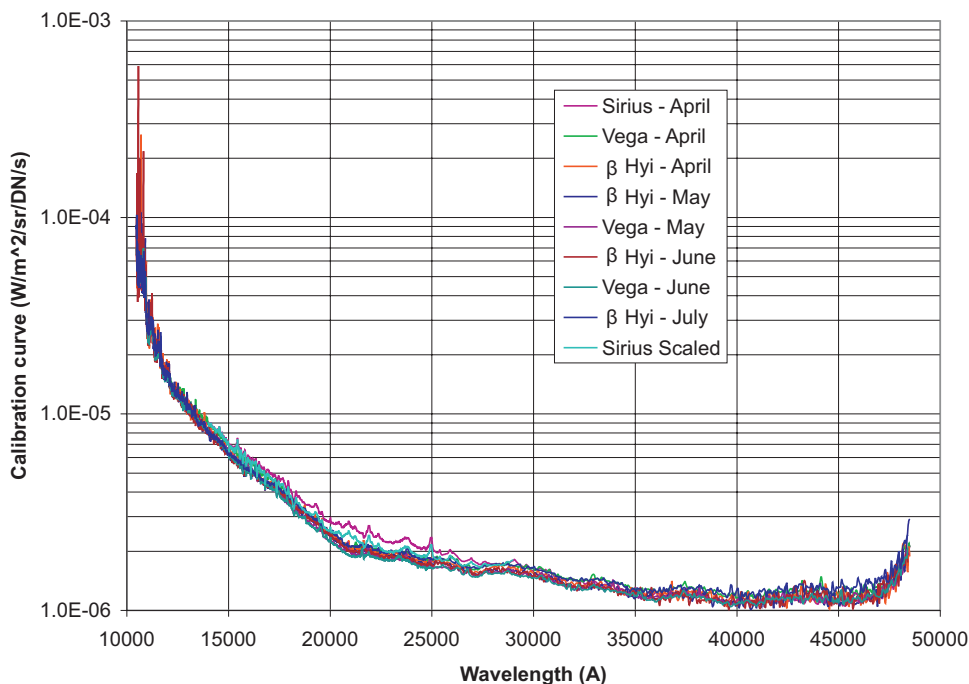


FIG. 81. (Color) IR spectrometer in-flight calibration curves; binned behind antisaturation filter on the top; unbinned outside antisaturation filter on the bottom.



mission scaled by a factor of 0.97, which fits the in-flight determination very well and eliminates the random noise effects.

To derive calibration curves that do not contain the high-frequency random noise effects that remain even in the best in-flight measurement, we took the approach of using the well-defined, low-noise prelaunch component-level measurements for the IR spectrometer (such as detector QE, telescope transmission, etc.) and adjusting a calibration curve derived from these component data to fit the in-flight measurements by multiplying it by smooth scaling curves versus wavelength. Unfortunately, we did not obtain a component-

level spectral transmission curve for the beamsplitter at the applicable in-flight temperature and incidence angle. We only obtained beamsplitter measurements at room temperature and the as-flown 45° incidence angle and at the flight temperature but at 90° incidence. These measurements gave us a general shape for the beamsplitter transmission curve, which we included in the curve derived from component data. They also showed the kind of ripples (amplitude and frequency spacing) in spectral transmission that we can expect from the beamsplitter in flight and showed that the nature of these ripples changes with both temperature and incidence angle. Therefore, the exact phasing and details of





FIG. 82. Best-fit averaged IR spectrometer calibration curves; behind antisaturation filter on the top; outside antisaturation filter on the bottom.

these ripples in the flight environment were not known or included in this initial model. We needed to estimate the higher frequency ripples in the beamsplitter spectral transmission curve using the in-flight results. Figure 84 shows the scaling required to fit the component-level curve to the in-flight measurement plotted versus the frequency of the incoming light. The ripples observed at $\sim 30/\mu\text{m}^{-1}$ and $\sim 5\%$ amplitude, growing to $\sim 20\%$ near $1\mu\text{m}$, are very similar in character to those observed in the preflight beamsplitter measurements.

We derived a smooth fit to the data in Fig. 84 that preserves the $30/\mu\text{m}^{-1}$ ripples but filters out the higher frequency noise spikes. The component-level-based calibration curve was then multiplied by this smooth curve to determine a final best-estimate in-flight calibration curve outside the antisaturation filter. The curve behind the antisaturation filter

was then derived by dividing by the scaled prelaunch transmission curve for the antisaturation filter. The resulting curves and the residual ratio between them and the best in-flight measured curve are shown in Fig. 85.

The calibration curves shown in Fig. 85 provide the conversion to units of spectral radiance for spatially resolved scenes. To convert to an I/F value, the radiance needs to be multiplied by π and divided by the solar irradiance at the proper range of the target from the Sun and the proper wavelength for each pixel based on the wavelength maps described in Sec. VI C 4

By virtue of eliminating star measurements with poor SNR from the data reduction and scaling the spectral irradiance of Sirius, the absolute IR spectrometer calibration is determined only by β Hya and Vega. The absolute accuracy is estimated at $\sim 10\%$ due primarily to uncertainties in the

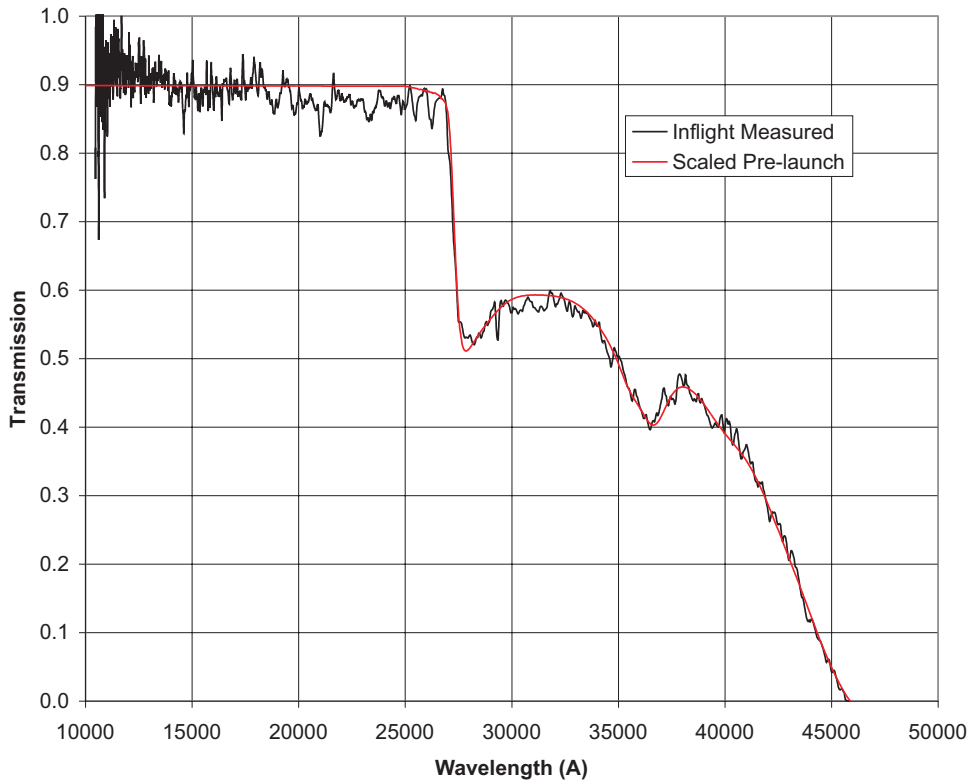


FIG. 83. (Color) IR spectrometer anti-saturation filter spectral transmission curve: black, inflight measurement; red, prelaunch measurement scaled to fit the inflight measurement.

stellar irradiance data. The residual noise in our fits is typically $\sim 3\%$ except below about $1.13 \mu\text{m}$, where the noise grows to typically 10% . The occasional spikes of several tens of percent below $1.13 \mu\text{m}$ are unlikely to be real calibration uncertainties; they are residual noise spikes in the star signal measurements. The in-flight calibration applies over the wavelength range of $1.0465\text{--}4.85 \mu\text{m}$ (but only to $4.585 \mu\text{m}$ behind the antisaturation filter).

7. Individual pixel response

a. *Bad-pixel maps.* IR arrays inherently have more defective and nonuniformly 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. 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. This group includes all pixels that have a response rate that is at least five times greater than the mean response rate of the column (spatial dimension) to which they belong.
- Cold pixels that always have a low signal. This group includes all pixels that have a response rate that is at least five times lower than the mean response rate of the column to which they belong.
- Pixels with atypical or very erratic responses. This group includes all pixels with
 - a root-mean-square error in a linear fit to the linearized response exceeding 2% ,

- a maximum residual in a linear fit to the linearized response exceeding 15% ,
- a maximum difference between the slope term of a linear fit to the linearized response of the pixel and the mean slope of its column exceeding 200% , and
- A response that could not be linearly fit (only two or fewer data points were unsaturated).

As an initial analysis, we used unbinned, full-frame, mode 6 darks taken at increasing integration times during ground calibrations (TV2 and TV4) for linearity analysis.




FIG. 84. Residual ratio between the component-level IR spectrometer calibration curve (excluding beamsplitter ripples) and inflight measured curve.

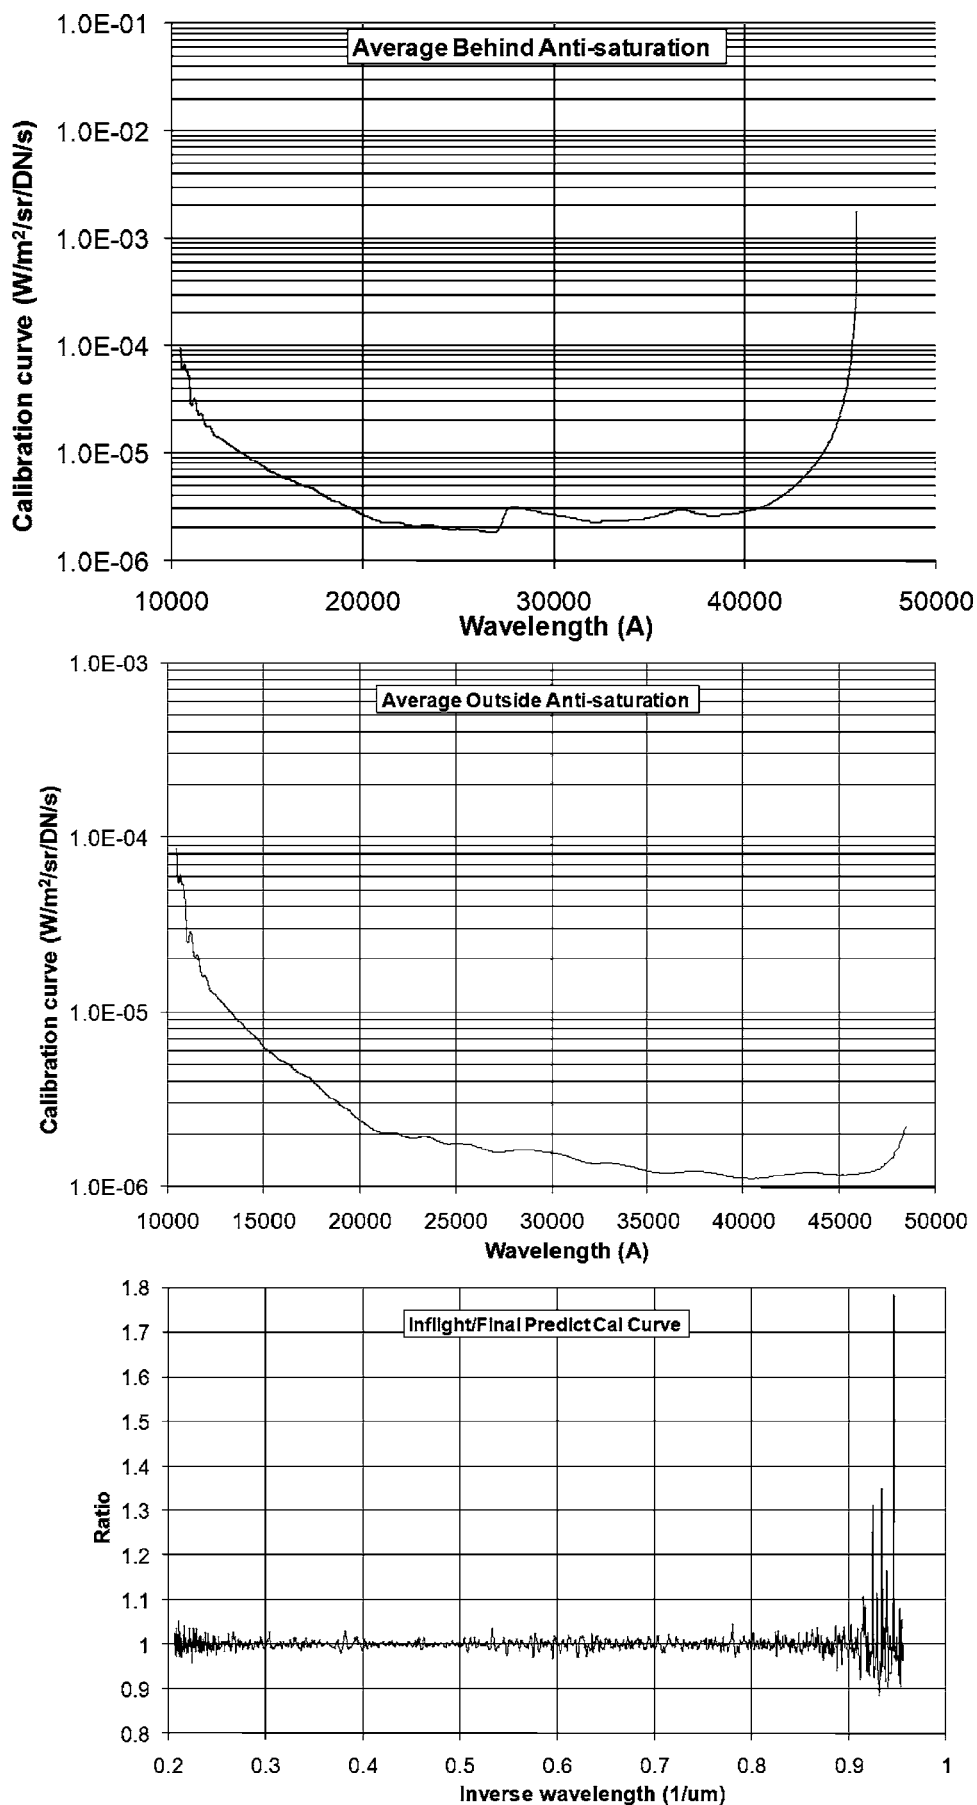


FIG. 85. Best-fit IR spectrometer calibration curves and their ratio to the in-flight measurements.

TABLE XVII. Number of IR spectrometer bad pixels identified using unbinned, full-frame dark data from ground calibrations (TV2 and TV4) and inflight calibrations in 2005.

Calibration Set	IR FPA temperature (K)	Optical bench temperature (K)	No. of bad pixels	% of bad pixels	Mode
TV2	82.6	130	5 036	0.98	6
TV2	85.1	136	4 940	0.96	6
TV2	87.2	141	4 259	0.83	6
TV4	106.3	139	37 482	7.31	6
April 2005	84.97	139.52	5 518	1.08	4(UBFF)
May 2005	84.28	137.80	5 823	1.14	4(UBFF)
June 2005	84.07	137.05	5 463	1.07	4(UBFF)
July 2005	84.37	136.93	5 870	1.15	4(UBFF)

Fortunately, the optical bench in the spectrometer proved to be a blackbody source that uniformly illuminated the array. Therefore, we used these data to establish the selection criteria provided above for identifying bad pixels.

We developed the following algorithm to identify bad pixels in the ground calibration and in-flight dark data taken at increasing exposure times.

- Make a median dark cube, with one median frame for each integration time. For in-flight data, the first frame of each integration set was excluded, when possible, from the median to avoid the known problem of the background level overshoot in the first frame in a sequence as described in Sec. VI C 3 c.
- For each median dark frame, linearize the raw DN values by applying the polynomial correction function developed during linearity analysis. The linearization polynomials were somewhat different for the two ground-based calibrations and the in-flight calibrations. The linearization polynomials based on the April, May, and June science calibrations were used to produce the bad-pixel maps for the pipeline.
- For each median frame, exclude saturated responses from the analysis (linearized DN > 12 000).
- For each spatial column, perform a linear fit of its mean response. For each pixel in the column, perform a linear fit of its response. Repeat this step for all active columns in the array. All linear fits attempt to minimize the Chi-square statistic based on percentage errors. (Note: the TV data were fit to $y=a+xb$. However, flight data were fit $y=xb$, i.e., the y intercept is forced to 0, because in all except mode 6 the IR signal should be 0 for an integration time of 0 ms by virtue of the read-reset subtraction that occurs. This simple fit is more appropriate, and it flags about 20%–25% more pixels as bad.)
- Select bad pixels using the criteria listed above.
- Make a bad-pixel map by setting bad pixels to a value of 1. Set all other pixels in the active image area to a value of 0, indicating good pixels. Set pixels in the reference rows and columns to 1.

Table XVII provides the number and percentage of bad pixels in the active image area using unbinned, full-frame




FIG. 86. Bad-pixel map for IR spectrometer unbinned full-frame modes 4 and 6 based on data acquired during the June science calibration.

dark data from the ground calibrations (TV2, TV4) and the in-flight science calibrations performed in April, May, June, and July 2005. The table includes focal-plane and optical bench temperatures for each calibration. The elevated focal plane temperature of TV4 caused excessive noise and a large increase in the number of pixel classified as bad. For April and May data, only three unbinned full frames (UBFFs) were taken at each integration time, and all three frames were processed for a median dark. Since five UBFFs were taken at each integration time in June and July, the first frame, which has a higher signal, was excluded from the median dark. Figure 86 presents an image of the bad-pixel locations.

During ground calibration, dark/linearity data were not taken for the binned modes. After attempting to make bad-pixel maps for binned modes using in-flight darks, we determined the results were not reliable. Therefore, we binned the mode 4 bad-pixel frame to make maps for modes 1, 2, 3, and 5. If any pixel in the 2×2 -bin area was bad, then the entire binned pixel was flagged. For modes 2 and 3, the mode 1 map was cropped spatially to the appropriate size (512×128 or 512×64 , respectively).

For the IR calibration pipeline, we elected to make a separate bad-pixel map for each mode for each in-flight calibration (see Table VI). The map from the April calibration is applied from launch until the May calibration; the May map from that calibration until the June calibration; the June map from that calibration until $E-8$ days; and the July map from $E-8$ days to the end of the mission. Table XVIII provides the number of bad pixels for each map used in the pipeline.

Once bad pixels have been identified, they are reclaimed or recovered by linearly interpolating between adjacent non-saturated, good pixels—one pixel above and below the bad pixel(s). 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 interpolated so that the user knows that their values are not true measurements (Sec. VII B). Our bad-pixel reclamation program will be placed in the DI PDS archive for use by other researchers.

b. *Flat field*. Each pixel in the IR detector array does not produce the same output signal when the spectrometer views a spatially and spectrally uniform scene. The corrections for spectral response variations and bad pixels were discussed in Secs. VI C 6 and VI C 7 a. The remaining pixel-to-pixel re-

TABLE XVIII. IR spectrometer bad-pixel maps used in the pipeline.

Calibration set	Mode	No. of bad pixels	% of bad pixels	Effective date for pipeline
April 2005	BINFF, ALTFF	3995	3.13	2005-01-12
	BINSF1	1907	3.04	
	BINSF2	866	2.85	
	UBFF, DIAG	5518	1.08	
May 2005	BINFF, ALTFF	4115	3.23	2005-05-10
	BINSF1	1969	3.14	
	BINSF2	898	2.96	
	UBFF, DIAG	5823	1.14	
June 2005	BINFF,ALTFF	4117	3.23	2005-06-07
	BINSF1	1984	3.16	
	BINSF2	901	2.97	
	UBFF, DIAG	5463	1.07	
July 2005	BINFF,ALTFF	4178	3.28	2005-06-27
	BINSF1	2006	3.20	
	BINSF2	920	3.03	
	UBFF, DIAG	5870	1.15	

sponse variations can be attributed to spatially dependent variations in the optical throughput of the spectrometer and to pixel-to-pixel response differences within the detector. We refer to these remaining pixel-to-pixel response variations as the flat-field response.

Attempts to measure IR flat fields were first made during the TV1, TV2, and TV4 test series. In TV1 and TV2, flat-field illumination was obtained using several blackbody hot-plate sources. The algorithm in Fig. 87 was used to create flat fields from the data. Two problems were found in the resulting flat fields: incomplete illumination of the slit and low SNR at the shorter wavelengths of the array, $\lambda < 3 \mu\text{m}$, due to the low maximum temperature achieved by the sources, $T < 600 \text{ K}$. Preliminary tests with a tungsten lamp (W lamp) using a small gold integrating sphere showed good signal throughout the $1\text{--}5 \mu\text{m}$ wavelength range, although problems with illuminating the entire slit persisted.

Further testing of the HRI-IR uncovered another prob-

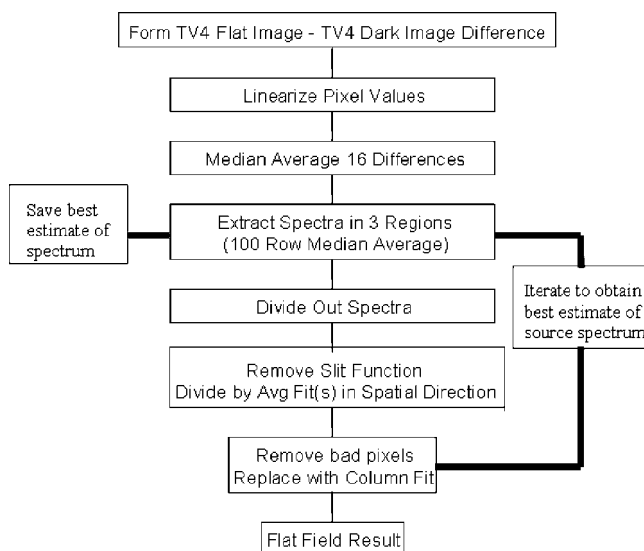


FIG. 87. Flow chart defining the algorithm for deriving the IR spectrometer flat field from images of a spatially uniform source radiance.

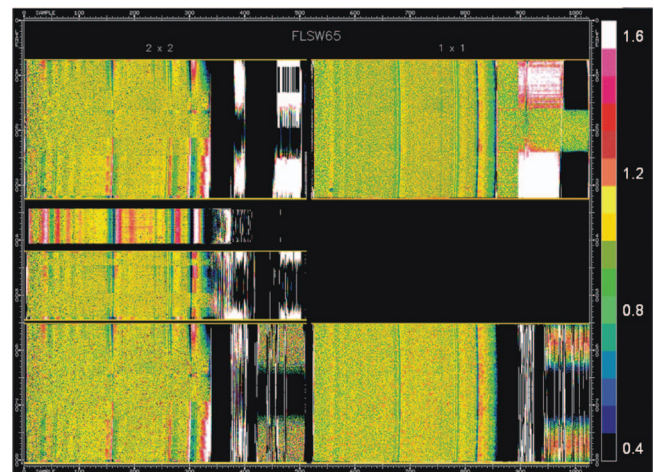


FIG. 88. (Color) Flat-field IR spectrometer images obtained during TV4 with the integrating sphere through the fused-silica window; brightnesses are normalized to the mean value. From top to bottom on the left are modes 5, 3, 2, and 1 flats, and on the right are modes 6 and 4. Wavelength varies from 1 to $5 \mu\text{m}$ from left to right in each image. The residual vertical structures on the right side of each frame result from regions of minimal input flux to the instrument due to atmospheric and fused-silica window absorptions.

lem with the instrument that made the TV1-TV2 flat fields problematic: a light leak, in which long-wavelength light entered the side of the HgCdTe chip and bounced internally, creating a series of “ghost bands” and an elevated background. Although this leak was blocked after TV2, it rendered the TV1 and TV2 flat fields unusable for calibration purposes.

After fixing the HRI-IR “ghost problem” by including a blocking shield around the edges of the detector, a second round of flat-field tests was devised for TV4. Conditions were not optimal, however, because the system was cooled only with liquid nitrogen (FPA temperature $\sim 105 \text{ K}$). Furthermore, the large fused-silica chamber window, while passing the optical and near-IR wavelengths from the large integrating sphere radiance source used, blocked wavelengths longer than $2.5 \mu\text{m}$, resulting in images for which only the short wavelengths were usable. A secondary chamber port was outfitted with a small, long-wavelength-transmitting ZnSe window, and light from a W lamp was reflected off a roughened gold scattering surface into the test chamber through a series of flat mirrors. Due to the limited room in the port, no focusing or light-shaping optics were included in the test setup, other than a bracketing aperture to eliminate straylight problems. This setup produced a high SNR for wavelengths from 1.7 to $5 \mu\text{m}$, with a modest SNR down to $1 \mu\text{m}$.

It was hoped that by combining the two sets of TV4 flats, a complete metaflat could be created. Good SNR images and flat fields with full illumination along the slit in the spatial (short) direction were obtained from $1\text{--}2.5 \mu\text{m}$ (covering approximately half the columns of the array) using the integrating sphere (see Fig. 88). No obvious structure was seen in the flat fields, other than atmospheric-absorption lines and high-frequency pixel-to-pixel gain variations with an rms variation of $< 5\%$. The W-lamp flats created by illuminating the array through the small ZnSe side window gave

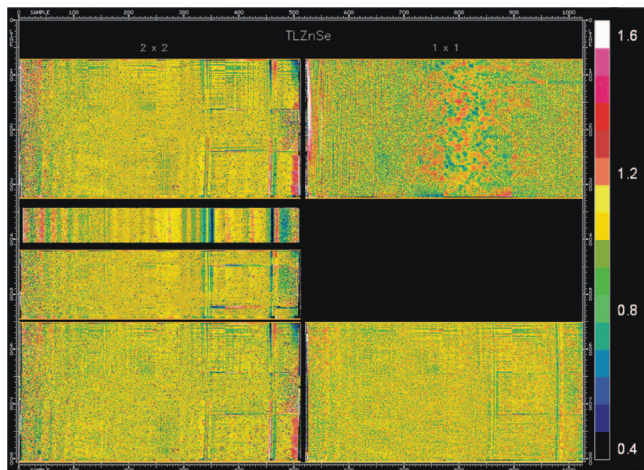


FIG. 89. (Color) Flat-field IR spectrometer images obtained during TV4 with the W lamp through the ZnSe window; brightnesses are normalized to the mean value. From top to bottom on the left are modes 5, 3, 2, and 1 flats, and on the right are modes 6 and 4. Wavelength varies from 1 to 5 μm from left to right in each image. Residual structures, some likely due to the illumination pattern of the W lamp reflecting off a rough gold surface, are apparent in both the dispersion (horizontal) and spatial (vertical) directions.

good SNR from 1.7 to 5 μm and moderate SNR from 1 to 1.7 μm . Unfortunately, large-scale, low-frequency structures were found in all of the W-lamp-ZnSe flats (see horizontal streaks near the top of the frame in Fig. 89). No corresponding structures were seen in the overlapping area of the array covered by the 1–2.5 μm fused-silica/integrating-sphere flats, proving that the observed structures were due to deviations from flat illumination in the ground support equipment (GSE) optics setup for the side port. Apparently the gold scattering surface with one reflection from source to instrument was not 100% efficient in randomizing, smoothing, and flattening the spatial structure in the lamp sources.

Because the slit was not illuminated uniformly in the side-port tests, a “slit illumination function” was adopted to provide a correction. This technique improved the 1–5 μm flats, but still left residual structure at the 10%–20% level—larger than the pixel-to-pixel variations estimated from the clean fused-silica window/integrating sphere tests.

After the creation of a metaflat field, it was tested by using it to flatten individual images obtained with Ar and Kr lamps through the ZnSe window. Results from these tests did not show a net improvement of the pixel-to-pixel rms variation (see Fig. 90) so that application of the flat fields derived in this way to the observational data was not found useful.

Since all the attempts at obtaining good quality IR flats from the ground-based calibration data were less than ideal, we turned to attempts to get flat fields from in-flight measurements using a large, relatively uniform celestial body as the illumination source.

The best object available for this purpose was determined to be the Moon, which was available only during the first 25 days of flight. The phase angle of the Moon during these times was $\sim 90^\circ$, providing a half-illuminated surface with variable lighting. The observing scheme used to provide spatially uniform illumination on the slit was to slew the slit rapidly across the moon’s surface in the direction of its long




FIG. 90. Application of IR flats to a TV4 HRI-IR Ar-lamp image. Signal from several selected rows is plotted. Squares, original data. Diamonds, data after flat-field division. Little to no improvement in the scatter of the data results from application of the TV4 flat.

axis during an integration so that all pixels would pass across the same chord of the Moon. In this way each part of the slit would see the same total light source.

The results from this technique were mixed. The first two attempts at obtaining lunar flats, at 4 and 16 days after launch, were compromised by saturation of the data in the first tests on $L+4$ days and erroneous pointing in the second test on $L+16$ days. The third and final test on $L+23$ days provided full unsaturated illumination of the slit with high SNR from 1.7 to 4.9 μm in modes 1 and 4. Although the optics bench temperature of ~ 147 K was elevated compared to the encounter temperature of ~ 136 K, the rms pixel-to-pixel variations were $< 5\%$, and no large-scale variations in the responsivity were seen across the chip (see Fig. 91). Very little signal was found in the shortest wavelength regions, but the artificial structure seen in the TV4 tests was verified as being due to GSE effects in the measurements through the ZnSe window.




FIG. 91. (Color) Flat-field IR spectrometer image obtained from the $L+23$ days lunar scan in mode 4; brightnesses are normalized to the mean value. Compared to the TV4 W-lamp/ZnSe window flat, there is almost no structure found in the long-wavelength (rightmost) regions. The horizontal green/red lines are at the boundaries of the central region of the array that is inside the antisaturation filter; some low-frequency structure is apparent at these boundaries. Low SNR at the shortest wavelengths and the longest wavelengths behind the antisaturation filter limits the usefulness of the flat at wavelengths < 1.7 μm ($<$ column 475) and > 4.1 μm ($>$ column 950) behind the antisaturation filter.




FIG. 92. (Color) Inflight IR spectrometer dark flat field produced using the thermal emission of the optical bench as the illumination source; brightnesses are normalized to the mean value. There is significant spatial structure in the derived flat compared to the TV4 and lunar measurements.

No complete flat was obtained from the tests conducted during the TV and in-flight measurements. However, by combining the results of the TV4 fused-silica window/integrating sphere and the $L+23$ days lunar flat measurements, it is clear there are no large-scale inhomogeneities greater than 10% in the response of the instrument. This finding allowed one other approach to measuring the flat field of the HRI-IR. The flat-field structure can be decomposed into effects due to the optics and effects due to the pixel-to-pixel response variations of the FPA

$$\text{flat field} = f(\text{optics}) * g(\text{FPA variations}),$$

where f and g are separable functions. Since the optical effects, which operate on large spatial scales (low spatial frequencies) appear small in the TV4/lunar tests, obtaining the FPA pixel-to-pixel response variation, g , should yield the majority of the variation in the flat, which occurs at high-spatial frequencies.

The approach taken was to use the passive thermal illumination of the FPA by the warm (~ 140 K) optics bench to determine the pixel-to-pixel response differences across the array. Although this background has a low temperature compared to the ~ 3000 K color temperatures of the W lamp and illuminating sphere sources, the large solid angle of illumination provides good count rates in all pixels, and since the optics bench source originates after the dispersing prisms, the array is illuminated at all wavelengths by large numbers of photons (the maximum response to the glowing bench is to wavelengths of $\sim 4.8 \mu\text{m}$ outside the antisaturation filter and $\sim 4.3 \mu\text{m}$ behind the antisaturation filter). HRI-IR dark images were used as the flat-field data. A few thousand DN's of signal were measured after a few seconds exposure during TV4. Admittedly, this technique makes the implicit assumption that the relative pixel responses are independent of the wavelength of the photons incident upon them, and it does not provide correction for any spatially dependent variations in the optical throughput of the spectrometer since the light collected at the FPA does not follow the same path and treatment as do sky photons through the HRI telescope optics. In addition, while this technique gave very flat response in the center of the array, it showed effects due to vignetting at the edges, indicating that the bench illumination pattern is not entirely uniform across the array (see Fig. 92).




FIG. 93. IR spectrometer linearity flat derived from a series of dark measurements in a linearity test. The low-frequency structure imposed by the instrument self-illumination pattern, seen in the dark flat image, has been removed by normalizing the response in each pixel to the median signal of its neighboring pixels in a 27×27 pixel area. The result leaves the high-frequency variations produced by the differences in gain from pixel to pixel. Relative-response values shown here range from 0.8 (black) to 1.2 (white).

Linearity tests were performed repeatedly in-flight by taking a series of dark exposures with integration times between 0.7 and 10 s, and these data were used to evaluate the pixel-to-pixel response variations. The signal provided by the self-emission of the instrument (i.e., the illumination from the warm optics bench), while nonuniform as noted above, is stable in time. The effects of illumination nonuniformity were removed by normalizing the response of each pixel to that of the median response value of its near-neighbor pixels in a 27×27 box. The pixel-to-pixel variations in the gain determined in this fashion were found to have a “salt and pepper” distribution across the array, with rms fluctuations on the order of 5%, similar to the variations seen in the best partial flat-field measurements (see Fig. 93).

Tests of the linearity flat field, conducted by dividing HRI-IR measurements by the flat field, resulted in slight improvements in the noise of the observations. Of all the flat-field tests that were performed, the linearity technique showed the best improvement when applied to the observational data. For this reason, it was adopted for use in performing the current version of the HRI-IR absolute calibrations. However, the remaining errors in the relative radiometric calibration at the individual pixel level due to flat-fielding uncertainties, including the low-spatial frequency components, can be as large as 5% at present. Work is likely to continue on improving the flat field, and in the future, it is possible that an improved version will be developed by piecing together the different wavelength regimes from the preflight measurements.

8. Noise

a. Random. The random noise in IR spectrometer frames consists of the detector read noise, signal shot noise due to photon statistics, and ADC quantization error. Measurements of the random noise levels in the spectrometer were made as part of the photon transfer analysis discussed in Sec. VI C 2. Figure 70 shows the relationship between the random noise level and the signal level illustrating that we are shot noise dominated over most of the detector's dynamic range. The measured read noise levels are 2 DN for unbinned modes and 1 DN for binned modes. Assuming a perfect ADC, quantiza-




FIG. 94. Histogram of IR spectrometer dark image showing modest uneven ADC bit weighting.

tion error would be at only the 0.3 DN level and would be insignificant. Actual ADC performance is discussed in Sec. VI C 8 c.

b. *Coherent.* The image noise in the IR spectrograph is dominated by random noise, and to the measurable level there was no obvious coherent noise in the IR system within a single frame. There are systematic changes in the dark signal level with time, as described in Sec. VI C 3 c. However, these changes are on a frame-to-frame basis. Interference from other electrical systems has not been seen.

c. *ADC.* Analysis of the accuracy of the digital encoding of the IR analog signals was performed using dark images from TV2. The signal at long wavelengths from the warm external room radiance transmitted through the chamber window results in images with a smoothly varying left-right signal gradient. For such an image, an ideal ADC would output data that have a smoothly varying histogram with approximately equal numbers of DN's within the neighborhood of each signal level in the image. Figure 94 shows a representative histogram from one of the TV2 dark images. The histogram indicates that the IR ADC performance is somewhat less than ideal, with encoding bin sizes that vary from level to level by perhaps as much as 10%. This small degree of uneven bit weighting is not significant for the IR spectrometer. At the minimum in-flight background signal levels, photon shot noise completely dominates the encoding errors.

d. *Compression.* Due to the need to preserve a significant amount of dynamic range for all IR data, compression of IR spectral images from 14 to 8 bits via a LUT generally results in noise no longer being dominated by random noise, but rather by digitization noise introduced by the compression process. Due to the higher than desired temperature of the IR spectrometer bench, there was always a significant background from the bench that limited the smallest external signal that could be observed. Almost all IR calibration data were taken uncompressed to preserve highest available SNRs, but due to limitations of onboard data storage and downlink capacities, much of the IR data taken at encounter was compressed. During the months preceding encounter, the expected range of possible operating temperatures of the IR spectrometer bench was characterized and the IR LUTs optimized to cover as small a range of input values as possible,

TABLE XIX. Comparison of encoding step size to signal shot noise for IR spectrometer data compressed using various LUTs

LUT	Uncompressed DN range	Compressed DN step/shot noise
0	245–11 500	3.2
1	1 040–4 018	2.43
	4 019–11 500	5.15
2	1 120–6 009	2.44
	6 010–11 500	7.21
3	2 500–8 517	2.2
	8 518–11 500	8.62

to minimize the amount of noise added by compression of the data, resulting in the final encounter IR LUTs described in Sec. III E. Because the IR LUTs use a square-root structure, for each square-root range the compression noise is a simple multiple of the photon shot noise. All of the IR LUTs except LUT0 are dual-slope power laws; that is, the uncompressed DN range is divided into two ranges each having a different conversion from $\sqrt{\text{DN}_{14}}$ to DN_8 . Table XIX summarizes the range of uncompressed DN's spanned by each LUT and the ratio of the compressed DN encoding step to the shot-noise value that applies for each LUT.

9. Scattered light

In order to support science objectives for the HRI-IR spectrometer, it is necessary to determine the extent of any scattered light in the instrument. This is particularly important for the DI mission, as one of the major science objectives is to determine the composition of the coma before and after impact. To do so, it is necessary to be able to measure a faint coma signal near the bright nucleus. Our calibration efforts to characterize both in-field and out-of-field scattered light included both ground-based laboratory measurements and specific in-flight data collections.

The primary in-field scattered light test preflight was a series of knife-edge tests. In these tests (which took place on 6 March 2003), a horizontal blocking mask was placed at 52 positions across the slit blocking the light in a spatial direction to create a horizontal knife-edge across the focal plane array. The light source for these experiments was a lamp that illuminated a white sheet, which was observed through a fused-silica window. Rather than a sharp vertical falloff, signal beyond the knife edge was observed, as shown in Fig. 95. These data indicate that there is $\sim 3\%$ – 4% scattered light 10 pixels from the knife edge and $\sim 1.5\%$ – 2% 20 pixels from the knife edge. No differences in scattered light as a function of wavelength (i.e., column number) are observed. In addition, no specific effects from the antisaturation filter are observed. These results provide an upper limit on the in-field scattered light in the HRI-IR spectrometer. Some or all of the observed effects could be due to the ground support equipment and ambient environment.

The in-flight scattered light test was carried out on 4 February 2005, using the Moon as a target. The slit was scanned at a constant rate across the Moon. The reconstructed images yield an image cube with two-dimensional images of the Moon in each of 512 wavelengths from

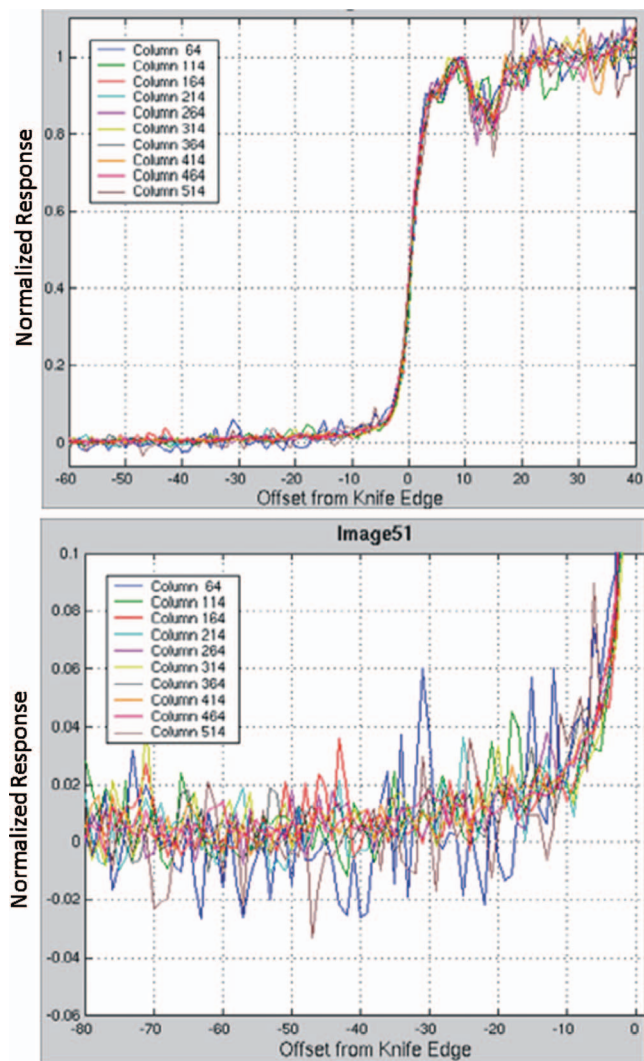


FIG. 95. (Color) IR spectrometer response profiles across a knife edge from the ground-based in-field scattered light test. Signal is detected 10–20 pixels away from the knife edge at a level of 3%–4%. Column numbers represent increasing wavelengths from 1 to 5 μm across the detector. No change in scattered light vs wavelength is seen.

1.08 to 4.8 μm . In the absence of any scattered light, the signal at the edge of the moon should fall off abruptly. However, as seen in Fig. 96, significant signal exists beyond the lunar limb and terminator. The spectral signature of the Moon is detectable in the dark sky up to 20 pixels from the limb. Images from three different wavelengths are presented (1.15, 1.90, and 3.00 μm) indicating no spectral dependence

to the scattered light components. In addition, a faint ghost image of the Moon is seen offset from the primary image.

The scattered light effects are quantified in normalized profiles taken both horizontally and vertically, as shown in Fig. 97. Here profiles from 2.5, 3.0, and 4.0 μm are plotted revealing in-field scattered light of 1% at 10 pixels in the horizontal direction and 2% at 10 pixels vertically. Neither shows a spectral dependence.

In addition to the scattered light off the limb, the lunar images also include a ghost image (see Fig. 96). This ghost image is the result of an internal reflection off the beamsplitter (see below). While the ghost is visible in the lunar images, it is best quantified from star images. For the out-of-field scattered light calibration, we collected saturated data of the star Canopus on 28 April 2005. A composite image created from three bands near 2.1 μm (the wavelength with maximum throughput in the IR system) from a scan of the star is shown in Fig. 98. In this saturated image, the effects of the defocused telescope (star pattern) and a single clear ghost image (at bottom) are clearly visible. The ghost level is quantified in vertical profile (Fig. 99) from the scan in Fig. 98 and a similar scan in the opposite direction. In the first case, the scan was in the $-y$ direction at a rate of 10 μrad (one slit width) per frame time. This scan produces a 4% ghost image 36 frames, or 360 μrad , after the star. The second scan was in $+y$ direction at a rate of 25 μrad (2.5 slit widths) per frame time producing a 2% ghost 13 frames, or 325 μrad , after the star. Since the two scans were taken at different scan rates, the amount of signal/pixel is different. Correcting for the difference in scan rates (25/10 μrad) increases the 2% ghost by a factor of 2.5, leading to an $\sim 5\%$ ghost at a nominal scan rate of 10 $\mu\text{rad}/\text{slit}$. The ghost image is also offset horizontally ~ 1 –2 pixels from the star. During the cruise phase of the mission, various other stars were observed with the HRI-IR spectrometer and showed similar results to those produced from the Canopus data.

In addition to HRI-IR spectral data, visible images of Canopus were also collected. In these data, a ghost of the star is detected 155 pixels vertically below the star. The HRIvis camera has a resolution of 2 $\mu\text{rad}/\text{pixels}$, placing the ghost image 310 μrad from the star position. Horizontally, the ghost is offset 6–12 visible pixels (12–24 μrad), very consistent with the 1–2 pixels offset in the HRI-IR data.

The ground-based out-of-field scattered-light test showed very similar effects. An off-axis pinhole test was set up for HRI on 31 August 2002. Here, the setup consisted of




FIG. 96. (Color) Gaussian-stretched, dark-subtracted IR spectrometer images of the Moon (composite RGB: 1.15, 1.90, 3.00 μm) showing scattered light around the lunar limb and terminator from two different observations. The arrows indicate the locations of the profiles plotted in Fig. 97.

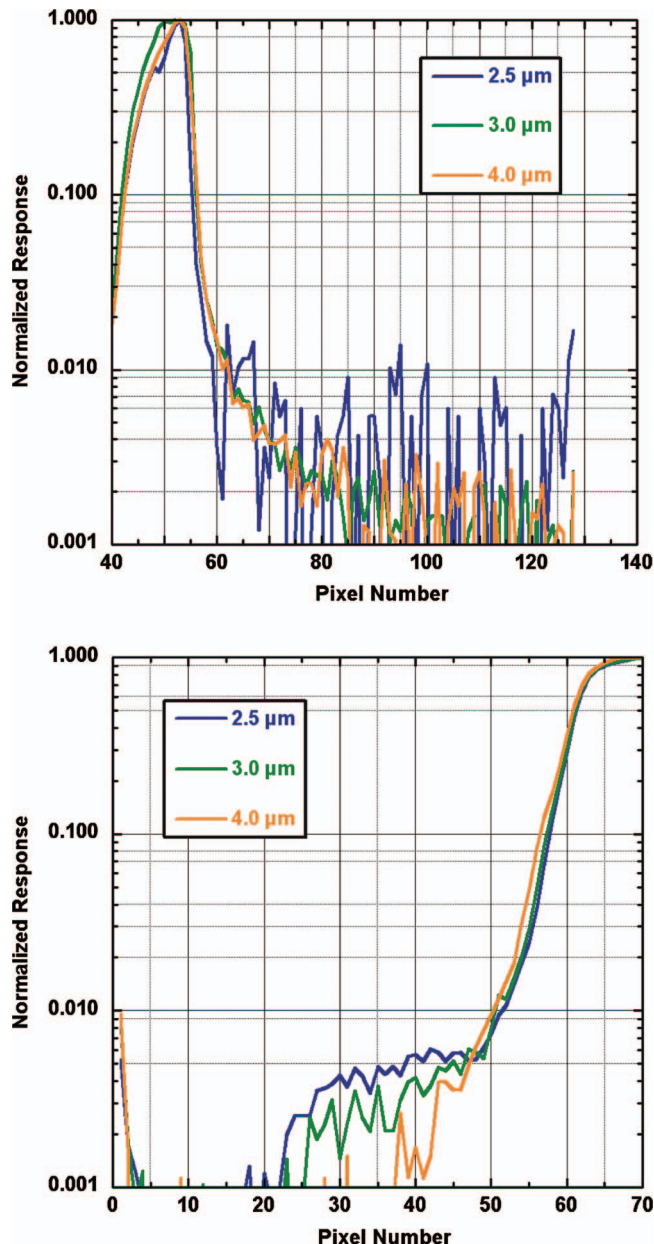


FIG. 97. (Color) Horizontal (top) and vertical (bottom) IR spectrometer profiles of scattered light off the lunar limb (see arrows in Fig. 96). Scattered light levels of 1% are detected 10 pixels from the limb in the horizontal direction and of 2% at 10 pixels from the limb in the vertical direction. The scattered light profiles at all three wavelengths (2.5, 3.0, and 4.0 μm) behave similarly.

a tungsten lamp illuminating the field of view through three 0.5 mm pinholes. Over the course of the experiment, the holes were moved horizontally with respect to the spectrometer slit in the cross-slit direction of -10 – $+5$ mm in 0.5 mm steps. (The slit was centered at the cross-slit position of -2.75 mm, and thus the actual positions are -12.75 – $+2.25$ mm relative to the IR slit.) Data were alternately collected with the HRI-IR and HRI vis at each position. Profiles of the HRI-IR response versus cross-slit source position for each of the three pinholes were generated. The HRI-IR data were saturated at the cross-slit positions near the center. The HRI-IR profiles versus slit position can therefore only provide qualitative evidence of scattered light.




FIG. 98. (Color) Gaussian-stretched, dark-subtracted, composite IR spectrometer image of three bands near 2.1 μm from a scan in the $-y$ direction of the star Canopus. As discussed in the text, a ghost image is clearly detected towards the end of the scan (bottom of image).

In addition, any scattered light observed could be due to the GSE. However, if the GSE were contributing to the scattered light in the HRI-IR, it should also produce scattered light in the HRI vis. To compare to the HRI-IR data, the HRI vis images were therefore also analyzed. When comparing the HRI-IR and HRI vis profiles versus cross-slit position, one peak of scattered light is seen in both data sets, suggesting it is from the GSE. However, a second peak is seen only in the HRI-IR at slit positions -3.5 and -4.0 mm with respect to the pinholes, which correspond to 350–400 μrad .

The position of the ghost image, now seen at the same location in ground and in-flight measurements, is very consistent with models of reflections off the beamsplitter. The ghost is not seen in HRI vis ground-based pinhole measurements because the transmission of the beamsplitter is ~ 0.999 in the visible and ~ 0.85 in the IR making the intensity of the reflection 150 times greater for the same input intensity. However, as discussed above, the ghost image is seen in in-flight HRI vis images of highly saturated stars such as Canopus.

Both in-field and out-of-field scattered light have been documented in the HRI-IR spectrometer. While scattered light in the HRI-IR is not insignificant, a full, quantitative measure of all the scattered light sources has not been done. The scattered light has a number of components: the spatial PSF of the telescope, the spatial PSF of the spectrometer, the spectral PSF of the spectrometer, reflections and scattering associated with the beamsplitter, and possibly other sources. Decoupling and characterizing all of these effects is beyond the scope of this work, and automated corrections have not been added into the pipeline processing. Thus, care must be taken during analysis of the HRI-IR data. In particular, researchers working on the science of the inner coma should be very cognizant of scattered light issues. Individuals may wish to use deconvolution techniques to diminish the effects of scattered light. Unfortunately, the level of background noise in the spectrometer due to SIM bench emission precludes measurement of the PSF out beyond about 30 pixels from the source (see Figs. 95 and 97), so a full-frame PSF cannot be determined, which limits the use of deconvolution techniques to fully correct for scattered light analytically.

10. Charge bleeding/residual

The pixels on the IR FPA are well-separated electrically so that even pixels saturated with charge up to several times full well were not seen to spill charge over to neighboring

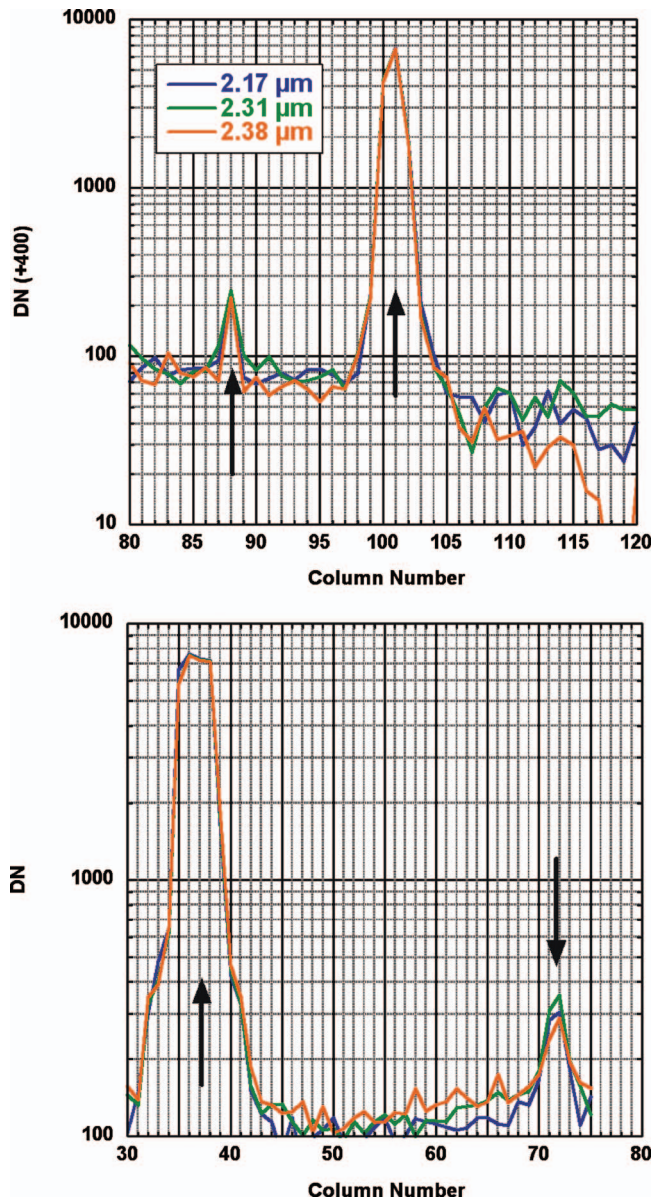


FIG. 99. (Color) Vertical profiles across the IR spectrometer Canopus scans. Top: Scan in the $+y$ direction showing a ghost image at 13 frames ($325 \mu\text{rad}$) after the star image. Bottom: scan in the $-y$ direction showing a ghost image at 36 frames ($360 \mu\text{rad}$) after the star image. As described in the text, when corrected for differences in scan rates, the position and magnitude of the ghost images are very similar.

pixels. This means that once a pixel reaches the ADC saturation level, there is no chance to reconstruct the total charge in a scene.

Ground tests designed to examine effects of charge residual showed no detectable residual in pixels that were saturated in a previous image. However, the results described in Sec. VI C 3 c indicate that the results of the residual test should be re-examined in terms of the overall detector's zero-level signal. A working hypothesis for the first-frame offset shift is that the total charge in the detector will shift the detector zero level. This has not been definitively shown by test, but many saturated pixels may result in a measurable shift.




FIG. 100. The number of clusters detected as CR events per s/cm^2 vs the number of pixels affected for out-of-peak solar activity in a typical calibrated dark IR spectrometer frame (1024×512) with $t=7.7$ s.

11. Radiation noise

IR images present some different problems in detecting CR events than are found in CCD images. The IR images have a much higher level of background noise due to the large signal produced by the SIM bench emission. In addition, the IR detector has many badly behaved pixels that produce either far more or far less signal than a typical pixel when exposed to a uniform illumination field. However, CR events are seen in IR images, and we have attempted to characterize them so they can be detected and corrected to some degree.

The most reliable way of detecting CR events in IR images is to difference two IR frames taken in rapid succession to find the changes, most of which will be due to CRs if they are above the noise floor. Most IR data sets consist of multiple frames, either from a spatial scan or from repeated exposures of the same scene, and one can therefore take advantage of this internal redundancy to isolate CR events.

In an alternative approach, IMGCLEAN was run on two successive calibrated IR frames (including bad-pixel and flat-field corrections), and the number of charge clusters that were different between the two runs was determined. IMGCLEAN classified about 2–3 events per s/cm^2 as CRs, consistent with the rates seen in visible images. Most of these CR charge clusters consist of only 1–4 pixels (see Fig. 100) consistent with the sizes seen in CCD images. Long CRs were seen less often on IR images than on visual images (at the same exposure times).

The number of charge clusters classified as CRs by IMGCLEAN run on a single calibrated IR image is typically greater by a factor of about 5 than the number of clusters on this image that are not present on an immediately sequential identical neighboring image. Most of the clusters deleted by IMGCLEAN from a calibrated IR dark image can also be found in the same locations on other images and are not true CRs. In many cases, they contain known bad pixels in the detector. Running a “bad-pixel” reclamation routine prior to IMGCLEAN eliminates about half of these false detections.

In our test runs with calibrated IR dark images, only IMGCLEAN worked correctly, and then only after bad-pixel and flat-field corrections. CRFIND and DI_{CRREJ} had problems running to completion with all parameter settings we tried.

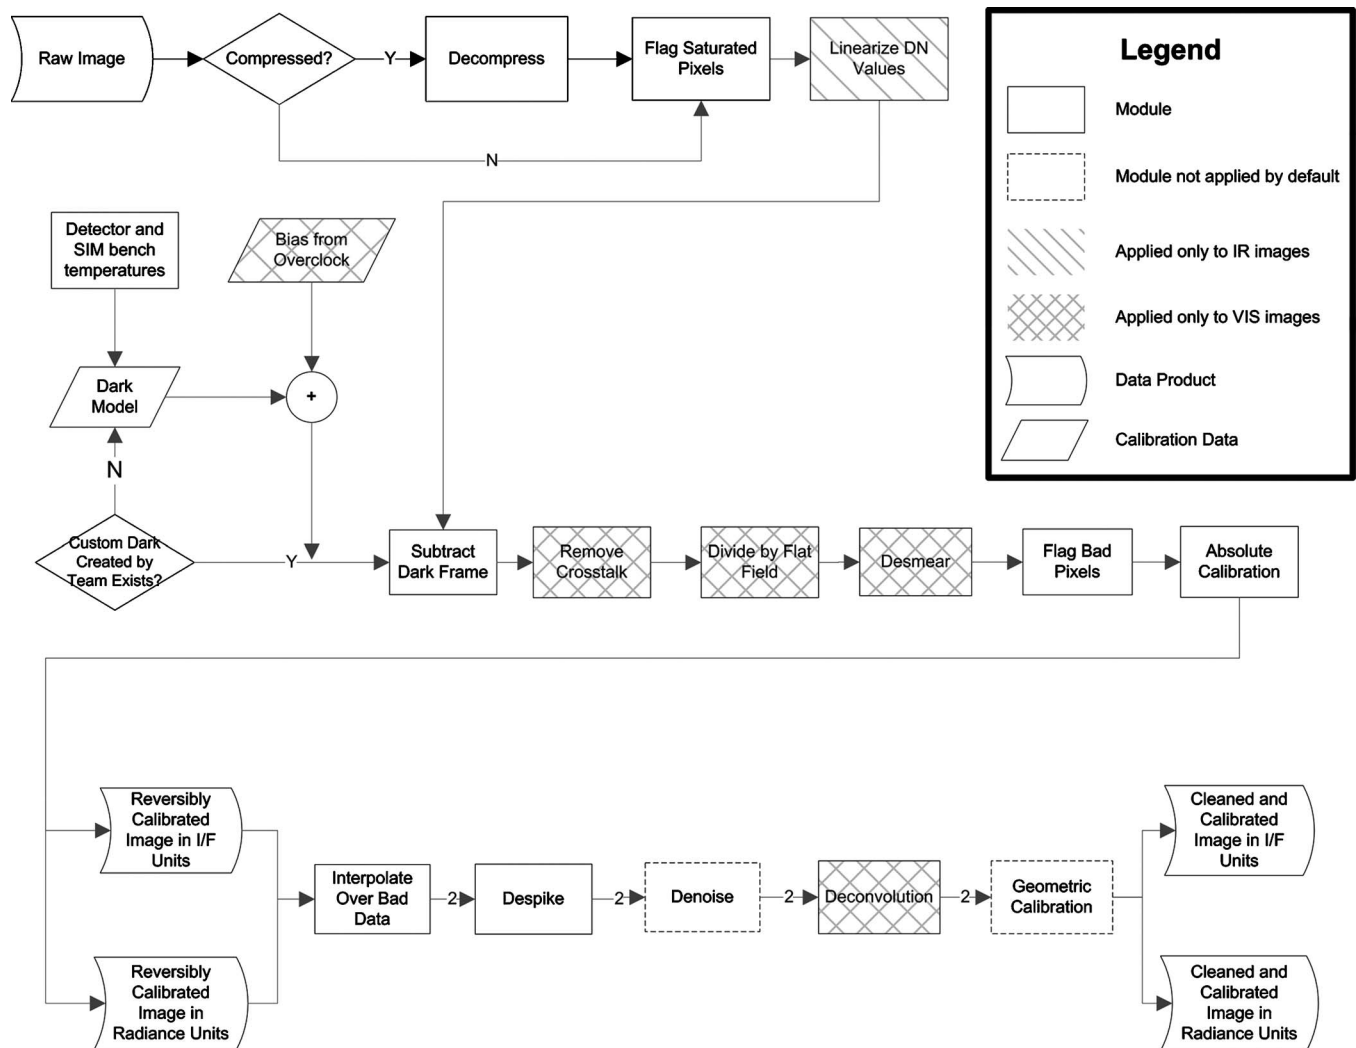


FIG. 101. A flowchart describing the image processing pipeline used to calibrate Deep Impact images. Note that some modules are not applied to all instruments.

RMCR needs a relatively constant background, which is not necessarily the case for IR frames. These programs may be useful for removing long cosmic rays if they are present on IR images, but in this case one must be careful not to remove any long lines that belong to the spectrum or to stars (i.e., remove only oblique CRs, but not vertical or horizontal lines). None of the programs have yet been made to work successfully on nondark (either calibrated or raw) IR images; the number of false CRs can exceed the real number by a factor of 100 in runs made to date. Work is ongoing to develop CR detection approaches that are generally applicable to IR images.

VII. PIPELINE PROCESSING

In order to produce calibrated archival data records, the DI science team developed a processing pipeline to systematically carry out the normal calibration steps using the best available calibration files and constants. This processing system was implemented at Cornell University. As part of the calibration process, auxiliary files were generated in image format and appended to each image record that map for each pixel the quality of the calibrated value and its SNR. For

each calibrated IR frame, a spectral registration map is generated and appended indicating the wavelength that applies to the calibrated radiance or reflectance for each pixel. Proper scientific analysis using calibrated DI data requires an understanding of the standard processing steps and auxiliary files outlined here.

A. Standard steps

For each image, there is a standard set of procedures and settings applied in our pipeline processing in order to calibrate the images automatically (see Fig. 101). In general, these default settings are the best the science team has been able to derive for the data set as a whole and thus do not necessarily reflect the best possible processing for any particular image. However, there are some observations around encounter, especially with the IR spectrometer, that contain very valuable scientific information but are not processed optimally by the default settings. For these cases, the automated pipeline has the ability to specify special settings for particular observations.

The standard pipeline begins by decompressing the image if it was compressed on the spacecraft. Images can be

compressed using one of four 14–8 bit LUTs optimized for different types of exposures. To uncompress the images, a reverse LUT is used that maps each 8 bit value to the average of all corresponding 14 bit values.

All saturated pixels are flagged in the quality map (Sec. VII B). Then an IR image is linearized using the correction described in Sec. VI C 1. Linearization is the first step for IR data because the IR output data represent a read-reset DN value causing all output DN values to represent signal collected in the detector (any fixed bias is subtracted out) and to be subject to the nonlinearity of its response. A vis image does not need this step because the instrument responds linearly (Sec. V C 1).

Next, a dark frame is subtracted from the image. If a dark frame has been created by the science team for the specific observation, then it is subtracted. Otherwise, a dark model is used to generate the frame (Secs. V C 3-vis and VI C 3-IR).

After the dark subtraction, a vis image undergoes a few extra processing steps not taken by every IR image. First, the electrical cross-talk (Sec. V C 7 e) is removed by subtracting a derived ghost frame. Each quadrant in this frame is a linear combination of rotated versions of the other three quadrants. Next, the image is divided by a flat field (Sec. V C 6) in order to account for variable responsivity across the detector. A flat field is only applied to unbinned IR images because the best binned-mode flat field does not seem to provide any noticeable improvement in SNR (and in the data products published as PDS version 1, unbinned IR images are not flat fielded either). Lastly, vis CCD frame transfer smear is removed using the parallel overclock rows if the image was taken in modes 1–6 or a column averaging approximation if the image is in modes 7 or 8 (Sec. V C 4).

After bad pixels are flagged, the image is radiometrically calibrated to produce a radiance image in $W/m^2 sr \mu m$ and an I/F image (vis cameras only). For a vis image, this is simply done by dividing the image by integration time and then multiplying by the appropriate conversion factor derived in Sec. V C 5 for the given filter and desired output. For an IR image, the procedure is more complicated as the absolute calibration is wavelength dependent, which in turn is temperature dependent. First, the wavelength and bandwidth for each pixel are calculated, as described in Sec. VI C 4. Then, each pixel is multiplied by the appropriate wavelength-dependent calibration factor (Sec. VI C 6) and divided by integration time and the pixel's spectral bandwidth. Once this radiance image is created, a copy is converted to I/F for the vis cameras by dividing by the solar spectrum at the target's distance from the sun and then multiplying by pi.

At this point, two reversible data products have been created, one radiance image (all instruments) and one I/F image (vis cameras only), and copies are run through the rest of the pipeline, which performs a series of nonreversible steps. First, the data are interpolated over the bad pixels and gaps. For a vis image, this interpolation is performed using thin plate splines anchored by the valid data around the edges of each hole. For an IR image, a linear interpolation is performed in the spatial dimension only.

Next, a despiking routine is applied in order to remove cosmic rays. This routine performs a sigma filter by calculating the median of each $N \times N$ -pixel box, where N is odd, and then replacing the central pixel with the median if it is more than M median deviations from the median. By default, both M and N are set to 3. The median deviation of a set S is defined as $\text{med}(|S - \text{med}(S)|)$.

Lastly, a vis image is deconvolved using the methods described in Sec. V B 3. This step is especially important for the HRI vis instrument which is out of focus.

B. Calibration quality map

Along with each calibrated image, a byte map is created that defines the data integrity for every pixel. For each byte in the map, representing 1 pixel, each bit acts as a flag that is set to 1 if the given criterion is met for that pixel. These flags are

MSB							LSB
7	6	5	4	3	2	1	0

(0) Bad pixel—this pixel is a known bad pixel. (1) Missing—the data for this pixel were not received from the spacecraft. (2) Despiked—this pixel was modified by the despiking routine. (3) Interpolated—this pixel has been reclaimed by interpolating from its neighbors. (4) Some saturated—the raw value for this pixel is above the point where some pixels are full well saturated. For vis instruments, this occurs at 11 000 DN, while for the IR spectrometer, this occurs at 8000 DN. (5) Most saturated—this raw value for this pixel is above the point where most pixels are full well saturated. For vis instruments, this occurs at 15 000 DN, while for the IR spectrometer, this occurs at 11 000 DN. (6) ADC saturation—the ADC was saturated for this pixel. (7) Ultracompressed—the pixel was in a compression bin so large that the value contains very little information.

For example, if the pixel is bad and has been reclaimed by interpolation, the decimal value in the quality map will be $2^0 + 2^3 = 9$. In the normal FITS format for the calibrated image, this map exists as the first image extension.

C. Signal-to-noise ratio map

In order to provide more information to the end user, the last extension of the image contains a map estimating the SNR for each pixel. The signal is taken to be the dark- and bias-subtracted image value in 14 bit DN, while the noise estimate consists of the root sum squared of three different noise sources: shot noise, read noise, and quantization noise. The shot noise in 14 bit DN is defined as

$$N_s = \sqrt{\frac{\text{Raw} - \text{Bias}}{K}},$$

where K is the gain in electrons/14 bit DN and is dependent on the instrument and mode, and raw and bias are in 14 bit DN. For the IR spectrometer, bias is 0 by definition except in mode 6. The quantization noise is defined as

TABLE XX. Noise parameters determined in ground tests of all instruments.

Instrument	$K (e/DN_{14})$	Uncompressed	
		$Q (DN_{14})$	Read noise (DN_{14})
IR unbinned	16	1	2.0
IR binned	64	1	1.0
HRI	27.4	2	0.7
MRI	27.2	2	1.0
ITS	30.5	2	1.2

$$N_q = \frac{Q}{\sqrt{12}},$$

where Q is the quantization step in 14 bit DN (N_q represents the rms error introduced by quantizing a signal whose true value could lie anywhere within the quantization bin with equal probability, i.e.,

$$\sqrt{\frac{\int_{-\Delta/2}^{\Delta/2} \epsilon^2 d\epsilon}{\Delta}},$$

where Δ is the quantization step size and ϵ is the quantization error). For uncompressed data, Q depends on the ADC performance of the instrument (see Secs. V C 7 c and VI C 8 c), while for compressed data, Q is set to the bin size in the decompression LUT that the pixel used or to the uncompressed Q value, whichever is larger. The parameter values needed for the noise calculation were determined from ground-based testing of the instruments and are shown in Table XX; the uncompressed Q values of 2 for the visible cameras are due to their less-than-perfect ADCs, as discussed in Sec. V C 7 c.

D. Spectral registration maps

In an IR image product, the second and third extensions are pixel-by-pixel maps of the spectral registration for the image. The second extension contains the effective wavelength of the pixel, while the third extension contains the spectral bandwidth. The calculations of these values are described in Sec. VI C 4.

E. Optional steps

Beyond the automated calibration pipeline described in Sec. VII A, a manual calibration can be performed where the user can specify his/her own settings and calibration files for each step. Also, any processing module can be disabled, and there are two extra ones that can be enabled. The first such module is a noise-reduction module that is applied after the despiking routine. This applies the BayesShrink wavelet thresholding algorithm³⁰ with a robust mean noise estimator³¹ to remove some of the noise. The other step that can be enabled applies a rubber-sheet geometric distortion correction. This step is not normally applied, as the optical distortion through the telescope is minimal.

ACKNOWLEDGMENTS

The work described herein was supported by the National Aeronautics and Space Administration through Con-

tract No. NASW00004 from the Discovery Program to the University of Maryland and through Task Order No. NMO710389 from the Discovery Program to the Jet Propulsion Laboratory, California Institute of Technology. We appreciate the contributions of James Baer, Ivo Busko, Lucas Kamp, Donald Lindler, Helen Mortensen, and William Owen to various aspects of this work. Figures 1–5, 56, and 71 reproduced from Ref. 3 with kind permission from Springer Science+Business Media.

APPENDIX A: DATA COMPRESSION LOOKUP TABLES PRIOR TO 1–16 days (18 JUNE 2005)

The data compression LUTs loaded into memory at the time of launch were the best estimates of the desired conversions based on prelaunch calibration data. Inflight measurements of actual instrument performance caused the science team to request updates to the LUTs. For the vis cameras, the LUT updates were driven by changes in the inflight bias levels along with a difference in the LUT number designations between what was assumed when developing the encounter command sequence and what was actually assigned in the launch load. For the IR spectrometer, the updates were driven by a realization that the IR launch LUTs were incorrectly derived and later by a need to match the LUTs to the higher background level resulting from warmer-than-expected IR bench temperatures.

Figures 102 and 103 show the LUTs for the HRI vis and MRI as loaded at launch. These LUTs applied until they were updated in flight on 26 May 2005. No inflight changes were made to the ITS LUTs; the LUTs shown in Fig. 12 were not changed during flight.

The IR spectrometer LUTs loaded at launch turned out to be incorrect. They were derived assuming that, as for the vis instruments, increasing numbers of photons on the detector produce increasingly positive signal voltages. In fact, for the IR spectrometer, increasing photon levels produce increasingly negative signal voltages. Therefore, the LUTs loaded at launch for the IR spectrometer were inverted, and Fig. 104 shows the result. Low DN levels are converted to high compressed DN's, while high DN levels are compressed to low DN values.

On 26 May 2005, a revised set of IR spectrometer LUTs was uplinked to the S/C. However, the true problem with the launch LUTs was not yet understood correctly at that time. While it was clear the IR LUTs were not behaving correctly, we mistakenly thought that the reason was that we were losing the first 2 bits of the encoded data, resulting in signal levels into the compressor being restricted to a range of 0–4095 DN. So LUTs were derived and uplinked that converted 14 bit DN's in the range of 0–4095 into 8 bit DN's in the range of 0–255. The LUT pattern was repeated three more times to fill in the full 16383 14 bit DN input table, and the sense of the LUTs remained inverted. The result is shown in Fig. 105.

After finally properly diagnosing the problem, we uploaded a corrected set of LUTs for encounter on 18 June 2005. These plots are shown in Fig. 13.




FIG. 102. HRIvis data compression LUT loaded at launch.




FIG. 103. MRI data compression LUT loaded at launch.




FIG. 104. IR spectrometer data compression LUT loaded at launch.




FIG. 105. IR spectrometer data compression LUT loaded on 26 May 2005.

APPENDIX B: ACRONYMS

ADC	= analog-to-digital conversion
AU	= astronomical unit
CCD	= charge-coupled device
CMOS	= complementary metal oxide semiconductor
CR	= cosmic ray
DI	= Deep Impact
DN	= data number
DOY	= day of year
ESD	= electrostatic discharge
FITS	= Flexible Image Transport System
FOV	= field of view
FPA	= focal plane array
FWHM	= full width at half maximum
GSE	= ground support equipment
GMT	= Greenwich Mean Time
HRI	= high resolution instrument
HRI vis	= HRI visible camera
HST	= Hubble Space Telescope
I/F	= reflected intensity/incident flux = reflectance
IFOV	= instantaneous field of view
IR	= infrared
ITS	= impactor targeting sensor
JPL	= Jet Propulsion Laboratory
L	= launch
LED	= light emitting diode
LUT	= look-up table
LVDS	= low-voltage differential signaling
LVPS	= low-voltage power supply
MRI	= medium resolution instrument
MWIR	= medium wave infrared
PDS	= Planetary Data System
POC	= parallel overlocked pixels
PSF	= point spread function
QE	= quantum efficiency
rms	= root mean square
RSS	= root sum square
SIM	= spectral imaging module
SNR	= signal-to-noise ratio
SOC	= serial overlocked pixels
STIM	= internal stimulus lamp
STSDAS	= Space Telescope Science Data Analysis System
TLM	= telemetry
TV	= thermal vacuum
UV	= ultraviolet
vis	= visible
VTC	= virtual time code

¹ W. H. Blume, *Space Sci. Rev.* **117**, 23 (2005).² M. F. A'Hearn *et al.*, *Space Sci. Rev.* **117**, 1 (2005).³ D. L. Hampton *et al.*, *Space Sci. Rev.* **117**, 43 (2005).⁴ T. L. Farnham, D. G. Schleicher, and M. F. A'Hearn, *Icarus* **147**, 180 (2000).⁵ J. T. Montroy, J. D. Garnett, S. A. Cabelli, M. Loose, A. Joshi, G. Hughes, L. Kozlowski, A. Haas, S. Wong, M. Zandian, A. Chen, J. G. Pasko, M. Farris, C. A. Cabelli, D. E. Cooper, J. M. Arias, J. Bajaj, and K. Vural, *Proc. SPIE* **4721**, 212 (2002).⁶ IAU FITS Working Group, Goddard Space Flight Center, <http://fits.gsfc.nasa.gov/iaufwg/> (2005).⁷ E. Hog, C. Fabricius, V. V. Makarov, U. Bastian, P. Schwekendiek, A. Wicencec, S. Urban, T. Corbin, and G. Wycoff, *Astron. Astrophys.* **357**, 367 (2000).

- ⁸A. S. Fruchter and R. N. Hook, *Proc. SPIE* **3164**, 120 (1997).
- ⁹H. C. Andrews and B. R. Hunt, *Digital Image Restoration* (Prentice-Hall, Englewood Cliffs, NJ, 1977).
- ¹⁰R. J. Hanisch and R. L. White (Eds), *The Restoration of HST Images and Spectra-II* (Space Telescope Science Institute, Baltimore, MD, 1993).
- ¹¹D. Lindler, I. Busko, M. F. A'Hearn, and R. L. White, *Publ. Astron. Soc. Pac.* **119**, 427 (2007).
- ¹²M. Morel and P. Magnenat, *Astron. Astrophys. Suppl. Ser.* **34**, 477 (1978).
- ¹³A. J. Pickles, *Publ. Astron. Soc. Pac.* **110**, 863 (1998).
- ¹⁴V. I. Burnashev, *Abastumanskaya Astrofiz. Obs. Bull.* **59**, 83 (1985).
- ¹⁵A. V. Kharitonov, V. M. Tereshchenko, and L. N. Knyazeva, *Spectrophotometric Catalogue of Stars* (Alma-Ata, Nauka, 1988), p. 484.
- ¹⁶N. V. Kharchenko, *Kinematics and Physics of Celestial Bodies* **17**, 409 (2001).
- ¹⁷D. G. Monet, S. E. Levine, B. Canzian *et al.*, *Astron. J.* **125**, 984 (2003).
- ¹⁸Th. Neckel, G. Klare, and M. Sarcander, *Astron. Astrophys. Suppl. Ser.* **42**, 251 (1980).
- ¹⁹See <http://redc.nrel.gov/solar/spectra/am0> for ASTM E490 solar spectrum.
- ²⁰W. M. Owen, personal communication (5/11/05).
- ²¹A. W. J. Cousins, *Mon. Notes Astron. Soc. South Afr.* **33**, 149 (1974).
- ²²G. Drolshagen, W. C. Carey, J. A. M. McDonnell, T. J. Stevenson, J. C. Mandeville, and L. Berthoud, *Adv. Space Res.* **19**, 239 (1997).
- ²³E. Deutsch, R. White, D. Lindler, and S. Ipatov, see <http://www.astro.umd.edu/~ipatov/CRcodes.html> for cosmic ray correction codes.
- ²⁴See http://pdssbn.astro.umd.edu/sbnhtml/comets/by_mission.html for the DI data archive.
- ²⁵See <http://www.iso.vilspa.esa.es/> for the NGC 7027 spectral emission lines.
- ²⁶R. R. Treffers, U. Fink, H. P. Larson, and T. N. Gautier, III, *Astrophys. J.* **209** 793 (1976).
- ²⁷D. Y. Gezari, M. Schmitz, P. S. Pitts, and J. M. Mead, *Catalog of Infrared Observations*, 3rd Edition (NASA Reference Publication 1294, Greenbelt, Maryland, 1993).
- ²⁸R. M. Cutri *et al.*, *The 2MASS All-Sky Catalog of Point Sources* (University of Massachusetts and IPAC/California Institute of Technology, Pasadena, California, 2003).
- ²⁹IRAS Catalog of Point Sources, Version 2.0, Infrared Processing and Analysis Center, California Institute of Technology, 1986.
- ³⁰S. G. Chang *et al.*, *IEEE Trans. Image Process.* **9**, 1532 (2000).
- ³¹I. Johnston and B. Silverman, *J. R. Stat. Soc. Ser. B (Methodol.)* **59**, 319 (1997).