

Imager for Mars Pathfinder (IMP) image calibration

R. J. Reid,¹ P. H. Smith,¹ M. Lemmon,¹ R. Tanner,¹ M. Burkland,¹ E. Wegryn,¹ J. Weinberg,¹ R. Marcialis,¹ D. T. Britt,¹ N. Thomas,² R. Kramm,² A. Dummel,³ D. Crowe,⁴ B. J. Bos,⁴ J. F. Bell III,⁵ P. Rueffer,⁶ F. Gliem,⁶ J. R. Johnson,⁷ J. N. Maki,⁸ K. E. Herkenhoff,⁸ and R. B. Singer⁹

Abstract. The Imager for Mars Pathfinder returned over 16,000 high-quality images from the surface of Mars. The camera was well-calibrated in the laboratory, with $<5\%$ radiometric uncertainty. The photometric properties of two radiometric targets were also measured with 3% uncertainty. Several data sets acquired during the cruise and on Mars confirm that the system operated nominally throughout the course of the mission. Image calibration algorithms were developed for landed operations to correct instrumental sources of noise and to calibrate images relative to observations of the radiometric targets. The uncertainties associated with these algorithms as well as current improvements to image calibration are discussed.

1. Introduction

The Imager for Mars Pathfinder (IMP) returned over 16,000 images from the surface of Mars from July 4, 1997, to September 28, 1997. Images acquired of the local and distant terrain permitted generation of high-resolution color, stereo, and multispectral panoramas, superresolution frames, and multispectral spots of the landing site, facilitating the geomorphologic and spectrophotometric characterization of the landing site geology. Images of the Sun, sky, stars, Phobos, and Deimos permitted modeling of aerosol scattering properties and atmospheric water vapor content. Images were also acquired of magnetic targets and windsocks on the lander, permitting measurement of the magnetic properties of the airborne dust and wind profiles at the landing site. The IMP also provided tracking and movie images of the Sojourner rover. Finally, images acquired of the two radiometric targets (RTs) permitted the initial spectrophotometric calibration of the raw image data.

The IMP was originally conceived to support what was then called the Mars Environmental Survey (MESUR) Mission. Its heritage derives from the Descent Imager/Spectral Radiometer (DISR) [Tomasko *et al.*, 1997] aboard the Huygens probe of the Cassini mission to the Saturn system. The heart of the IMP is a 512-square CCD array that rapidly transfers the image charge to the storage section in 0.5 ms (no mechanical shutter is required). The chip is then read out through a 12-bit analog-

to-digital converter (ADC) over the next 2 s. This 12-bit encoding results in a data range of 0–4095 data numbers (DN).

The IMP has a resolution of about 1 mrad/pixel and a field of view of $14.4^\circ \times 14.0^\circ$ (256×248 pixels). There are two eyes for stereoscopic imaging, with a 15-cm separation. The 12 filter positions for each eye are used for a total of 8 solar, 15 geology, and 1 diopter filter split between the two eyes. A more detailed description of the IMP is provided by Smith *et al.* [1997].

The primary requirement for mission success was that IMP provide a color panorama of the landing site. Other science goals included contour mapping of the local terrain, multispectral imaging of the surrounding rocks and soils to study local mineralogy, viewing of three windsocks, measuring atmospheric opacity and water vapor content, and estimating the magnetic properties of wind-blown dust.

In this paper we review the baseline IMP performance, expanding on the prelanding experiment report [Smith *et al.*, 1997]. We then discuss several experiments that were performed postlaunch that verify the nominal performance of the camera system. Finally, we describe the algorithms used to calibrate raw image data and review the differences between the level of calibration performed during operations versus more current and pending calibration algorithm. Note that this document is primarily intended to review IMP calibration and performance. The full IMP calibration report [Reid *et al.*, 1998] will be archived with the raw data by the Planetary Data System (PDS).

2. Laboratory Calibration

The Imager for Mars Pathfinder was extensively calibrated prior to launch at the Lunar and Planetary Laboratory (LPL) at the University of Arizona as reported by Smith *et al.* [1997]. This section supplements that work with additional calibration information acquired in the laboratory and some revision of previously published parameters. The principal raw laboratory calibration data will be made available to the community through the PDS.

2.1. CCD Properties and Dark Current

The charge-coupled device (CCD) detector for IMP was produced at Huntington Beach, California, at the former Ford

¹Lunar and Planetary Laboratory, University of Arizona, Tucson.

²Max Planck Institute for Aeronomy, Katlenburg-Lindau, Germany.

³DLR, Institute of Planetary Exploration, Berlin, Germany.

⁴Optical Sciences Center, University of Arizona, Tucson.

⁵Department of Astronomy, Cornell University, Ithaca, New York.

⁶Technical University of Braunschweig, Braunschweig, Germany.

⁷U.S. Geological Survey, Flagstaff, Arizona.

⁸Jet Propulsion Laboratory, California Institute of Technology, Pasadena.

⁹Department of Mines and Geological Engineering, University of Arizona, Tucson.

Copyright 1999 by the American Geophysical Union.

Paper number 1998JE900011.

0148-0227/99/1998JE900011\$09.00




Figure 1. The quantum efficiency of the IMP flight CCD as measured in the laboratory; replaces Figure 2 of *Smith et al. [1997]*. FWHM, full width at half maximum.

Aeroneutronics plant. The device was developed for the DISR experiment under a memorandum of agreement to the Max-Planck-Institut für Aeronomie (MPAE). The CCD is a front-side illuminated frame transfer device that uses buried channel technology with two-phase multiphase pinned (MPP) clocking. The pixel spacing was $23 \times 23 \mu\text{m}$, with a pixel size of $23 \times 17 \mu\text{m}$ and $6 \mu\text{m}$ antiblooming channel. The full-well capacity was $>1.5 \times 10^5$ electrons. The image section comprised 512×256 pixels. An additional eight columns at the edge of the image section were covered with an aluminum coating to allow measurements of the dark current for exposed frames. The optics of the IMP were constructed so that an image in each “eye” comprised 248×256 pixels with a gap (12 pixels wide) between the detector sections used for each eye [*Smith et al., 1997*]. The quantum efficiency of the CCD selected for IMP is shown in Figure 1. Note that Figure 1 replaces Figure 2 of *Smith et al. [1997]*.

The electronics were designed to provide a rapid parallel shift ($2 \mu\text{s}/\text{line}$) to terminate the exposure, yielding a total shift time of 0.5 ms. Typical exposure times were ~ 100 ms, but with shorter exposure times the excess charge (smear) introduced during the parallel shift could be significant. A 12-bit analog-to-digital converter (ADC) was used at a rate of $16 \mu\text{s}/\text{pixel}$. The system gain was set at 30.6 electrons/DN. Because the effective read noise of the CCD and pre-amplifier was between 5 and 8 el, the system noise was digitization noise limited at low intensity levels. The CCD and the associated electronics were fully described by *Kramm et al. [1998]*.

2.2. Geology Flat Field Uncertainty

The uncertainty in the flat fields is estimated by two methods: a lower limit calculation based on the combination of the shot noise in the light and dark frames and the read noise, and a comparison of two subsets of the images used to generate the flat fields. Ten frames had dark current removed using a common dark and dark shutter frame for each geology filter. The 10 calibrated frames were then averaged together to produce the flat field. The subsequent normalization and bad pixel correction do not affect the image uncertainty.

The calculation of the shot noise is based on typical DN

values in the center of the image, the typical DN values from dark current, and the read noise of the system. The typical shot noise from the image is 10 DN for a 3000-DN image. The typical shot noise from the dark current is 2.5 DN from the 70 DN of read out dark current. Since the typical exposure is much less than 1 s and the dark current rate is about 40 DN/s (at room temperature), most of the dark current results from the average 1-s readout time. The read noise is 0.5 DN. The DN contributed by the shutter effect is proportional to the shift time of 0.5 ms/exposure time, making it significant only for the shortest exposures. The calculation is done by adding the mean squared noise in each image that is used, then taking the square root, and converting to percent of the original DN value. The mean squared noise is equivalent to the number of electrons in each image. The data and the results of the calculation are summarized in Table 1. For the average of 10 frames, the mean squared noise from raw flat and shutter images were divided by 10, while the dark and dark shutter images were left the same.

As a check on this calculation, the corrected flat fields for the right eyes were averaged in two subsets of five images each, and the average difference between each pixel for these two subsets was expressed as a percent of the average. The uncertainty for a 10-frame average flat field is estimated by multiplying this uncertainty by 0.707. This estimated value is about 23% greater than the calculated values for an average of 10 frames. Both estimates of the uncertainty are significantly below the shot noise except for filter 0, where the long exposures more significant dark current increase the uncertainty. These results are shown in Figure 2.

In conclusion, the uncertainty estimated from the comparison of the subsets is reasonably close to the estimates based on the known noise sources. The uncertainty is typically much smaller than the uncertainty due to shot noise in a 3000-DN image, so flat fielding IMP images should not add significantly to the noise.

2.3. Responsivity

The absolute calibration was briefly discussed by *Smith et al. [1997]* and tables were presented showing the filter responsivities at -20°C , the typical operating temperature for IMP’s CCD detector. However, as will be shown in section 3.6, tem-

Table 1. Flat Field Uncertainties

Filter and Eye	DN in Central Area of Flat Field Frames	Exposure, s	Single Image Uncertainty, %	Flat Field Uncertainty 10-Frame Average, %
0L	2650	3.200	0.39	0.17
0R	2900	3.200	0.37	0.16
5L	3100	0.100	0.35	0.13
5R	3000	0.100	0.35	0.13
6L	3200	0.050	0.34	0.13
6R	3200	0.060	0.34	0.13
7L	2100	0.020	0.44	0.18
7R	2800	0.010	0.38	0.14
8L	3050	0.035	0.35	0.13
8R	3150	0.130	0.34	0.13
9L	3200	0.085	0.34	0.13
9R	3050	0.210	0.34	0.12
10L	3200	0.180	0.34	0.13
10R	3150	0.630	0.34	0.13
11L	3150	0.100	0.34	0.13
11R	3100	0.100	0.35	0.13




Figure 2. Right eye flat field uncertainties. Solid squares represent RMS difference between two sets of five flat fields, open triangles represent the projection of uncertainty from the average of 10 flats, and asterisks represent the RMS error calculated from 10 averages.

peratures were often not nominal during the mission. The laboratory data taken for all environmental temperatures that were expected during the mission have now been analyzed and fit with quadratic functions, having the form

$$R = A_1 + A_2T + A_3T^2 \quad (1)$$

where R is the absolute responsivity in units of $(\text{DN/s}) W^{-1} m^{-2} \mu m^{-1} s^{-1}$ and T is temperature in $^{\circ}\text{C}$. Table 2 shows the coefficients for the quadratic fits for these functions. In addition, while analyzing this data, we found that one of the data sets had been corrupted leading to an error in the L7-858.4 nm filter responsivity in Table 2 of *Smith et al.* [1997]. This error has now been corrected in Table 2.

Table 2. IMP Responsivity Parameters

Filter	λ , nm	Bandwidth, nm	A1	A2	A3
L0	443.3	26.2	128.8	-0.387	-0.0007
L1	450.3	4.91	0.246	-0.0025	-0.00001
L2	883.4	5.60	14.55	0.0233	0.00002
L3	924.9	5.03	5.389	0.0193	0.00004
L4	935.4	4.84	10.42	0.0378	0.00006
L5	671.4	19.7	575.3	-0.570	-0.0013
L6	801.6	21.0	872.2	0.237	-0.0029
L7	858.4	34.4	1435.	2.491	0.0035
L8	897.9	40.8	1120.	3.006	0.0059
L9	931.1	27.0	478.7	1.928	0.0050
L10	1002.9	29.1	213.4	1.606	0.0052
L11	968.0	31.4	395.8	2.027	0.0051
R0	443.2	26.2	117.9	-0.392	-0.0006
R1	669.8	5.30	2.238	-0.0058	-0.00002
R2	945.5	43.7	25.94	0.0780	0.00005
R3	935.6	4.91	9.738	0.0280	-0.000003
R4	988.9	5.39	1.857	0.0064	-0.000006
R5	671.2	19.5	557.3	-0.575	-0.0014
R6	752.0	18.9	787.1	-0.247	-0.0019
R7	—	—	7596.9	9.057	-0.0235
R8	599.5	21.0	592.7	-0.598	-0.0013
R9	530.8	29.6	578.6	-0.893	-0.0020
R10	479.9	27.0	368.1	-0.668	-0.0020
R11	966.8	29.6	393.5	2.185	0.0065

Table 3. Camera Model Parameters

	Left Eye	Right Eye	Error	Dimensions
Half baseline W	0.075	0.075	(0.00005)	m
Baseline to entrance pupil V	0.0623	0.0623	(0.0005)	m
Exit pupil to focal plane h	0.0234	0.0234	...	m
Pixel spacing	2.30E-5	2.30E-5	...	m/pixel
Height of optic axis above elevation axis g	0.012	0.012	(0.0005)*	m
Scaling θ	0.000981	0.000985	(0.003)	rad/pixel

Read 2.30E-5 as 2.30×10^{-5} .
*Estimated.

The uncertainties in these calibrations arise from the systematic errors in calibrating the standard lamps that we used. The values in Table 2 were obtained using a standard spectralon target illuminated by a standard lamp that is calibrated to National Institute of Standards and Technology (NIST) standards. Estimated errors are 1–2% for the lamp calibration. However, additional errors can be introduced through lamp aging and scattered light entering the camera from the walls of the room and from the interior of the camera. Our best estimate of the final absolute calibration is <5%. Values obtained at room temperature without any extra optics between the camera and the spectralon were extended to other temperatures by placing the camera inside a thermal vacuum chamber with high-quality BK-7 windows. All the measurements in the chamber were tied to the room temperature value so that reflections from the windows could not introduce errors.

2.4. Pointing

A camera model for the Imager for Mars Pathfinder [*Burkland et al.*, 1997], initially developed at NASA’s Jet Propulsion Laboratory (JPL), was used for calibration of the camera’s optical and mechanical pointing parameters at the Lunar and Planetary Laboratory of the University of Arizona.

The optical parameters (Tables 3–5) were determined by comparing laboratory images of calibrated geometric targets with the camera model predictions. The value of the IMP translation vector (the location of the base of the camera’s mast canister with respect to the lander coordinate system), reported in Table 6, was likewise determined by imaging surveyed targets during system testing at JPL and comparing the results with the camera model predictions.

Pointing curves for the measured camera motion in azimuth and elevation as a function of motor step (MS) are shown in Figures 3 and 4 for two-step incremental motion. The vertical axes represent the difference between the measured angle of rotation and the predicted angle of rotation based on the MS value. The perceived angle of rotation of the camera about either gimbal axis, determined from the position of the target

Table 4. Absolute Toe In for All Color Stereo Filters

Filter	Toe In		
	Left Eye	Right Eye	Boresight
0	12.44 ± 0.37	-24.66 ± 0.33	-0.82 ± 0.10
5	13.05 ± 0.37	-24.68 ± 0.32	-0.13 ± 0.10
11	12.55 ± 0.37	-24.45 ± 0.32	-0.13 ± 0.10

Table 5. Hardstop Values

	Hard Stop		Torque Offset	
	Stowed	Deployed	Low Hard Stop	High Hard Stop
Azimuth	240.95°	210.95°	0.146°	0.060°
Elevation	6.1°	6.1°	not known	0.075°

in the images, is not equivalent to the actual angle of rotation of the camera due to parallax. The actual angle of rotation is determined as follows:

$$\text{Angle} = \frac{\text{scale} \cdot \Delta\text{Pixel}}{\text{Parallax}(\text{Pixel}) \cdot 180^\circ/\pi}, \quad (2)$$

where ΔPixel is the shift in pixels of the target between two successive gimbal positions. Parallax is quantified as the ratio of the actual angle of rotation to the angle observed in the images as a function of pixel position. This value of parallax is determined from the camera model based on the surveyed target coordinates and the camera's optical parameters.

The difference between counterclockwise (CCW) and clockwise (CW) curves in azimuth (Figure 3a) and the up and down curves in elevation (Figure 3b) is a measure of the backlash in the gear head. For motion in azimuth the backlash is a relatively constant 1.0° for all MS. The chaotic regions, 0 to ~ 80 MS CCW and 646 to ~ 560 MS CW, are due to a nonequilibrium condition in the cable torque occurring close to the low hard stop. The backlash in elevation is nonlinear for MS 70 to ~ 235 , then reaches a constant value of $\sim 1^\circ$ for the rest of the steps. Since imaging in elevation was impossible for 0 to ~ 70 MS due to obscuration by the camera yoke, the up and down curves were set to coincide at MS 70. Pointing is predictable in azimuth for MS >80 when moving CW and for MS <550 when moving CCW, and in elevation for MS >235 .

In order to achieve an accurate prediction of camera pointing, a few rules must be observed: (1) a minimum motion of two steps, in either axis, to ensure actual camera motion (2) knowledge of the previous history of pointing is necessary for backlash correction, in either axis. In addition, the chaotic regions in azimuth (CCW: 0 to ~ 80 and CW: 646 to ~ 560 MS) will have larger pointing uncertainty. Images that were not acquired in adherence to these rules may have significant pointing uncertainty.

2.5. Point Spread Function

The IMP's modulation transfer function (MTF) was measured on the ground for filters 0, 5, and 11. One-dimensional MTF profiles were obtained in three directions relative to the CCD array; parallel to the antiblooming gates, perpendicular to the antiblooming gates, and at 45° relative to the antiblooming gates. For an exterior reference the camera mast is perpendicular to the antiblooming gates (the antiblooming gates

are roughly parallel to the Martian horizon). The MTF profiles were taken at three different longitudinal distances for filters 5 and 11 and at two distances for filter 0.

IMP data users should be aware that our testing showed the IMP system will pass, with measurable amplitude, spatial frequencies higher than the array Nyquist frequency. An example of this can be seen in Figure 4.

Since our measurements resulted in only one-dimensional slices of the optical transfer function (OTF), we needed a method of producing the two-dimensional, system point spread function (PSF) at all the wavelengths of interest. To do this, we modeled the lens system as being diffraction limited, suffering only from defocus, and modeled the array as a $17 \times 23 \mu\text{m}$ rectangle function. Since the lens system was a 23-mm Cooke triplet operated at approximately $f/18$, modeling the lens as diffraction limited was very reasonable. Our diffraction model followed *Gaskill* [1978] and *Goodman* [1968].

To compare our model with the MTF data, we took the convolution of the lens diffraction PSF, the array PSF, and a slit to obtain line spread functions. The model gave good agreement with the data that we took at the short wavelength. Unfortunately, for the longer two wavelengths the model was found to be inadequate.

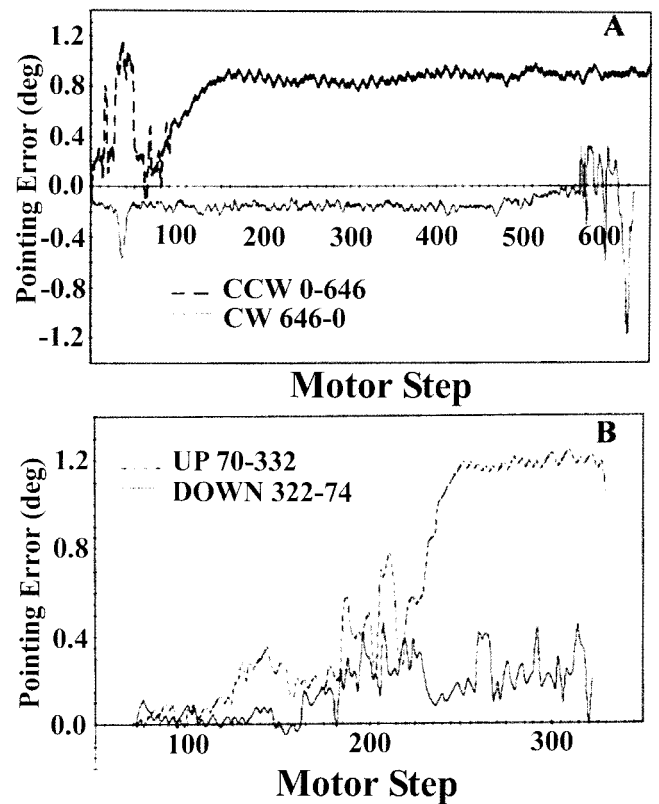


Figure 3. (a) Pointing error in azimuth as a function of motor steps (MS). The vertical displacement between the counterclockwise motion (CCW) and the clockwise motion (CW) is the measure of backlash, about 1.0° over most of the range of motion. (b) Pointing error in elevation as a function of MS. The data range begins at 70 MS due to visual impairment of the camera by the azimuth outer housing. The vertical displacement between the up and down motion is a measure of the backlash. The backlash is a complicated function of MS for 70 to ~ 235 MS, then becoming relatively constant ($\sim 1^\circ$) for the remaining steps.

Table 6. IMP Translation Vector

	Corrected*	Initial	Difference
x	0.2030	0.1984	0.0046
y	0.0134	0.0090	0.0044
z	0.4431	0.4367	0.0064

In meters.

*Corrected values are determined from model fit.

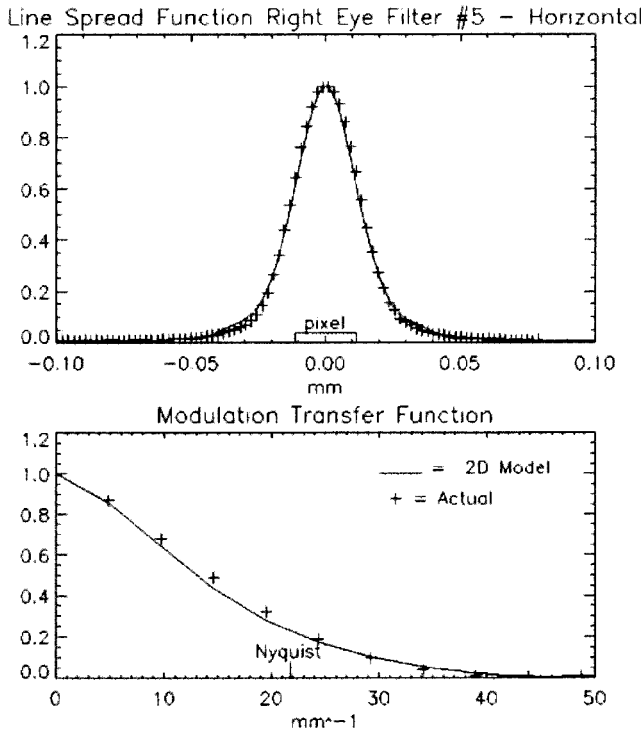


Figure 4. Comparison of the measured line spread function and MTF of the IMP with our diffraction and crosstalk array model at 673 nm. The object to camera distance was 0.5 m.

We believe that significant pixel cross talk at the longer wavelengths was responsible for the disagreement between the data and the ideal model. Pixel cross talk effects have been seen in solid-state imagers since soon after their introduction, and they are expected to become more pronounced as the size of pixels goes down [Seib, 1974]. One of the causes of pixel cross talk, which is strongly dependent on wavelength, is the diffusion spreading of minority carriers through the bulk semiconductor.

To account for this diffusion of minority carriers, we replaced our ideal CCD array response with a more realistic

model, consisting of 9 pixels. The center pixel has a normalized response of 1.0, and the 8 remaining pixels have a lower response that depends on the amount of crosstalk. The pixel pitch is $23 \mu\text{m}$, and the pixels are $17 \times 23 \mu\text{m}$ rectangle functions. The gaps in the pixels are caused by $6\text{-}\mu\text{m}$ -wide antiblooming gates. The square root of two terms in the corner pixels is simply due to the greater distance of the corner pixels from the center pixel than the others. The 0.7 term is a value we found to agree well with the data to account for the barrier provided by the antiblooming gates.

This new array model convolved with the diffraction pattern of the lens agreed quite nicely with our measurements at the two longer wavelengths. The required amount of cross talk for agreement between our model and the data is also reasonable according to the literature [Lavine *et al.*, 1985]. Our model also accounts well for the change in system PSF for different object distances. The amount of cross talk is dependent on wavelength. It is zero for wavelengths below 574 nm and increases linearly for wavelengths above 574 nm.

2.6. Calibration Target Photometry

The radiometric targets (RTs) and color targets (CTs) (Figure 5) were designed to survive the harsh environments of cruise and the Martian surface, but still to emulate, as closely as possible, Lambert surfaces with known reflectances. Deviations from Lambertian, however, require that the reflectance properties of the targets be carefully measured. To evaluate the scattering behavior of the targets, we used the Gonio Photometry Laboratory (GPL) equipment at the DLR Institute of Planetary Exploration (Berlin, Germany).

The GPL is divided into three parts: the lamp unit, the sample holder, and the detector unit. The units are combined by rotating arms with stepping motors and are mounted on a platform. The entire instrument was placed in a temperature-stabilized room ($\pm 1^\circ\text{C}$) with near-clean room conditions. The incident beam was produced by a 250-W quartz-halogen lamp and had a homogeneity footprint in the sample plane of $\pm 2.5\%$ within a 40 mm radius. The slit width of the monochromator was 5 mm resulting in a spectral resolution of 14 nm, similar to the bandwidth of the IMP. A photomultiplier was used to




Figure 5. Black and white image of lower RT acquired on sol 55.




Figure 6. DLR target measurement geometry. The azimuth angle (Φ) is defined as the angle between the scattering plane and the incidence plane. Both planes are defined as the plane that includes the surface normal and the incidence/emergence direction. The symbols are i for the incidence angle, e for the emergence angle, and ϕ for the azimuth angle. The sign indicates the left or right quadrant of the incidence plane.

detect the light in the wavelength range from 400 nm up to 1000 nm in steps of 20 nm.

To investigate the RTs, we used two fixed emergence angles 26° and 47° , which is similar to the geometry of the upper and lower RTs on the lander. The incidence angle varies between $\pm 50^\circ$ in steps of 10° at the following azimuth positions (defined in Figure 6): $0^\circ, 180^\circ; 60^\circ, 240^\circ; 140^\circ, 320^\circ$. All known systematic and statistical errors of the equipment lead to a maximum error over the entire wavelength range of $\pm 3\%$.

The quantity that is generally measured in the laboratory is the reflectance coefficient (r_c) [Hapke, 1981]. This parameter is defined as the ratio between the radiance of the sample measured at fixed geometry to the radiance of a Lambert surface with identical illumination. As a near-Lambertian reference standard, we used Spectralon (the trade name of polytetrafluorethylene (PTFE)).

We can calculate the reflectance coefficient from the measurements in the following way:

$$r_c(\text{sample: } \lambda, i, e, \phi) = \frac{I(\text{sample: } \lambda, i, e, \phi)}{I(\text{Std: } \lambda, i_{\text{ref}}, e_{\text{ref}}, \phi)} \cdot r_c(\text{Std: } \lambda, i_{\text{ref}}, e_{\text{ref}}, \phi) \quad (3)$$

The white RT reflectance data show a strong opposition surge in the principal plane, which is characteristic of low




Figure 7. Reflectance characteristics of the white calibration target (WB3) at $e = 26^\circ$. A strong opposition surge ($i = 26^\circ$) and a specular reflection ($i = -26^\circ$) are evident.




Figure 8. Reflectance characteristics of the gray calibration target (GB3) at $e = 26^\circ$. The opposition surge ($i = 26^\circ$) is smaller compared to the specular rise ($i = -26^\circ$) than for the WB3 target.

absorption materials (Figure 7). A broad rise in reflectance is seen near the specular geometry for $e = 26^\circ$. This behavior is similar to specular reflections of playas discussed by *Shepard et al.* [1993]. Indeed, microscopic images of the RT surfaces reveal that the RT surfaces resemble a playa surface.

In Figure 8 the reflectance coefficient of the gray RT shows similar behavior as the white RT. The strong opposition surge is smaller compared to the specular rise than for the white target.

The black RT exhibits a very small opposition effect but a large specular spike (Figure 9). This is the dominant effect in the reflectance characteristics of the black RT and doubles the reflectance when viewed in the specular geometry.

In all these measurements we can observe that the RTs have, as a result of the surface roughness diffusing the specular reflection, a broad specular lobe of $\sim \pm 30^\circ$ in width, which is most pronounced for the black target and least pronounced for the white target. All laboratory measurements shows that specular reflection is a significant factor in the reflectance characteristic of the radiometric and color targets.

2.7. Compression

Since the data transmission rate was limited, some image compression was required for all image data. Two kinds of compression were used, lossless, with no degradation of image quality, and lossy, where some image information was sacrificed to improve the compression ratio.

It is well known that a block-based discrete cosine transform (DCT) produces high levels of aliasing. At high compression ratios, aliasing will be removed poorly during reconstruction and becomes visible at block boundaries (blocking artifacts). Blocking artifacts can be minimized using a block overlapping transform. Therefore two block overlapping transform coders and a non-block-oriented wavelet coder have been considered for IMP image data compression: (1) a task-oriented modification of the widely used sequential DCT based mode of the JPEG standard, enhanced by local cosine transform (LCT) prefiltering [*Aharoni et al.*, 1993], (2) a JPEG-like scheme using lapped orthogonal transform (LOT) [*Malvar*, 1992], and (3) an embedded zero-tree wavelet coder.

Controlled output rate was a principal IMP requirement. Output rate control is favored by wavelet image data coding (option 3) in two or more passes. In contrast, the one pass transform coders (options 1 and 2) execute output rate control via an iterative search process. After the block-based individual DCT of the image is finished, the subsequent compression steps are iterated during a binary search algorithm in order to fix that quantization level which produces the best fit with the commanded output rate.

Option 1 was selected. The decisive argument was that this option is characterized by an add on to the proven JPEG standard. The additional prefilter function can be switched off at lower compression ratios and in case of operational problems. Then, standard JPEG is the trusted fallback position.




Figure 9. Reflectance characteristic of the black calibration target (BB) at $e = 26^\circ$. The opposition surge ($i = 26^\circ$) has nearly vanished, while specular reflection ($i = -26^\circ$) can double the radiance coefficient value at some geometries.

IMP lossless image data compression was realized following the well-known Rice algorithm [Rice, 1991].

3. Verification of Laboratory Calibration Postlaunch

In order to verify that the IMP system survived launch, cruise, and landing and that it continued to perform according to its specifications, several data sets were obtained and analyzed throughout the course of the mission. Several of these verification tests are discussed in more detail by Reid *et al.* [1998] (available at http://imp.lpl.arizona.edu/imp_Team/report/).

3.1. Bias Voltage and CCD Substrate Grounding

The electronics were designed so that an additional voltage was applied to the signal from the detector prior to analog to digital conversion. This additional voltage is normally called the bias, dc offset, or hardware offset and is necessary to prevent underflow of the ADC. The bias was set preflight to a value equivalent to 8.7 DN. The serial register contained four extra pixels at the beginning of each line that monitored the serial register charge in addition to the dc offset applied to the ADC. These extra pixels were separately stored in the so-called "null strip" data frames containing 4×256 pixels each. The first data returned by the IMP from the surface of Mars showed that the null strip was very close to 8.7 DN and that

operation of the device was nominal. However, later in the mission, some changes became evident.

Loral constructed the substrate grounding of the CCD by an aluminum metallization on the back side of the detector die. During the evaluation of several detectors in the MPAE laboratory, an extra offset above the bias was found. The effect could be explained by poor substrate grounding caused by oxidation of the aluminum surface and the subsequent failure of the bond between the aluminum and the ground. Poor substrate grounding causes an exponentially decreasing offset amplitude upon the first pixels in each line. As a consequence, in 1994, Loral modified their process to use gold metallization on the back side. However, these devices could not be produced in time to supply the IMP camera, and therefore the IMP has a detector with aluminum metallization.

A total of 91 null strip frames were received. The first 18 frames (obtained between sols 1 and 3) give clear indications that the substrate grounding was perfect. However, all remaining null strip data received between sols 12 and 83 contain an extra offset that suggests that substrate grounding degraded. While there are no null strips available from sols 3 to 11, the most likely event to have resulted in this small change in electronics behavior was the deployment of the IMP on the evening of sol 2.

After sol 12 the first pixel in each line is truncated to 0 DN by negative offset. The offset amplitude on the next three




Figure 10. The readout region of the IMP CCD, showing the null pixels (columns 0–3), the dark strip (columns 4–11), and a portion of the imaging section (columns 12–21). The staircase effect is due to charge accumulation during CCD readout.

pixels is exponentially decreasing from +0.5 DN to less than +0.3 DN so that only minor offset effects can be expected on pixels from the image. The hardware offset value from sol 3 onward is now calculated to be 8.3 DN. This effect remains stable during the rest of the mission.

3.2. Dark Current

Because the readout region of the IMP CCD serves as the pixel gateway for the entire chip (through which all electrons must pass), we monitored this region during operations as part of the camera health check. The readout region includes 4 “null” pixels, discussed above, an 8-pixel wide “dark strip,” and the neighboring image area. The shaded 8×256 pixel dark strip is at the physical edge of the chip and provides a measurement of the CCD dark current.

Figure 10 shows the a section of the readout region of the IMP CCD, acquired during a zero second exposure. The 4 null pixels are shown on the left edge of the image. To the right of the null pixels are the 8-pixel-wide dark strip and a 9-pixel-wide portion of the CCD imaging section. The pixel to pixel variation in the null “region” is due to CCD readout noise. The staircase effect seen on the dark strip and image region is due to the accumulation of dark current as the charge sits in the memory section of the CCD during readout. The pixels nearest to the horizontal shift register (the first row in Figure 10)

contain the smallest amount of dark current, while the pixels spending the most time on the chip (those in the back row) have the highest amount of dark current. This charge accumulation effect can be removed from an image by subtracting a zero second exposure from an exposed image, called shutter subtraction (or correction). The example shown in Figure 10 was chosen because the dark current was relatively high ($T = 16.2^\circ\text{C}$); during surface operations this effect was much smaller.

Dark images were obtained twice during the cruise phase to Mars. During check-out (December 16, 1996, 12 days post-launch), full-frame dark images from both eyes were obtained at 17°C . It was noted at this time that the detector appeared to show a large number of warm spots during this period. Data acquired at a lower temperature on June 20, 1997, showed similar effects, although only subframes of the detector were transmitted to ground.

On the surface of Mars a total of 9 dark frames from each eye were obtained. The data taking was scheduled for the early afternoon when the camera head was expected to be at its warmest. Consequently, the dark frames covered a temperature range of -15°C to -7°C . Dark patterns obtained on different sols are identical, indicating that the new warm spots are reproducible and are not the direct result of energetic particle impacts during the exposure and/or readout.

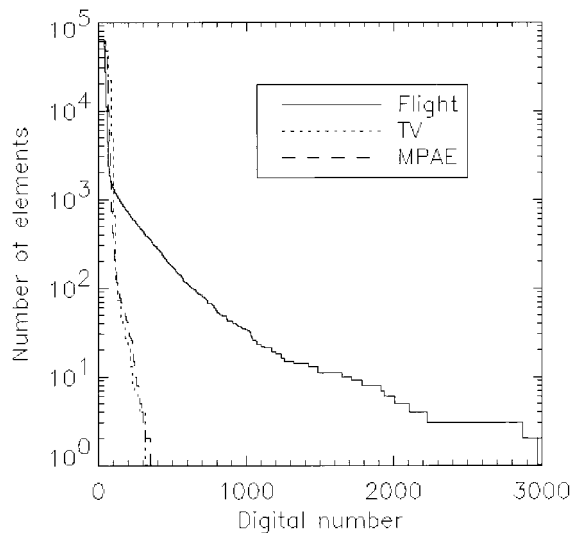


Figure 11. The number of elements in an array greater than a specified data number. This type of plot allows one to see instantly a change in the number of warm spots. Note that the flight data (solid line) has more than 1000 more warm spots than observed during MPAE CCD selection testing and during the system thermal vacuum test.

The investigation of the energetic particle flux is complicated by the limited number of images available and the wide range of temperatures that they cover. A hot spot will provide a large jump in intensity over a relatively small temperature range, and therefore for frames taken over a range of temperatures a method must be devised to separate hot spots from energetic particle events. The temperatures of the two frames used (-9.2°C and -9.8°C) are similar, and therefore large jumps in the data numbers are not expected. Assuming that dark current doubles every 7°C , the difference in temperature between the two frames should produce a 9% difference in the dark level. An energetic particle event was suspected if the difference between the two frames exceeded 3 times this value (27%). To ensure that hot spots were eliminated, pixels with a level 3 times greater than the detector mean were correlated with the suspected events and eliminated from the suspect list. The two frames produced an average of 74 affected pixels within the 256×248 pixel array in a 32.7675-s exposure. (It should be noted that the number of events does not scale directly with the exposure time because of the additional time the image rests on the detector during the 2-s readout.) We estimate an error bar on this value of 15% from varying the thresholds in our automated procedure. It has been previously estimated that the detector showed three events per second during the cruise phase of the mission (D. Crowe, personal communication, 1997). Our results are in reasonable agreement with this value. This, in turn, gives a maximum value for the flux of energetic particles of $0.78 (\pm 0.12)$ particles $\text{cm}^{-2} \text{s}^{-1}$. This is an upper limit because one particle may (and probably will) affect more than one pixel. For comparison, a detailed analysis of the CCD in the Halley Multicolour Camera recorded 1.75 particles $\text{cm}^{-2} \text{s}^{-1}$ during the Giotto flyby of comet Halley in 1986 at a heliocentric distance of 0.89 AU [Kramm et al., 1993].

The evaluation of the numbers of energetic particle events inflight allows us to make a direct comparison of the dark matrices produced during on ground testing and the flight

data. Energetic particles were removed from the flight data following the above procedure. The data were then compared with CCD selection tests conducted at the Max-Planck-Institut für Aeronomie (MPAE) in November 1994 and with system level thermal vacuum test data in October 1995 (14 months prior to launch). We have selected data obtained at approximately -8°C with a 2-s integration and scaled the data to simulate a 32.7675-s integration.

Figure 11 gives the number of elements in an image with a digital number greater than that specified by the ordinate. The width of the peak at low digital numbers is an indication of the temperature of the detector. It can be seen in Figure 11 that the thermal vacuum test data (TV) shows a slightly broad peak that the other two data sets indicating that the TV data were obtained at a slightly higher temperature (as expected). Both the TV data and the MPAE test data show the device to be very clean and the two plots track each other well. On the other hand, the flight data show a significant number of pixels with additional charge, which clearly lead to the new warm spot pattern. Figure 11 shows a breakpoint in the distribution at around 85 DN. Around 1500 pixels have data values greater than this, corresponding to 2.4% of the detector. We note that the number of affected pixels is 20 times the number expected from energetic particle events. The mean data value of these pixels is 270 DN or $0.2 \text{ DN s}^{-1} \text{ pixel}^{-1}$ averaged over the device.

Typical exposure times on Mars were of the order of 200 ms. The additional charge would then be around 2 DN. Thus, for typical exposures of the daytime surface through the geological filters this dark current pattern would be insignificant for signal levels greater than about 150 DN. Long exposures (e.g., through the blue filters) of shadowed areas or during twilight may, however, require calibration.

3.3. Flat Fields

To examine the stability of the flat field correction, several images of the sky were examined before and after flat field correction was applied. For this study, several images containing the horizon were used that were either losslessly compressed or had a low compression ratio of 2:1. An example of images used for this study is shown in Figure 12. Images used were acquired during superpan sequences taken late in the mission on sol 75, such that any degradation of the calibration due to exposure to the Martian environment should be detectable.

Visual analysis indicates that the raw images show evidence for the crosshatched pattern characteristic of the manufacturing process of the CCD. The design of this system also creates a falloff in intensity that goes as the \cos^4 of the angle away from the optical axis, which is also evident in the data. The magnitude of the falloff and the pixel-to-pixel variation can be seen in Figure 13 for the R5 filter (671 nm), showing a pixel-to-pixel variation of $>1\%$ and a falloff of $\sim 5\%$ for this particular row. Removal of the flat field using the instrument correction procedure discussed below shows significant improvement in the visual quality of the images. The profile through the same row of the corrected image shows that both the low- and high-frequency variations have been reduced. In general, the \cos^4 falloff has been completely removed, with additional low-frequency structure representing actual intensity variations in the sky. The pixel-to-pixel variation has been reduced to $\sim 0.4\%$, representing a signal-to-noise (SNR) of 250:1, which compares favorably to the maximum system SNR of 300:1 from




Figure 12. R5 (671 nm) lossless superpan image of horizon on sol 75 (i12533192351.img_0183070036). Left image is raw data enhanced to show flat field cross-hatching and \cos^4 falloff. Right image is data corrected for dark, dark shutter, and flat field by CCDCAL v1. Note removal of cross-hatching and \cos^4 falloff. Spots above horizon on left side of both images are indications of dust accumulation on IMP windows.

Poisson noise at comparable exposures of ~ 3000 DN. Results are comparable for all filters.

The low responses in the first ~ 15 columns require further study. Since the toe-in of the IMP has not changed within measurable limits (M. Lemmon, personal communication, 1998), this effect is likely due to vignetting by internal baffles, rather than a change in the optical axis. Other effects of the prolonged exposure to the Martian surface can also be seen. There are two spots on just above the horizon on the left side of the image that are visible in both the raw and CCD-corrected images. These spots are out-of-focus dust accumulating on the front windows of the IMP.

3.4. Relative Responsivity

To examine the stability of system responsivity, stereo images were compared for consistency in response, and several images were calibrated using independent reference standards to study the band-to-band stability of the responsivity.

Tests were done using images containing the horizon, as was

done for flat fields. The first tests measured left to right eye stability by comparing the absolute radiance of a patch of sky in both eyes in the stereo filters, using the laboratory values for absolute system response. Since the same patch of sky was imaged simultaneously by both eyes, these measurements will observe identical intensities. Results can be seen in Figure 14. There is excellent ($<1\%$) agreement in the IR stereo filters and good (3%) agreement in the red stereo filters. Disagreement in the blue stereo filters is significant, however, with a disagreement of $\sim 20\%$.

Examination of the stability of spectral response of the IMP was also facilitated by observations of the same images used for the flat field analysis. In this case, measurements of the radiance of the sky just above the horizon were calibrated with two independent reference standards. The data were normalized by the unattenuated solar flux at the top of the Martian atmosphere and normalized by the measured intensities of RTs acquired concurrently with that imaging sequence. As can be seen from Figure 15, both methods of calibration produce




Figure 13. Demonstration of removal of flat field from raw image. Dotted line shows profile along row 220 of raw image, with large pixel-to-pixel variation and \cos^4 falloff evident. Solid line shows same profile across corrected image, showing reduction in flat field effects. Residual drop-off in first few columns is being studied.




Figure 14. Intensity of sky just above horizon in sol 75 superpan sequence 183. Good agreement in red (671 nm) and IR (967 nm) stereo filters indicates little change occurring between the two eyes. Significant disagreement ($\sim 20\%$) in the blue stereo filters (443 nm) suggests a change in responsivity, possibly due to degradation of the R0 filter.




Figure 15. Verification of responsivity by using independent calibration methods. Normalization of the sky intensity by the solar flux at the top of the atmosphere provides a sky spectrum that includes the effects of responsivity variations (dotted line). Normalization of the sky intensity by the target brightness produces a sky spectrum that is independent of responsivity (solid line). The agreement in shape of the two lines suggests that the responsivity of the instrument has remained stable, except perhaps for the higher of the two blue stereo filters (R0).

similar results. Since any changes in responsivity would be indicated by a severe difference in the shapes of these lines, the similarity between them indicates that no major changes in spectral response have occurred since laboratory calibration,

with the exception of one of the 440-nm stereo filters. The high apparent value in the 480-nm filter is seen in both the target-corrected and solar-flux-corrected curves. This indicates that this spectral feature is a real property of the atmosphere and not due to changes in the responsivity of the IMP.

By examining the shapes of these curves, it is evident that the discrepancy between the blue stereo filters is due to a change in responsivity of the right eye at this wavelength. Since this effect is not seen in the other filters, it is unlikely that this difference is due to changes in the CCD. The evidence suggests that a change in the R0 filter is the cause of the apparent change in responsivity. The nature of this discrepancy will be investigated in future work.

3.5. Radiometric Color Targets Cross-Check

Several sets of color targets were located near both radiometric targets and the magnetic arrays. These targets were composed of pigments cast in silicone binders in five colors including the iron oxides hematite (red), maghemite (brown), and goethite (yellow) as well as chromite (green) and cobalt (blue) commercial paint pigments.

Images containing the CTs were calibrated with observations of the RTs to provide a check on the ability of IMP to discern spectral features. While the apparent lack of strong spectral features in most of the Martian scene raised concerns about IMP's capabilities, spectra extracted from the calibrated CT images show that the visible and near-infrared spectral features associated with each color chip replicate those in the laboratory spectra (Figure 16). This indicates that spectral features are discernable using IMP.

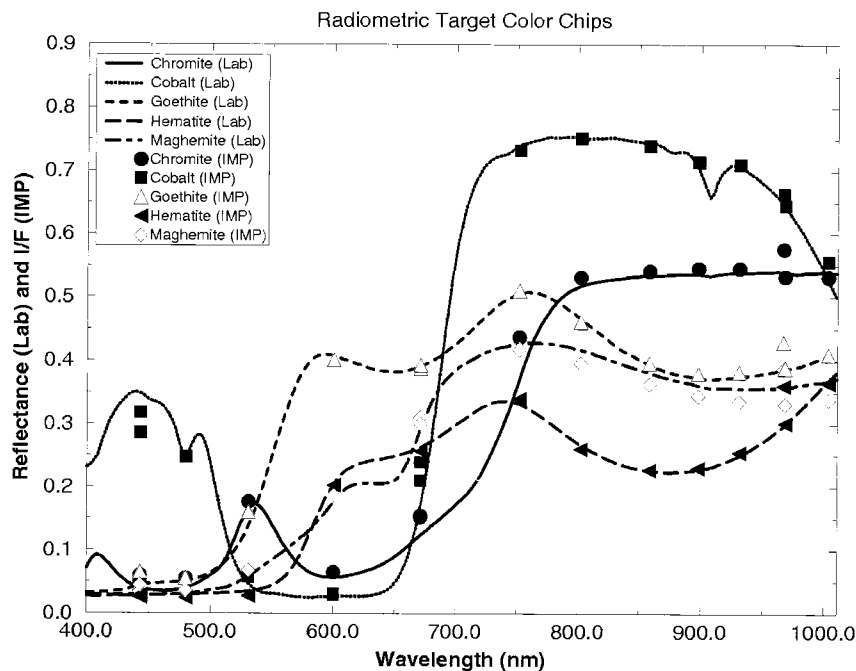


Figure 16. Laboratory spectra of color targets (lines) compared to spectra extracted from calibrated IMP images of the targets on Mars (symbols). Right eye points were obtained from lower RT (sol 1, 0950 LST); left eye points were obtained from upper RT (sol 3, 1000 LST). Because the dark side of the hematite color chip was measured in the laboratory, while the brighter side was placed on spacecraft RTs, IMP-derived hematite points have been multiplicatively scaled to match lab reflectance to compensate for differences in albedo.

3.6. CCD Heater Performance

A CCD heater was included on a tab on the back of the IMP CCD in order to maintain a CCD temperature of -20°C , primarily to ensure consistent dark current in atmospheric water vapor measurements. The heater was simply a 1-W resistor bonded to a tab on the rear of the CCD package, with a thermistor mounted on the opposite side of the tab for temperature control. In addition, another thermistor mounted on the optical bench reported temperature that closely follows the ambient temperature of the IMP. The control algorithm simply turned on the heater if the CCD temperature was lower than the control temperature and turned off the heater if the CCD temperature was greater than the control temperature. The CCD temperature algorithm is activated once every 30 s.

A typical IMP temperature cycle is shown in Figure 17 for sol 13. The temperature reaches a minimum near -80°C shortly before sunrise. When the heater is turned on in the morning, it cannot initially provide enough heat to maintain a control temperature of -20°C . Rather, it provides $\sim 40^{\circ}\text{C}$ of temperature boost. Later, as the ambient temperature warms to above -60°C , the heater is able to maintain the control temperature. As the temperature warms above the control temperature, the heater is turned off and the CCD temperature once again follows the ambient optical bench temperature.

Two effects cause the measured variability around the control temperature. The first is the coarseness of the control of the heater; temperature measurements were reported in 30-s intervals, during which time the CCD temperature could overshoot the control temperature. Also, due to the proximity of the thermistor to the heater, the fluctuations are more severe than those on the front of the CCD. These factors suggest that the actual CCD temperature is much more stable when being controlled, within a few degrees of the average. In addition, due to the thermal path to the front of the CCD and the proximity of the sensor to the heater rather than the front of the CCD, the average temperature is slightly lower than the control temperature. In general, when the temperature is being controlled, the actual CCD temperature is approximately




Figure 17. IMP temperature cycle on sol 13. Optical bench temperature (dotted line) is ambient temperature, while CCD temperature (solid line) shows effect of CCD heater when the heater is on. Note that CCD heater can provide a maximum of $\sim 40^{\circ}\text{C}$ temperature boost.

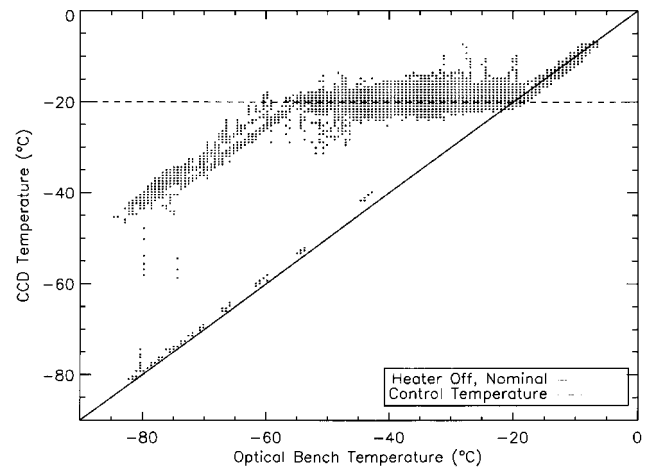


Figure 18. IMP temperature measurements for all 16,000+ images acquired. CCD and optical bench temperatures follow each other when CCD heater is off (solid line). Data demonstrate scatter of temperatures about control temperature (dashed line) and inability of CCD heater to provide more than $\sim 40^{\circ}\text{C}$ of temperature boost.

-21°C , rather than the CCD temperature value reported in the header.

Another way of examining the performance of the CCD heater is to plot the CCD temperature versus the optical bench temperature as in Figure 18. Temperature data reported for all images are shown in Figure 18. Figure 18 demonstrates the ability of the heater to provide as much as 40°C of temperature increase and shows that both CCD and optical bench temperatures are in close agreement when the heater is off.

3.7. Pointing

Initial observations of atmospheric images suggested a discrepancy between calculated and actual IMP pointing. Images intended to be centered on the Sun were not, and images of the solar aureole also showed a discrepancy in pointing, specifically that the brightness contours in the aureole images would not line up with contours in sky brightness models. Since the position of the Sun in the sky can be calculated with great accuracy for a given time at which an image is taken, this discrepancy indicates a significant source of error in the header azimuth and elevation values.

However, once some systematic errors in the initial ground processing of the raw data were accounted for, it was found that relative pointing accuracy on Mars was similar to that measured during calibration on Earth. Each of the 42 observations of the upper RT through the left eye filter 5 was examined. A fiducial point in the images was chosen, and its position was noted for each of the images. In azimuth the position of the fiducial had a standard deviation of 4.3 pixels around its mean position. In elevation the position of the fiducial had a standard deviation was 4.5. In nearly all of these images the pointing command and the pointing history were the same. Typically, backlash causes pointing uncertainties to be higher.

Absolute pointing was measured first by a sequence imaging six fiducial points on the lander frame and second by imaging the Sun. The fiducial sequence (S0025) obtained stereo images aimed at the six fiducial points. The high gain antenna blocked the view of one point, and various cables blocked the view of

several others. In addition, during the system test, different pointing commands were used, and the high gain antenna blocked the view of a different fiducial. Of the three best determined measurements the camera model determination of the fiducial location was off by -4 , 4 , and 9 pixels in azimuth and -15 , -9 , and -1 pixels in elevation. During the sequence execution the elevation errors were moved in small amounts repeatedly, causing the backlash to be unpredictable of the order of 1° (18 pixels).

The camera imaged the Sun over 2000 times. The bulk of these observations were atmospheric science images using the ATMOS flag. In this setup a 31×31 pixel subframe centered on the Sun was downlinked. Unfortunately, the image header information does not indicate the location of the subframe on the CCD. The position of the Sun on the CCD was recorded in the engineering data record, where it was overwritten with each new image, but 500 distinct position were downlinked during the mission. These Sun positions were associated with images for which the IMP camera model had been used to predict the Sun position on the CCD. The measured deviations (observation minus model) are $8.3 + -0.5$ pixels in x (azimuth) and $4.6 + -0.5$ pixels in y (elevation). The standard deviations of the distribution were 8.3 and 10.4 in x and y respectively. In addition to these images, images of the Sun in a known subframe were taken on five occasions. The errors (observation minus model) for these images fell within the above distribution. Currently, errors in the spacecraft event time determination are being investigated. Timing errors, an error of 0.5° in spacecraft orientation, or an error in the camera model could explain the deviations.

4. Image Calibration Algorithms

Several image calibration pipelines were developed in order to facilitate rapid processing of raw image data for immediate scientific analysis during the landed mission. The first versions of these pipelines have been used to calibrate most IMP data acquired during the mission, on which many preliminary scientific reports are based [cf. J. F. Bell et al., unpublished manuscript, 1998; Johnson et al., this issue]. There are two pipelines, one to perform basic instrument level corrections, such as bad pixel replacement, dark current subtraction, and flat fielding; and one to facilitate calibration of images to reflectance relative to measurements of the RTs. An improved instrument-level correction algorithm has been developed that reduces the uncertainties associated with the algorithm used during the landed mission, which will be archived with the IMP data set by the PDS [Reid et al., 1998].

4.1. Flight Software Corrections

The level of flight software (FSW) correction varies from image to image. Generally, bad pixel correction was applied to every image that was compressed, and electronic shutter correction was applied to all subframed images but not to full-frame images. A small number of images have had both FSW dark current and flat field corrections applied. Owing to limitations of the small amount of nonvolatile memory on board the spacecraft, the FSW versions of these corrections were less accurate than those that could be applied on the ground. Since downlink rates were higher than expected, the small increase in compression efficiency afforded by applying these corrections in FSW was outweighed by the larger benefit of performing these operations on the ground. Owing to the large number of

individuals designing imaging sequences, there are deviations from the above generalizations that should be checked by verifying FSW processing in the raw image headers.

Subframing and pixel averaging were also performed by FSW. FSW pixel averaging first smoothed the subframed image data horizontally and then vertically with a linear filter. All pixels within the block were then averaged, and the average value was returned. Images were restored on the ground by using bilinear interpolation.

Lossless and lossy compression algorithms were also applied to the raw data. Lossless compression clearly adds no additional noise to the signal, but lossy JPEG compression removes high spatial frequency information from the image and adds noise in a way that is dependent upon both the scene and compression factor used. It is impossible to remove the noise generated by lossy compression.

In addition, whenever FSW corrections were applied to the data (except for bad pixel corrections), a software offset of 16.0 DN was added to the image to prevent the signal from underflowing during processing.

4.2. Ground Instrument Corrections

The ground instrument calibration software is intended to correct the major sources of noise omitted by FSW processing and return images in units of radiance ($\text{W m}^{-2} \mu\text{m}^{-1} \text{s}^{-1}$). Full correction of instrumental sources of noise requires the generation of a bias and a scaling frame. The bias frame removes sources of noise that are independent of illumination of the detector, that is, DN due to thermally generated electrons and the hardware offset, and the scaling frame removes pixel-to-pixel variations in the responsivity of the system. Both frames are then subframed and pixel-averaged identically to the raw image they are meant to correct.

The bias frame is generated by using the laboratory-derived dark model [Reid et al., 1998]. This frame includes signal from thermally generated electrons and the hardware offset that is included in the raw data. The dark model has the form

$$\text{DN} = A_d t e^{B_d T} D(x, y) + A_s e^{B_s T} S(x, y) + A_n e^{B_n T} + H_{\text{off}}, \quad (4)$$

where $A_d = 3.016$, $B_d = 0.105$, $A_s = 2.845$, $B_s = 0.105$, $A_n = 4.05$, $B_n = 0.144$, $H_{\text{off}} = 8.27$, and T is temperature in $^\circ\text{C}$.

The first term models the dark current generated in the active region of the CCD during an exposure of t seconds, with the term $D(x, y)$ representing the normalized laboratory-measured dark pattern. The second term models the dark current generated during readout of the image from the storage region of the CCD, with the term $S(x, y)$ representing the normalized laboratory-measured dark shutter pattern. These frames will be archived by the PDS. The third term models the temperature-dependent null pixel offset, and the fourth term represents the constant hardware offset. In addition, a software offset of 16.0 DN must be removed from the raw data if any FSW processing (other than bad pixel replacement) was performed on the data.

Not every image will require that all of the above terms be used in the generation of the bias frame, since some FSW corrections will typically have been applied to each image. The level of FSW calibration must be verified in each raw image that is to be calibrated so that the proper terms can be included.

The scaling frame is simply the normalized, bad pixel cor-

rected flat field response pattern for a given eye and filter. These patterns were measured as discussed above in laboratory calibration and will be archived by the PDS.

In addition, the smear induced by electronic shuttering will be removed for images that contain the row adjacent to the readout section of the CCD, for which its magnitude has an analytical solution. While not removable from every image, the electronic shutter smear is typically a very small (few DN) effect.

Finally, the corrected DN/s for the image is converted to radiance values by using the temperature-dependent responsivity discussed above. This was an important step omitted from the operations correction algorithm.

The original version of the ground instrument corrections software (CCDCAL) was intended to provide rapid basic calibration of all raw images. As such, some corrections were excluded to save processing time or reduce program complexity. In addition, variations in imaging sequence design and telemetry processing led to unexpected results for some imaging sequences being uncalibratable with the original algorithm.

Bad pixel replacement was not performed, since it was to be generally done in FSW. While there are very few bad pixels that have been identified on the CCD, their existence could result in misinterpretation of the data.

Shutter smear removal was not done, as this was done when FSW shutter correction was applied. However, a large number of images were returned that did not have FSW shutter correction applied.

CCDCAL did not perform pixel-averaging identically to FSW, so that the correction frames differ slightly from the images they are meant to correct.

As discussed above, the CCD temperature behavior is believed to have been much smoother than was reported in the headers. The operations software, however, used the CCD temperature value directly from the header, as input to the dark model.

CCDCAL returned images in units of corrected DN, instead of units of intensity.

Finally, due to a software error, the software offset was not removed, resulting in an additional uncertainty of $\sim 0.5\%$ for well-exposed (~ 3000 DN) regions.

4.3. Ground Target Corrections

Calibration relative to the RTs is more complicated than CCD corrections. The ground target correction pipeline, SPECTCAL, performed the fundamental algorithms necessary for relative spectral calibration, although to facilitate rapid calibration of image data, several more thorough corrections were excluded.

First, both scene and target images (in the same eye and filter) were corrected using CCDCAL in order to remove the instrumental sources of noise and are then divided by exposure time, resulting DN/s values that were linear with intensity. Radiance (in units of DN/s) were measured in the sunlit portions of the white, gray, and black rings and fit to RT reflectance values measured at UA [Reid, 1997] at the corresponding wavelength (Figure 19). For this level of correction the RTs were assumed to be Lambertian. Ideally, the fit of the radiance to reflectances should pass through the origin, but an offset of up to 150 DN is evident in some of the fits. The offset in this fit is due to a number of contributions that will be discussed below. The offset was discarded and only the slope was used. The slope (in reflectance/(DN/s)) was then applied

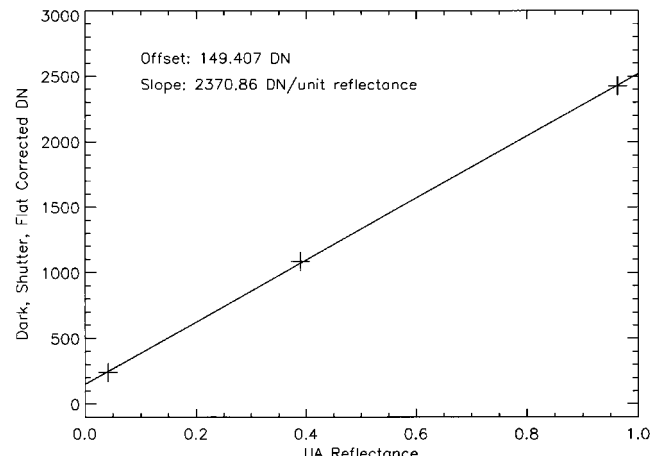


Figure 19. Demonstration of generation of calibration function for R5 (671 nm) filter on lower RT. Fit is excellent through the three sunlit rings, but significant (150 DN) offset indicates that additional factors, such as target photometry and scattered light, need to be taken into account. Offsets were discarded in version 1 of calibration pipeline.

to the relevant scene image, converting the corrected DN/s into reflectance for each pixel.

Since this calibration technique (total intensity of scene normalized by total intensity of target) differs somewhat with conventional definitions of reflectance, we are using the notation R^* (J. F. Bell III, personal communication, 1997) to represent the reflectance values quoted in preliminary analyses. R^* is defined as the brightness of the surface divided by the brightness of an RT scaled to its equivalent Lambert reflectance. It is approximately equivalent to radiance coefficient as defined by Hapke [1981], but since the diffuse sky component is included in the calibration, this is not a true measurement of bidirectional reflectance.

SPECTCAL did not take the photometric properties of the RT into account, assuming them instead to be Lambertian. By plotting the variation in target reflectivity as a function of the solar incidence angle and the azimuth angle between the IMP optical axis and the incident solar vector, the effect of the target photometry on calibration can be seen (Figure 20). The most severe variations in reflectivity occur near the opposition surge at 26° incidence and 0° delta azimuth. No data on Mars could be taken at such illumination conditions due to shadowing by the spacecraft. The majority of target data, however, were acquired in the region of large delta azimuth angle. In this region the target has a significant variation in reflectivity surrounding the specular portion of its phase curve.

In addition, each eye sees the target at a significantly different angle, differing by 24.4° in azimuth for the upper target and 16.6° for the lower target. The result is that each eye sees the target at a significantly different reflectivity due to photometric variation. While the total surface can have as much as 10% variation over the measured geometries, the difference in eye azimuth alone can account for a 5% change in brightness of the target. The photometric properties of the targets are relatively independent of wavelength, however, such that band-to-band uncertainties are typically $\sim 1\%$.

This effect is important when reflectance (R^*) measurements are made for a set of scene images using both eyes. Since measurements of the targets can vary by as much as 5% in




Figure 20. Upper white RT reflectivity (in terms of bidirectional (D-D) reflectance) values as a function of incident and delta azimuth angle. Figure shows strong opposition effect at small delta azimuth, and broad specular rise at large delta azimuth. Much of the target data were acquired in the region of large delta azimuth.

brightness between the two eyes, calibration factors generated from RT measurements will also differ by up to 5%. Spectra that include this discrepancy between calibration factors consequently show a significant offset between the two eyes. This effect was observed early in the mission, causing misinterpretation of the reflectance calculated for the R11 (968 nm) filter. An anomalously large apparent absorption band was due to this property of the target observations and not a property of Martian materials.

The baseline calibration procedure assumed that both the RT and scene images were taken concurrently. In many cases this is a reasonable approximation, since sequences requiring target calibration were often designed with target imaging built in. Later in the mission, however, there were often no target images associated with a scene sequence, or scene sequences were of such duration that illumination conditions changed greatly during their acquisition. At times, target and scene sequences were even acquired on different sols when low-downlink availability prevented concurrent acquisition.

The total illumination on a scene can change by as much as 20%/h for typical 0900 local solar time (LST) ($\sim 45^\circ$ solar incidence angle) imaging sequences (Figure 21). Typical uncertainties introduced by this effect depend on the separation in time between target and scene image but can be as high as 20% absolute, although for most sequences this uncertainty is $<10\%$. Likewise, changes in color of illumination can result in band-to-band uncertainties as high as 5% relative, but for most sequences this uncertainty will be $<2\%$.

The baseline target calibration procedure assumes that the RT sees the entire sky identically as a flat patch of the Martian surface, free from scattering off of spacecraft components. In reality, the target experiences both scattering and sky occultation from the RT shadow post as well as nearby spacecraft components.

Owing to the location of the shadow post at the center of the

RT, the post will occult an increasing solid angle of the sky with decreasing radius on the target. This has the effect of reducing the incident illumination on the target from the sky, resulting in a decrease in target brightness close to the shadow post (Figure 22a). This effect is most pronounced in the shadow, where all illumination is from the sky only, but its effect is also important in the sunlit portion of the target. This effect results in a reduction in brightness by $\sim 5\%$ for sunlit regions close to the shadow post but can be much larger for the shaded regions, decreasing brightness by as much as 50% for the white ring




Figure 21. Brightness of upper white RT as a function of incidence angle. Figure shows strong change in target brightness, especially during typical morning imaging sequence conditions ($\sim 45^\circ$ incidence, 0900 LST). Target and scene sequences that span a significant range of incidence angles must be corrected for these brightness variations.

adjacent to the shadow post. Measurements at that extreme were not generally used during operations.

In addition, since the sky intensity is anisotropic, having a large, bright solar aureole, the target experiences an azimuthal decrease in brightness as more and more of the aureole is blocked by the shadow post (Figure 22b). Light scattered from the shadow post back onto the target also contributes a small amount to the increase in brightness in the sunward direction. This effect is much more pronounced, resulting in a variation in target brightness by as much as 10% for measurements of the white ring.

These shadow post effects are important in that they not only affect correction to absolute reflectance, but they also affect the relative calibration. Since the sky has not only an anisotropic intensity distribution but also an anisotropic color, extraction of target DN values from different regions of the RT may introduce color differences that are due solely to the calibration process and not the properties of the surface materials being measured. This effect can be minimized by extracting RT DN values in a consistent manner band to band, but since the magnitude of this effect was not realized until the late in the mission, many early calibrations relative to the targets will include this uncertainty. These shadow post effects typically reduce the brightness of the RT, which will result in overestimating the apparent reflectivity of surface materials by as much as 5%. Since consistent measurement of the targets to reduce coloration changes have not been done in preliminary calibration, there is also a $\sim 2\%$ uncertainty in band-to-band reflectance values.

As with the shadow post, scattering and sky occultation from the spacecraft are also significant contributors to the measured target brightness. This effect has not yet been extensively studied but is likely to be comparable in magnitude to the shadow post effects. Spacecraft scattering, however, is a much more complicated effect and will be explored in future work.

4.4. Calibration Uncertainties

Owing to the high radiometric accuracy of IMP and extensive laboratory measurements of dark current and flat field properties, IMP image data have great scientific value. Images corrected using the most recent calibration algorithm are accurate to $<5\%$ absolute response and display $<0.4\%$ pixel-to-pixel variability at 3000 DN, close to theoretical limits. Images corrected using CCDCAL, however, in general, do contain an additional 16.0 DN of software offset, resulting in a systematic overestimation of scene brightness by $\sim 0.5\%$ for a 3000-DN exposure, and include several other small systematic uncertainties.

Correction to reflectance in terms of R^* facilitated by observations of the RTs provided good results during landed operations for preliminary science assessment, but significant reductions in uncertainty will be facilitated in the next version of the algorithm and other calibration techniques. Combining the errors discussed above, the uncertainties associated with computation of R^* are typically $<10\%$ absolute, $<3\%$ band to band, and $\sim 5\%$ eye to eye, with a $<5\%$ systematic overestimation of scene reflectance due to sky occultation on the RT by the shadow post. The magnitudes of these uncertainties are very dependent on the circumstances of individual observations, and are often significantly less than the upper limits described here.

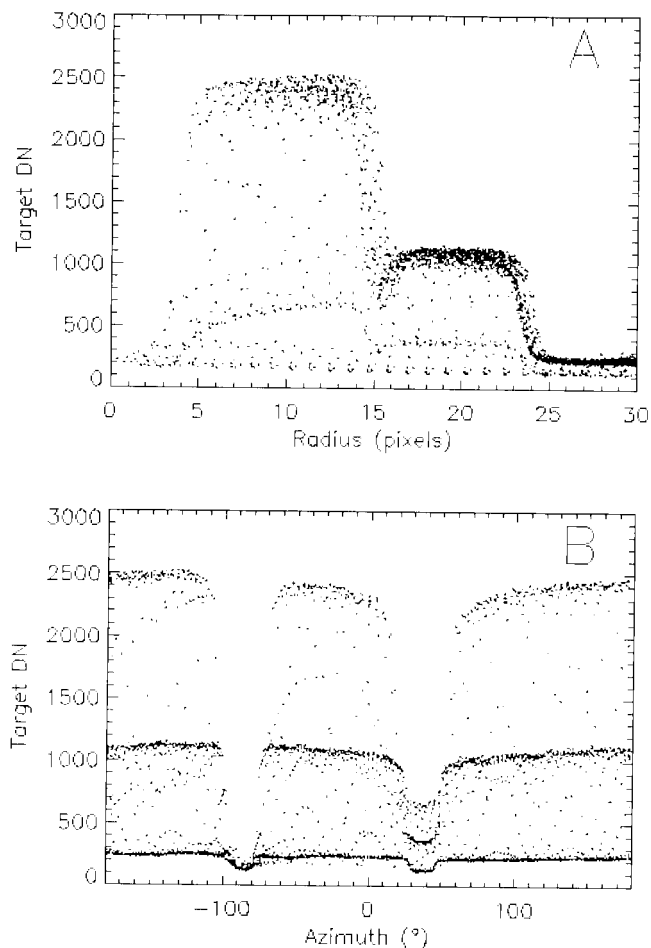


Figure 22. (a) Radial variation in target brightness is caused primarily by the increasing solid angle of sky blocked by the shadow post with decreasing radius on the target. Effect is most pronounced in shaded pixels (lower sets of points for each ring) but is also significant in sunlit pixels. (b) Azimuthal variation in target brightness is caused primarily by occultation on the anisotropic sky by the shadow post. Occultation of the large, bright solar aureole causes broad decrease in brightness as the antisun direction (38° azimuth) is approached on the RT. This effect can cause a significant (5%) uncertainty in relative reflectance calibration if target measurements are not all taken from the same region of the target.

5. Discussion and Future Work

A significant source of uncertainty in calibration of scene images is due to variations in the color and intensity of illumination on surfaces that are oriented differently from the RTs. The target calibration procedure assumes that there is identical illumination on both the RT and the scene, specifically that both are horizontal and see the full hemisphere of sky. Since the sky is anisotropic, and the solar illumination depends on the local incidence angle, the total intensity of irradiance on a surface changes greatly as its orientation changes. Color changes are also induced with changing facet orientation due to the anisotropic color of the direct and diffuse illumination on the facet.

A horizontal facet has a total downward flux curve that is symmetric about local noon, while a facet tilted toward the east sees an asymmetric downward flux curve that peaks in the morning (Figure 23). Since the facet and horizontal surface

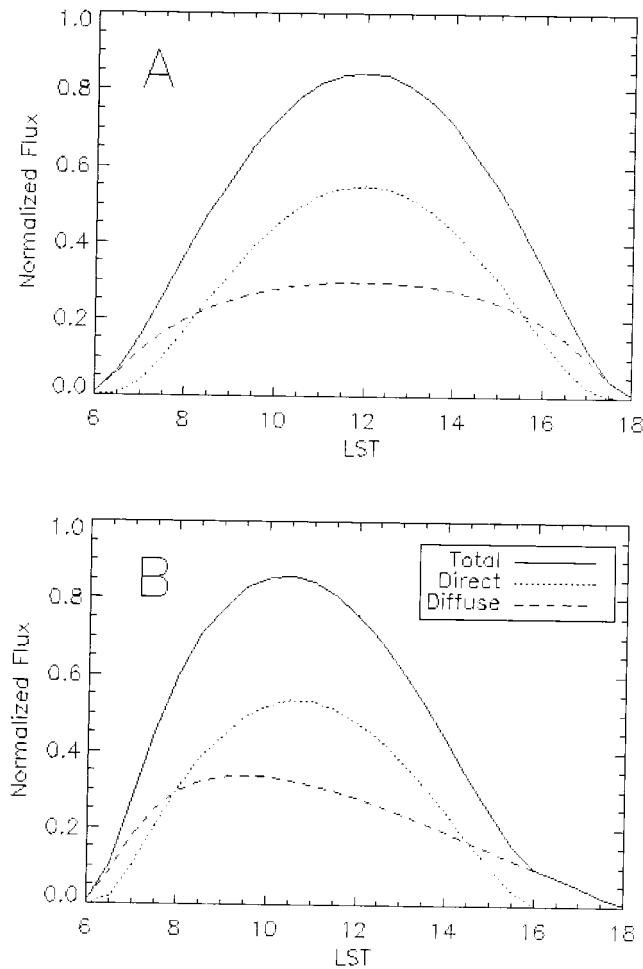


Figure 23. (a) Simulated downward flux on a horizontal facet is symmetric about local noon. (b) Simulated downward flux on facet that is tilted 45° down and 45° South of east has an asymmetric distribution peaked in the morning hours and showing the facet going into shadow at 1600 LST.

(e.g., an RT) see different fluxes of incident radiation at the same time of day, calibration relative to the targets will result in an incorrect determination the reflectivity of the surface. Likewise, color differences between the Sun and sky result in the color of downward flux being different on different facets (Figure 24). For a horizontal facet, the color ratio stays constant, since the flat facet is being calibrated by a flat RT. For a facet tilted toward the east, however, the color of illumination changes such that the object appears more blue in the morning, when it is receiving a disproportionately large contribution from the Sun, and more red in the afternoon, when it is receiving a disproportionately large contribution from the sky. This effect is most pronounced when the surface goes into shadow, and all of the illumination has its source in the reddish sky.

While not an issue of instrument calibration, this effect is still extremely important for determining the true color of the surface of Mars. It is also important to consider when attempting to compare laboratory spectra of materials acquired under a direct incident beam with spectra measured on Mars, where the bright, red sky can cause brightness and color changes of surface materials that are not identically illuminated as the RTs.




Figure 24. Simulated reflectance ratios for same facets in Figure 23 using red reflectance of 0.2 and blue reflectance of 0.05. Apparent reflectances are determined by ratioing to simulated horizontal Lambert surface. Reflectance ratio is constant for a facet that is horizontal, since it is calibrated against a horizontal surface (dotted line). For a facet tilted (solid line) as described in Figure 23 the reflectance ratio varies throughout the day, being more blue early in the morning as direct sunlight dominates and more red later in the afternoon as diffuse skylight dominates. Effect is most pronounced at 1600 LST when surface goes into shadow.

Knowledge of the surface topography and improved atmosphere models permit modeling of the true illumination environment at any location at the landing site, providing a better technique for the photometric calibration of image data. Removing the effects of diffuse atmospheric illumination will be a primary task for future analysis of the IMP data set.

The Imager for Mars Pathfinder performed very well during the course of the Mars Pathfinder mission. Far more image data were returned than was expected, resulting in a great deal of preliminary scientific analysis. Using the extensive knowledge of instrument parameters provided by the IMP calibration team, the scientific value of this data set can be realized. Precise calibration of image data has facilitated interpretation, resulting in the determination of atmospheric water vapor and aerosol properties, the observation of the landing site geomorphology, and the creation of high-resolution panoramas. Measurements of the calibration targets have facilitated the initial determination of the surface reflectance properties, which will be improved as additional factors are accounted for. Lessons learned from the IMP experiment will be propagated to future missions, particularly the Surface Stereo Imager and Robotic Arm Camera aboard the 1998 Mars Polar Lander Surveyor mission.

Acknowledgments. There are numerous individuals without whose help the development and execution of calibration procedures for IMP data would have been impossible. We would like to thank T. Daley and S. F. Hviid for their assistance with support software during mission operations; H. P. Gunnlaugsson for his help in tying magnetic target observations with RT and CT measurements; H. McSween, N. Bridges, and R. Anderson for their assistance with scientific issues during operations; and all of the members of the IMP calibration team, whose many long hours and dedication resulted in the excellent characterization of the Imager for Mars Pathfinder.

References

- Aharoni, G., A. Averbuch, R. Coifman, M. Israeli, Local cosine transform—A method for the reduction of the blocking effect in JPEG, *J. Math. Imag. Vision*, 3, 7–38, 1993.
- Burkland, M., J. Wellman, P. Smith, D. Crowe, R. Tanner, and R. Reynolds, Computer modeling of the Imager for Mars Pathfinder, internal project report, Lunar and Planet. Lab., Univ. of Ariz., Tucson, 1997.
- Gaskill, J. D., *Linear Systems, Fourier Transforms, and Optics*, 499 pp., John Wiley, New York, 1978.
- Goodman, J. W., *Introduction to Fourier Optics*, pp. 121–124, McGraw-Hill, New York, 1968.
- Hapke, B., Bidirectional reflectance spectroscopy, 1, Theory, *J. Geophys. Res.*, 86, 3039–3054, 1981.
- Johnson, J. R., et al., Preliminary results on photometric properties of materials at the Sagan Memorial Station, Mars, *J. Geophys. Res.*, this issue.
- Kramm, J. R., N. Thomas, and H. U. Keller, The CCD detectors of the Halley multicolour camera after seven years in flight, *Rep. MPAE-W-100-93-22*, Max-Planck-Inst. für Aeron., Katlenburg-Lindau, Germany, 1993.
- Kramm, J. R., H. U. Keller, and T. Behnke, Performance degradations observed on scientific CCD imagers from Loral with erratic substrate grounding, *Tech. Rep. MPAE-T-42-94-21*, Max-Planck-Inst. für Aeron., Katlenburg-Lindau, Germany, 1994.
- Kramm, J. R., N. Thomas, H. U. Keller, and P. H. Smith, The CCD imager electronics for the Mars Pathfinder and Mars Surveyor cameras, paper presented at IEEE Instrumentation and Measurement Technology Conference, St. Paul, Minn., May 18–21, 1998.
- Lavine, J. P., et al., Monte Carlo Simulation of the Photoelectron Crosstalk in Silicon Imaging Devices, *IEEE Trans. Electron Devices*, ED-32(10), 2087–2091, 1985.
- Malvar, H. S., *Signal Processing With Lapped Transforms*, Artech House, Norwell, Mass., 1992.
- Patterson, W. R., III, F. O. Huck, S. D. Wall, and M. R. Wolf, Calibration and performance of the Viking lander cameras, *J. Geophys. Res.*, 82, 4391–4400, 1977.
- Reid, R. J., Spectrophotometric calibration of Imager for Mars Pathfinder datasets, thesis, Univ. of Ariz., Tucson, 1997.
- Reid, R. J., et al., Imager for Mars Pathfinder calibration report, Lunar and Planet. Lab., Univ. of Ariz., Tucson, 1998.
- Rice, R. F., Some practical universal noiseless coding techniques, III, Module PSI14, K+, report, Jet Propul. Lab., Pasadena, Calif., 1991.
- Seib, D. H., Carrier diffusion degradation of modulation transfer function in charge coupled imagers, *IEEE Trans. Electron Devices*, ED-21(3), 210–217, 1974.
- Shepard, M. K., R. E. Arvidson, and E. A. Guinness, Specular scattering on a terrestrial playa and implications for planetary surface studies, *J. Geophys. Res.*, 98, 18,707–18,718, 1993.
- Smith, P. H., et al., The Imager for Mars Pathfinder experiment, *J. Geophys. Res.*, 102, 4003–4025, 1997.
- Tomasko, M. G., et al., The Descent Imager/Spectral Radiometer (DISR) aboard Huygens, Huygens science, payload, and mission, *Eur. Space Agency Spec. Publ., ESA-SP-1177*, 109–138, 1997.
- J. F. Bell III, Department of Astronomy, Cornell University, Ithaca, NY 14853.
- D. T. Britt, M. Burkland, M. Lemmon, R. Marcialis, R. J. Reid, P. H. Smith, R. Tanner, E. Wegryn, and J. Weinberg, Lunar and Planetary Laboratory, University of Arizona, P.O. Box 210092, 1629 E. University Boulevard, Tucson, AZ 85721. (breid@lpl.arizona.edu)
- R. Kramm and N. Thomas, Max Planck Institute for Aeronomy, Katlenburg-Lindau, Germany.
- B. J. Bos and D. Crowe, Optical Sciences Center, University of Arizona, Tucson, AZ 85721.
- A. Dummel, DLR, Institute of Planetary Exploration, Berlin, Germany.
- F. Gliem and P. Rueffer, Technical University of Braunschweig, Braunschweig, Germany.
- K. E. Herkenhoff and J. N. Maki, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109.
- J. R. Johnson, U.S. Geological Survey, Flagstaff, AZ 86001.
- R. B. Singer, Department of Mines and Geological Engineering, University of Arizona, Tucson, AZ 85721.

(Received April 13, 1998; revised September 10, 1998; accepted September 15, 1998.)

