Earth and Space Science

RESEARCH ARTICLE
10.1029/2018EA000473

Key Points:
- An anomalous effect of the Pancam 1,009-nm filter has been successfully modeled and simulated based on a hypothesis for the physical origin
- A correction to the effect has been developed, based on the model, which successfully removes it, without adding any artificial signal
- The correction has been verified using independent, in-flight R7 images of the Martian surface, and of the calibration target

Correspondence to:
S. J. Jakobsen,
simonejuuljakobsen@gmail.com

Citation:
Jakobsen, S. J., Kinch, K. M., Madsen, M. B., Bell, J. F. III, Wellington, D., Dajose, L., & Alizai, K. (2019). An instrument anomaly in the Mars Exploration Rover Pancam 1,009-nm filter (R7): Characterization, simulation, correction, and preliminary verification. *Earth and Space Science*, *6*, 96–115. https://doi.org/10.1029/2018EA000473

Received 12 SEP 2018
Accepted 13 DEC 2018
Accepted article online 19 DEC 2018
Published online 25 JAN 2019

An Instrument Anomaly in the Mars Exploration Rover Pancam 1,009-nm Filter (R7): Characterization, Simulation, Correction, and Preliminary Verification

Simone J. Jakobsen[1], Kjartan M. Kinch[1], Morten Bo Madsen[1], James F. Bell III[2], Danika Wellington[2], Lorinda Dajose[3], and Khaled Alizai[4]

[1]Astrophysics and Planetary Science, Niels Bohr Institute, University of Copenhagen, Copenhagen, Denmark, [2]School of Earth and Space Exploration, Arizona State University, Phoenix, AZ, USA, [3]California Institute of Technology, Pasadena, CA, USA, [4]Astrophysics and Atmospheric Physics, Danish Technical University Space, Kgs. Lyngby, Denmark

Abstract During pre-flight calibration of the panoramic camera (Pancam) instrument on board the Mars Exploration Rovers MER A (Spirit) and MER B (Opportunity), a discrepancy was noted between 11-band spectra extracted from Pancam images of the camera's radiometric calibration target and reflectance spectra obtained with a spectrometer. This discrepancy was observed in the longest-wavelength filter of the camera (the longpass R7 filter with system λ_{eff} = 1,009 nm) and consisted of a reduction in contrast between bright and dark regions. Here we describe and characterize this effect. We propose that the effect arises because long-wavelength photons close to the silicon band-gap at 1,100 nm are allowed through the R7 filter, pass through the bulk charge-coupled device, scatter from the backside, pass through the charge-coupled device again, and are registered in a pixel other than the pixel through which they originally entered. Based on this hypothesis we develop a model capable of accurately simulating the effect, and correct for it. We present preliminary results from testing this correction on preflight, as well as in-flight, images. The effect is small, but in some specific cases in small regions of high contrast, the effect is significant. In in-flight images of Martian terrain we observed the signal in dark shadows to be artificially inflated by up to ∼ 33% and analysis of early-mission calibration target images indicated that the reduced contrast due to the artifact is equivalent to >100 DN (full well = 4095 DN) for a hypothetical perfectly dark pixel.

1. Introduction

The two rovers Spirit (MER A) and Opportunity (MER B) of the Mars Exploration Rover mission landed on Mars in 2004 where they began studying the surface using the remote sensing and in situ instruments of the Athena Science Payload (Arvidson et al., 2008, 2006; Squyres et al., 2004a, 2004b, 2006). From their respective landing sites on opposing sides of the planet, they have contributed to the increasing knowledge of the history of water on the Martian surface. After the termination of Spirit's mission in 2010, Opportunity continued to collect data up until 3 June 2018, when a dust storm cut off the connection to the rover. Efforts are presently being put into recovering the connection.

One of the main tools available on both rovers is the panoramic camera imaging system (Pancam; Bell et al., 2006; Bell, Squyres, Arvidson, Arneson, Bass, Blaney, et al., 2004; Bell, Squyres, Arvidson, Arneson, Bass, Calvin, et al., 2004; Bell et al., 2003). The system consists of two cameras on each rover, each equipped with an eight-position filter wheel. The eight filters per camera cover the wavelength range (432–1,009 nm) with the filters denoted R1–R8 for the right eye and L1–L8 for the left eye. Two of the filters (L8 at λ_{eff} = 440 nm and R8 at λ_{eff} = 880 nm) are used for solar imaging, and one is a broadband filter (L1 at λ_{eff} = 739 nm). λ_{eff} here denotes the "effective" or central wavelength, defined as the weighted mean wavelength of the band, where the weighting is by the sensitivity of the whole camera system (filter, optics, and detector). The remaining 13 filters are narrowband "geology" filters with band widths of 17–38 nm. Two of these are overlapping bands for the right eye and left eye, making it possible to obtain stereoscopic images at these specific wavelengths. This leaves the possibility of obtaining 11-point spectra spanning the full wavelength range (Bell et al., 2003).

In addition to being a navigation and assessment tool used as a way to identify interesting sites encountered along the traverse, the Pancam system provides image data used for a variety of scientific data analyses. In order to most accurately convert acquired images from units of radiance to units of radiance factor I/F (where I is the measured scene radiance at the sensor and πF is the incident solar spectral irradiance at the top of the martian atmosphere at the time of the observation), it is necessary to take several images of the onboard radiometric calibration target (hereafter referred to as the cal target) each Sol (a Sol is a day on Mars, which is approximately 24 hr and 37 min long). The calibration process and the conversion from radiance to I/F is described more carefully in Kinch et al. (2015) and Bell et al. (2006).

During preflight calibration activities in 2002, a discrepancy was noted between I/F spectra extracted from Pancam images of the cal target, and equivalent spectra obtained with a spectrometer (Bell et al., 2006). The discrepancy was observed as a small, yet noticeable, decrease in contrast between the brightest areas and the darkest areas of the cal target at the R7 filter wavelength ($\lambda_{\text{eff}} = 1{,}009$ nm). At the time, this observation was hypothesized to be related to absorption properties of the cal target material (mentioned in Bell et al., 2003, on page 24 line 20, and again in the caption of Figure 20), but a full explanation was not found, and the effect was deemed small enough to not affect the image data analysis on the tactical timescales needed for planned rover operations. We present here evidence of the same effect seen in in-flight geology images obtained in the R7 filter and argue that the effect can be significant in images covering high-contrast regions and that the effect may influence the spectra extracted from such regions. The image shown in the left panel of Figure 1 was obtained by the Opportunity rover on Sol 954, and the I/F spectra presented in the right panel show an example of the R7 filter effect. The I/F values from the darkest shadow in the R7 image show substantial increases relative to the neighboring filter (R6 at $\lambda_{\text{eff}} = 934$ nm). One shadow (shadow at cliff edge: dark red) shows an R7/R6 radiance ratio of 1.39 and another (shadow in cliff face: dark green) shows an R7/R6 radiance ratio of 1.33. After conversion to I/F the ratios are 1.31 and 1.26, respectively. These ratios are very large and unlikely to reflect real wavelength-dependent differences in the radiance from these regions that are just shadows but not otherwise expected to be significantly different in composition or mineralogy from other materials in the scene. That this dramatic effect is observed in images of Mars indicates that it is in fact not related to properties of the cal target but is rather a general effect related to the instrument itself.

Several previous works point out anomalies in R7 images relative to other filters of the Pancam system. In a study of dust deposition on the Pancam radiometric calibration target (Kinch et al., 2015), it was found that derived dust albedo values showed a dramatic upturn at the longest wavelength (Figure 8 and section 4.1.2 of Kinch et al., 2015). The authors also find that the ratio between direct and diffuse illumination showed a dramatic downturn at the longest wavelength (Figure 10 and section 4.2.2 of Kinch et al., 2015). The authors hypothesize that these observations are caused at least partly by an instrumental artifact and mention under "future work" that this issue is in need of further characterization (section 5.4 of Kinch et al., 2015).

In Wang et al. (2008), the authors reported the existence of a characteristic spectral downturn from filter R6 to R7 for regions on the Martian surface known to be silica rich. The authors of Rice et al. (2010) use this idea to look for a correlation between the magnitude of the spectral slope and the amount of silica in the soil, in order to look for hydrated minerals. They report a detection of false hydration signatures at very low Sun angles and attribute it to inaccuracies in the cal target calibration model (p. 384, second column in Rice et al., 2010). They further report a detection of hydrated minerals on surfaces at high local emission angles and argue that this signal could be caused by the Martian dust itself (p. 393, first column in Rice et al., 2010).

In this paper, we describe and characterize the R7 anomaly, discuss a possible explanation for it, and develop a model to use in an automated algorithm for the purpose of correction. The focus here is on *showing* that an instrumental artifact exists, explaining the *cause* of this artifact, developing a *correction* for this artifact, and *demonstrating* the performance of this correction on a small data set, specifically chosen to show most clearly the effect of this artifact. It is an obvious implication of the present work that a number of previous studies, including those mentioned above, are in need of re-evaluation in light of the results presented here. We are committed to fully reprocess the relevant data and re-evaluate the conclusions of these previous studies in light of the present work. This work is ongoing and will be presented in a follow-up paper but is not covered here.

In section 2 we describe the method of developing a model and algorithm that satisfactorily simulates the R7 anomaly; in section 3 we present the results for the model parameters, and their interpretation, as well

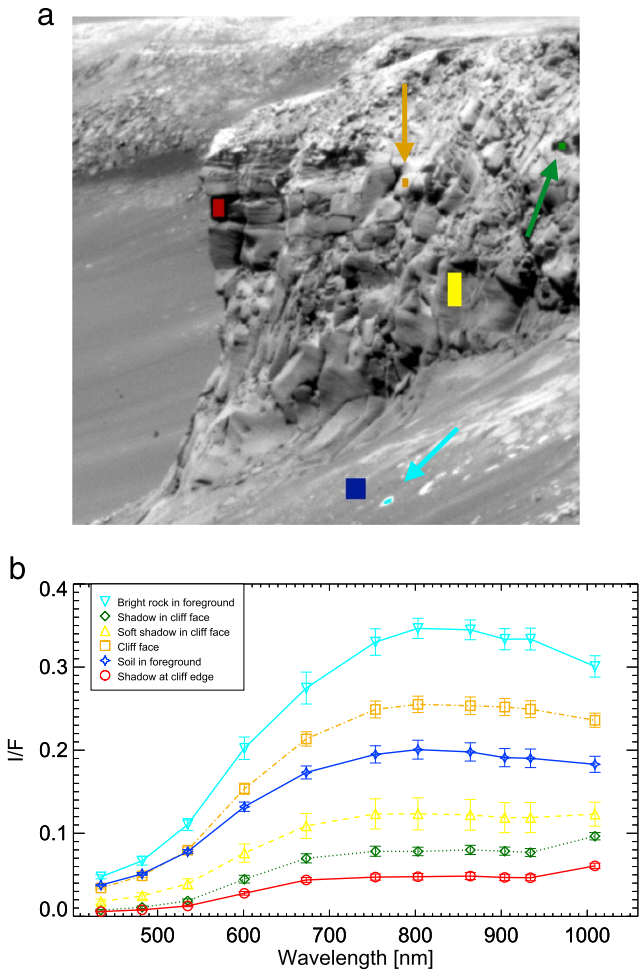

Figure 1. An example of the effect seen in images obtained with the R7 filter. (a) R7 image obtained on Sol 954 by the Opportunity rover. (b) The extracted spectra, after conversion into I/F values. Spectra from very dark areas exhibit a dramatic upturn at R7 ($\lambda_{\text{eff}} = 1{,}009$ nm). Image retrieved from the Planetary Data System (PDS) archive for MER data. No scaling was performed between the right and left eye values. Values at overlapping wavelengths (L7R1 and L2R2) are presented as an average of the two eyes. The error bars show the standard deviation between individual pixels within each region of interest. Image ID: 1P212872229RAD76EVP2586R7C1.

as examples of in-flight images before and after correction. In section 4 we discuss the model as a physical description of the R7 effect as well as the performance of the algorithm in the process of correcting data.

2. Methods

In order to develop a correction for the R7 images, it was necessary to characterize the nature of the effect and propose a hypothesis for how it arises. When a model had been constructed that successfully simulated the effect, the model was applied to the inverse problem in order to correct for the effect. The hypothesized model was tested under known and well-defined conditions before it was introduced on independent image data from the two rovers Spirit and Opportunity. In this section we provide an overview of the entire process of developing a correction for the R7 filter images.

This section is divided into five separate parts. Section 2.1 describes and characterizes the observed effect, section 2.2 treats the development of a model that describes the effect in a satisfactory way, section 2.3 describes how the developed model was implemented into an algorithm that can be applied to image data, section 2.4 looks at the inverse problem of actually correcting for the effect, and, finally, section 2.5 introduces the correction on in-flight image data, giving us the opportunity to test its performance.

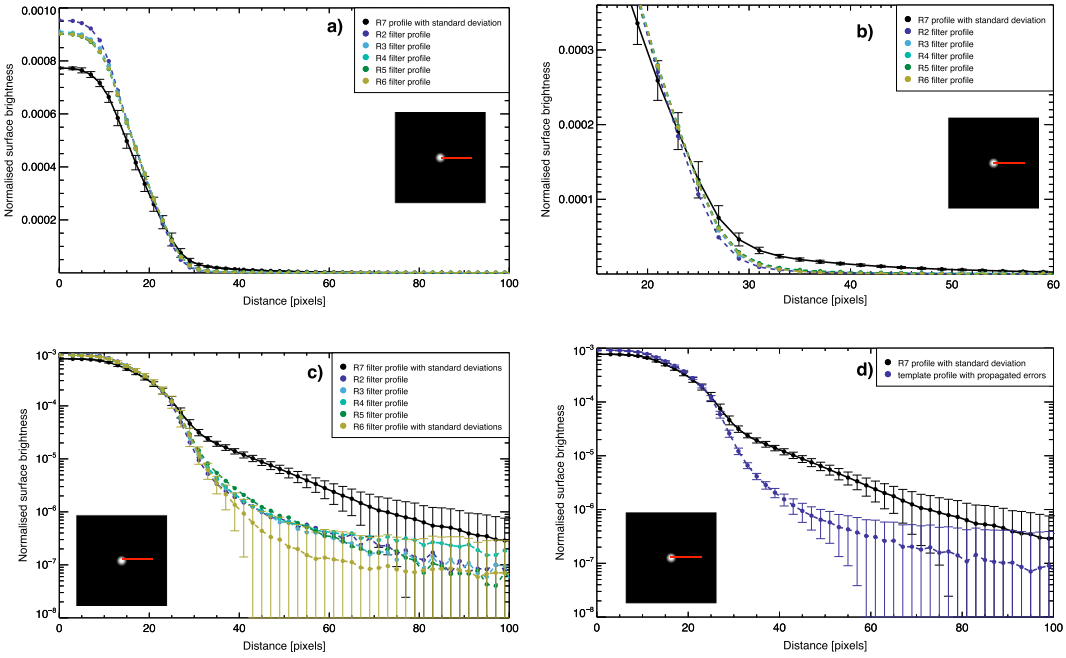

Figure 2. The plot shows six different radial profiles, each representing the light distribution of the source in a filter mounted in front of the right eye of the Pancam on board the Spirit rover. Black curve: R7; colored curves: R2–R6. Note how the R2–R6 profiles are generally very similar and have a distinctively different shape than the R7 filter profile. (a) Linear representation of all filter profiles; (b) linear zoom-in of the region where the difference between the R2–R6 profiles and the R7 profile is largest; (c) logarithmic representation of all filter profiles; (d) the template profile taken as the average of the R2–R6 profiles, and the R7 profile. Image ID: 020715222717_ 0001989_ R7_ 103_ rad.

2.1. Characterization of the Anomaly

We performed a systematic search of the large data set of images acquired during preflight calibration of the Pancam system in order to identify image sequences that most clearly display this subtle artifact in the R7 filter. We identified a sequence of images acquired using the Spirit Pancam at the beginning of a stray-light test as the ideal data set for visualizing and characterizing the effect.

During the test, a series of images of a broad spectrum light source on a dark background were collected using all 13 spectral filters R1–R7 and L2–L7. These test images were obtained in order to look for ghost images and stray-light effects arising as a consequence of charge-coupled device (CCD) overexposure. However, a few images acquired at the beginning of the test sequence only showed the bright, circular light source on a dark background, with no overexposed pixels. The high symmetry, sharp edge, and high contrast between the bright source and the extremely dark background in these images provided the ideal base for characterization of the anomalous effect. These images gave us the means of comparing images between filters, by relying on the sensible assumption that when normalizing the images by dividing each image by the sum of all pixels, the light distribution in the image across all filters should be the same. The assumption is that there is no physical effect that would make the light distribution substantially different from filter to filter in this simple image of a bright, broad-spectrum source on a dark background. Any difference in light distribution between filters would be due to properties of the instrument. Thus, after normalization, the image should be very similar in all filters. This is indeed the case for the images acquired in the filters R2-R6, whereas R7 images look very different (see Figures 2a–2c).

We made use of these nonoverexposed, radiance calibrated preflight, stray-light images from the Spirit rover for characterization of the effect, and subsequent modeling and simulation. We looked for similar images from the Opportunity rover but were unfortunately not able to locate any nonoverexposed images that were suitable for use (the point of the stray-light test, for which these images were acquired, is overexposure). However, since the right eye cameras on the two rovers are designed to be identical, and indeed are effectively identical in all respects as borne out by published preflight calibration parameters (Bell et al., 2003), we

argue that a correction developed based on images from the Spirit camera can be applied to Opportunity images, as long as the images have been calibrated according to established procedures.

The source center in each filter image was identified by fitting the source with a two-dimensional Gaussian. An estimate of the background in each image was calculated as the average radiance value in an annulus with an inner ring radius of 258 and a width of 2 pixels, well outside the source area. Each background-subtracted image was then cropped to a size of 481 pixels by 481 pixels with the source centered, and normalized by dividing each pixel value with the pixel sum of the entire (cropped) image. Lastly, each image was binned by applying 60 concentric annuli centered on the source, each with a width of 2 pixels, and the largest annulus stretching a distance of 120 pixels from the source center. By determining the average radiance in each bin, it was possible to extract a radial profile representing the light distribution in each filter. A plot showing the resulting profiles is presented in Figure 2.

Comparing the profiles for the right eye of Pancam (filters R2 at 754 nm, R3 at 803 nm, R4 at 864 nm, R5 at 904 nm, R6 at 934 nm, and R7 at 1009 nm) reveals that the filters R2–R5 show a very similar distribution. The R6 profile also shows a very similar distribution except for some deviation visible on a logarithmic scale at distances beyond 40 pixels from the center, where all these filters show very low signal. We believe that this small difference between the R6 profile and R2–R5 is related to the background subtraction, and we stress that although distinctly visible on a logarithmic plot the difference in absolute values is much smaller than the difference between the R7 profile and the rest of the filters (see Figure 2b). The R6 profile is lower than R2–R5 only by a value comparable to the standard deviation of pixels in each annulus of the R6 image (see Figure 2c). In order to compare the distribution of the R7 profile to that of the other filters, we averaged the normalized, stray-light images from the filters R2–R6 and extracted a new profile using the same binning procedure. This averaged image is hereafter referred to as the template image, and the profile extracted from the image is hereafter referred to as the template profile (see Figure 2d). This template profile is used throughout this work as a representation of how the light is expected to be distributed when incident on the CCD. The error bars of the template profile are the root-mean-square of the individual errors for the profiles R2–R6 in each extraction bin. By treating the template profile as "truth," we can think of the deviation of the R7 profile as solely a consequence of the instrumental effect, making it possible to model it. This is only possible because of the simplicity and symmetry of the stray-light images.

The most dramatic difference between the R7 profile and the template profile is found between pixel ~ 30 and pixel ~ 60 from the source center, just on the dark side of the transition between the bright and the dark areas of the image. In relative terms, the difference is at a maximum at 49 pixels from the source center (see Figure 2d), where the R7 value is a factor of ~ 10 larger than the template value.

2.2. The Model

In this section we propose a detailed hypothesis capable of describing the instrumental artifact by analyzing the difference between the R7 filter image and the rest of the filters. Unlike the R2–R6 filters, the Pancam R7 filter is a longpass filter, relying on the rapidly declining sensitivity of the CCD itself to provide the long-wavelength cutoff. In contrast, the conceptually similar Mastcam system on board the Mars Science Laboratory (MSL), all filters are band-pass filters (Bell et al., 2017). Since the described effect is only observed on MER, not on MSL, and only in the R7 filter, this leads us to believe that the effect arises from light at wavelengths very close to the edge of the silicon band gap at 1,100 nm. Perhaps the fact that this effect is not observed on MSL Mastcam is related to the long-wavelength cutoff of the longest-wavelength filter on the MSL Mastcam ($\lambda_{\text{eff}} = 1{,}035 \pm 50$ nm).However, the CCDs are different devices, and other differences in the construction and material of the two systems may also be important, and therefore, it may be difficult to directly compare the two.

In any case, assuming that the effect is an intrinsic occurrence of the CCD, we propose that the effect is related to the very long penetration depth in silicon attained by light with wavelengths close to 1,100 nm. We propose that the artifact arises from light penetrating the CCD and reflecting off the backside of the CCD bulk material at the longest wavelengths let in by the R7 filter. Figure 3 shows how instead of being completely absorbed, some of the light hitting a pixel of the CCD will be transmitted through the bulk silicon layer, scatter from the back edge, travel back through the CCD, and reach the active layer at a different, nearby pixel location. Some of the radiation will be registered in this nearby pixel, with the amount determined by the distance x from the pixel receiving the incident light. This ultimately results in a broadening of the signal that was intended to be registered in one single pixel. In effect, this small artifact will manifest itself as an

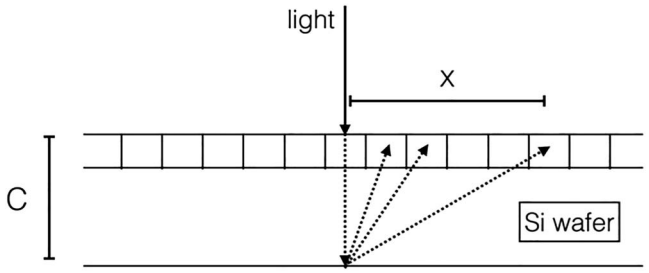

Figure 3. The figure shows a schematic representation of our hypothesis for the nature of the observed effect and how it affects the registration of radiation in one pixel. Part of the light incident on the charge-coupled device (CCD) penetrates the pixel layer without giving rise to a signal and instead reflects off the backside of the CCD eventually resulting in an additional signal being registered at a distance x from the pixel receiving the incident light. C is the bulk thickness of the CCD.

additional very long, very faint, tail in the point spread function for the camera when using the R7 filter. The result in the case of a symmetric source, such as the stray-light image source, will be a broadening of the light distribution across the image as seen in Figure 2.

In order to correct for the effect, we first develop a model capable of simulating the transformation of the template profile into the R7 profile, using our hypothesis as a good description of the behavior of the light let in by the R7 filter ($\lambda_{\text{eff}} = 1,009$ nm).

The total distance traveled by light hitting normal to the surface of the CCD, penetrating the bulk silicon layer, getting backscattered, and eventually reaching the active layer is given by $L = C + \sqrt{C^2 + x^2}$, where C is the thickness of the CCD and x is the distance between the pixel receiving the initial radiation and the pixel registering the backscattered radiation. For convenience we measure distances in units of pixels (1 pixel = 12 μm). The light will be attenuated on its path through the CCD according to the Beer-Lambert law: $I(r) = I_{in} * e^{-B(\lambda)*r}$, where $I(r)$ is the intensity of the light after traveling a distance r through the CCD chip, I_{in} is the intensity of the light incident on the CCD, and $B(\lambda)$ is the absorption coefficient of the CCD silicon, which is wavelength dependent.

For simplicity we assume the scattering off the backside of the CCD to be Lambertian and we therefore introduce a cosine factor into the model: $I(x) \propto I_{in} * e^{-B(\lambda)*L} * \cos(\theta)$, where θ is the angle between the incoming light normal to the plane of the backside of the CCD, and the scattered light, which gives us $\cos(\theta) = \frac{C}{\sqrt{C^2+x^2}}$, and, in turn, $I(x) \propto I_{in} * e^{-B(\lambda)*(C+\sqrt{C^2+x^2})} * \frac{C}{\sqrt{C^2+x^2}}$.

After scattering, the light is no longer collimated and the amount of light reaching a pixel at a particular distance x is approximated by the inverse-square law, giving us the final form of the function describing the scattering of light from the backside of the CCD for one specific wavelength λ:

$$I(x) \propto I_{in} * e^{-B(\lambda)*(C+\sqrt{C^2+x^2})} * \frac{C}{\sqrt{C^2+x^2}} * \frac{1}{C^2+x^2} = I_{in} * e^{-B(\lambda)*L} * \frac{C}{(C^2+x^2)^{3/2}} \quad . \tag{1}$$

We now assume that the chance of generating a signal in the CCD at a given wavelength within the R7 band can be treated as inversely proportional to the penetration depth—in effect this means that we assume that attenuation of the light as it penetrates the bulk layer of the CCD is dominated by electrons lifted into the conduction band and ignore other causes of absorption. This assumption is likely to be good over much of the R7 band but will break down very close to the band gap at 1,100 nm, where the probability of excitation becomes very small and absorption begins to be dominated by other processes.

Making this assumption and introducing an overall scaling factor A_{scale} that contains the intensity of light incident on the CCD, the efficiency of the reflection off the CCD back surface as well as some factor for the conversion efficiency in a pixel, we can write the signal generated in a given pixel by this process due to light that originally reached the CCD at a pixel a distance of x away:

$$f(x)_{\lambda} = A_{scale} * B(\lambda) * e^{-B(\lambda)*L} * \frac{C}{(C^2+x^2)^{3/2}} \quad , \tag{2}$$

with $L = C + \sqrt{C^2 + x^2}$.

Equation (2) expresses how the artificial signal varies as a function of distance x from the pixel where the light originally reached the CCD. It also describes the dependence on wavelength. At shorter wavelengths (high $B(\lambda)$) the signal will disappear because of the exponential term as the absorption is too high for the light to penetrate the CCD. At very long wavelengths the signal is depressed by the low sensitivity of the CCD (low $B(\lambda)$).

Since $B(\lambda)$ varies dramatically over the R7 filter band, we do not expect a single wavelength approximation to be sufficient to describe the observed effect. We therefore expand the model to describe the total effect across the whole wavelength band of the R7 filter. We argue that we can use an exponential profile as an

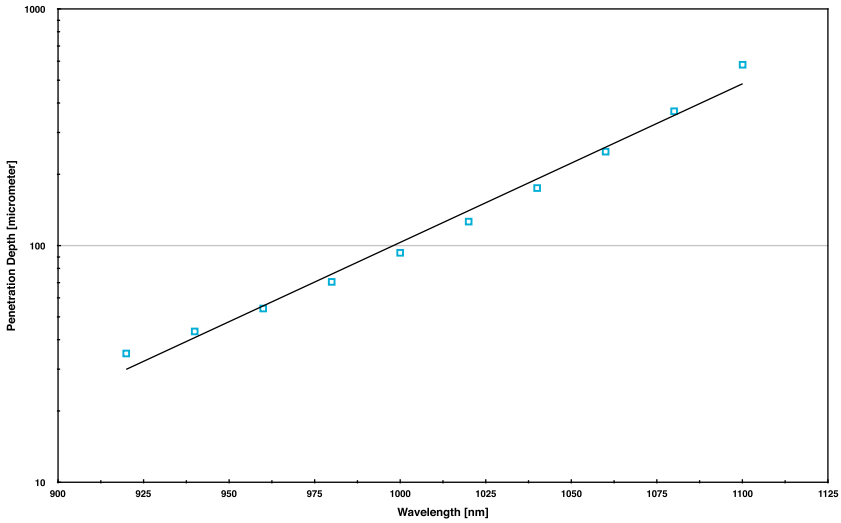

Figure 4. The plot shows the penetration depth in silicon as a function of wavelength in the R7 filter wavelength range. The black line is an exponential fit to the data. The data were adopted from Table 3.1 on page 172 in Janesick (2001).

approximation of how the silicon absorption coefficient decreases with increasing wavelength at the wavelengths of interest (~ 900–1,100 nm). By integrating equation (2), we find a model describing the integrated effect across the R7 filter band.

The validity of such an approximation is supported by the plot seen in Figure 4. The plot has been generated based on a table listing the penetration depths in silicon across a broad wavelength range (p. 172, Table 3.1 in Janesick, 2001).

We use the approximation $B(\lambda) = B_0 * e^{-\alpha*\lambda}$ and integrate equation (2). The lower limit of the integration, λ_{min} is assumed to be at a wavelength where the effect disappears because the light does not penetrate far enough (i.e., $e^{-B(\lambda_{\text{min}})*L} = 0$). This can be thought of as 960 nm at the approximate lower cutoff of the R7 filter, where the penetration depth is approximately ($\sim 100\mu$m), whereas the CCD thickness is at least 400 μm (J. N. Maki, January 12, 2016; see Janesick, 2001, for standard silicon thickness for a front illuminated CCD). The upper wavelength limit, λ_{max}, will be used as a free parameter and can be thought of as the point close to 1,100 nm where silicon becomes almost transparent to the radiation and absorption is dominated by other processes that do not generate electrons in the conduction band and thus do not generate signal in the CCD. Thus, the signal generated in a given pixel can then be expressed as

$$f(x) = \int_{\lambda_{\text{min}}}^{\lambda_{\text{max}}} f(x,\lambda)\mathrm{d}\lambda = A_{\text{scale}} * \frac{C}{(C^2+x^2)^{3/2}} \int_{\lambda_{\text{min}}}^{\lambda_{\text{max}}} B(\lambda) * e^{-B(\lambda)*L}\mathrm{d}\lambda \tag{3}$$

in which we assume that A_{scale} can be treated as wavelength independent. We integrate by substitution. Differentiating $B(\lambda) = B_0 * e^{-\alpha*\lambda}$ with respect to λ gives us $\frac{dB(\lambda)}{\mathrm{d}\lambda} = -\alpha * B(\lambda)$ and by changing limits in our expression we get

$$f(x) = A_{\text{scale}} * \frac{C}{(C^2+x^2)^{3/2}} * \frac{1}{\alpha} \int_{B(\lambda_{\text{max}})}^{B(\lambda_{\text{min}})} e^{-B(\lambda)*L}\mathrm{d}B \quad . \tag{4}$$

Carrying out the integration and assuming $e^{-B(\lambda_{\text{min}})*L} = 0$, we get

$$f(x) = A_{\text{scale}} * \frac{C}{(C^2+x^2)^{3/2}} * \frac{1}{\alpha} * \frac{1}{L} * e^{-B(\lambda_{\text{max}})*L} \quad . \tag{5}$$

Introducing $A = A_{\text{scale}}/\alpha$, writing B as shorthand for $B(\lambda_{\text{max}})$, and inserting $L = C + \sqrt{C^2+x^2}$, we finally get

$$f(x) = A * \frac{1}{C+\sqrt{C^2+x^2}} * e^{-B*(C+\sqrt{C^2+x^2})} * \frac{C}{(C^2+x^2)^{3/2}} \quad . \tag{6}$$

Equation (6) expresses the total contribution to the R7 signal in a given pixel from light that entered the CCD at another pixel at a distance x. The equation has three free parameters: A is an overall scaling factor governing the strength of the effect, B is the absorption coefficient at the point close to 1,100 nm where silicon becomes almost transparent to the radiation and absorption is dominated by other processes that do not generate electrons in the conduction band and thus do not generate signal in the CCD, and C is the physical thickness of the CCD wafer.

The above exercise provides us with an expression that should be able to describe the transformation of the template profile into the R7 profile, as long as our hypothesis is a good approximation for the nature of the effect. Applying the above model pixel by pixel to the template image and extracting a profile from the resulting image should result in a profile that is comparable to the R7 profile. By running this simulation as a fitting routine, we are able to optimize the model and attain approximated values for the model parameters, which can then later be used in the correction.

2.3. Simulating the Effect

Assuming our model is a good approximation of the effect, and arguing that the template image should be a good representation of how the light would be distributed in the R7 image if the effect was not present, we can recreate the R7 image from the template image. As mentioned in section 2.2 the large-scale effect seen in the stray-light images (shown in Figure 2) is assumed to be a representation of the pixel scale effect described by the hypothesis (Figure 3). Hence, by applying equation (6) to each pixel of the template image, it should be possible to simulate the behavior of the radiation incident on the CCD, and recreate the effect.

If we define the R7 image as X_{R7} and the template image as X_T, we can write up the process as follows:

$$X_{R7}(i,j) = X_{\text{Template}}(i,j) + X_{\text{Template}}(i,j) * D + \sum_{k\neq i}^{K}\sum_{l\neq j}^{L} X_{\text{Template}}(k,l) * f(x) \quad , \tag{7}$$

where the parameter D is negative and introduced to keep the normalization of the image, and where (K, L) is the size of the image. This assures that the transformation does not change the total signal in the image (the sum of all pixels will be the same). The variable x still defines the distance from the pixel (i,j) to the pixel (k,l), and $f(x)$ is as given in equation (6). The double sum runs over all pixels within 120 pixels of pixel (i,j). The function $f(x)$ is very small at $x > 120$ (see equation (6)), so this truncation at $x = 120$ does not influence the final result significantly.

The result of this simulation will be an image that is directly comparable to the R7 image through its radial profile. The opportunity of comparing the simulated image directly to the R7 image through their radial profiles makes it possible to fit the model parameters. This in turn provides us with what is needed to develop a correction with the purpose of removing the effect from R7 filter images.

In order to be able to apply the model to each pixel of the template image through equation (7), an algorithm containing the model was formulated using the Interactive Data Language, and a Levenberg-Marquardt weighted least squares fitting routine (MPFIT by Craig Markwardt Markwardt, 2008) was used in order to determine the best fit parameter values for the model. The script reads in the template profile, and the function containing the model is called by determining a set of start values for the model parameters. A radial profile is extracted from the output and evaluated against the R7 filter profile, until an acceptable fit has been achieved.

The first guess for the parameter values was based on their expected physical interpretation according to our model and hypothesis (see section 2.2). This was done to avoid being stuck in a local minimum when running the fit with most of the parameters free. An acceptable first guess for the parameters was determined by running the simulation with different values of the parameters without invoking the fit, and evaluating the result by eye. It proved difficult to get convergence to a stable result when all four parameters (A, B, C, and D) were kept free. For this reason the parameter C was fixed at 33, corresponding to 33 pixels = 396 μm (the minimum thickness of the CCD reported by the manufacturer is 400 μm; J. N. Maki, January 12, 2016; see Janesick, 2001, for standard silicon thickness for a front illuminated CCD). When this was done the result was a very satisfactory fit as can be seen in section 3.1.

2.4. The Inverse Process—The Correction

Equation (6) is only strictly valid for pixels that are more than 120 pixels from the edge of the image. For pixels closer to the edge, a weighting function is applied to equation (7), which takes into account the fact that the

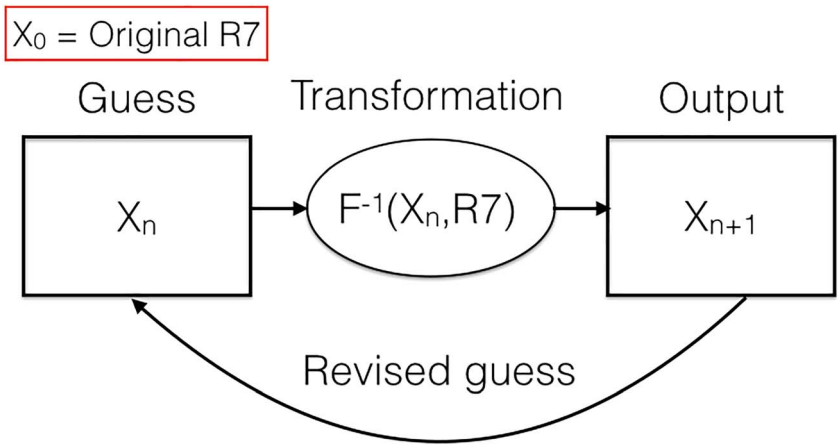

Figure 5. This schematic shows how the R7 images are corrected using an iterative process. Each iteration consists of a full circle in the schematic and results in an output which is a fully transformed image one pixel at a time. This output is then used as the new guess in the following iteration. The output will converge, and when a certain predefined cutoff value is reached, the iterative process will stop.

pixels closest to the edge of the image will be affected only by the pixels inside the image and therefore will be artificially dimmed if no weighting is included.

By simulating the effect as described in section 2.3, we found an estimation of the parameter values in our model. This gave us the means of reversing the transformation, so that we can actually correct for the effect we have so far been simulating. The inverse problem is not as simple as the direct simulation, as can be seen from rearranging equation (7)

$$X_{\mathrm{True}}(i,j) = X_{R7}(i,j) - X_{\mathrm{True}}(i,j) * D - \sum_{k\neq i}^{K}\sum_{l\neq j}^{L} X_{\mathrm{True}}(k,l) * f(x) \quad . \tag{8}$$

Equation 8 describes how to derive the best estimate of the "true" light distribution from a recorded R7 image, that is, how to correct the image. Since X_{True} appears on both sides of the equation and particularly also inside the sum we have to perform the correction through an iterative process. We can achieve a corrected version of the R7 image by letting the corrected output converge toward the true light distribution. Each iteration toward convergence uses the best fit parameter values determined in the simulation of the effect (i.e., $f(x)$ is unchanged). The model is applied to the original R7 image by making a guess for the corrected version of the image. The first guess is the original R7 image itself, which is an appropriate initial guess, since the effect is generally very small. Following each iteration, the output image is used as a new guess until the output converges toward a corrected result. A schematic showing the procedure is provided in Figure 5.

Writing it up in a similar way to the one used in equation (8), and using the terminology from the schematic in Figure 5, we get a correction for each pixel:

$$X_{n+1}(i,j) = X_{R7}(i,j) - X_{n}(i,j) * D - \sum_{k\neq i}^{K}\sum_{l\neq j}^{L} X_{n}(k,l) * f(x) \quad . \tag{9}$$

The correction is performed pixel by pixel, as the simulation was done pixel by pixel. The correction to the (i,j) pixel of the R7 filter image is the summed correction from all neighboring pixels within a distance of 120 pixels from the (i,j) pixel, which is the maximum distance we include in the correction.

The correction from each bin is weighted before it is added to the (i,j) pixel to avoid an artificial brightening of the edge pixels in the corrected image, since these pixels have fewer neighboring pixels. The corrected output from the iteration is used as a new guess for the corrected image. This process is repeated until a convergence is reached. The number of iterations needed to obtain sufficient convergence is determined by the change in the mean of the pixel values in the image:

$$test - value = \sum (X_N - X_{N-1})^2/(K * L) \quad , \tag{10}$$

where X_N is the output image from the current iteration, X_{N-1} is the output image from the previous iteration, and $K * L$ is the size of the image in units of pixels. The correction was invoked 20 times on the stray-light R7 image, and the test value for each iteration was logged and evaluated.

2.5. Correction of In-Flight Images

As a first check of the correction method, the preflight, stray-light R7 image was corrected and the result inspected. This way we were able to choose a cutoff value to use for consecutive corrections as described in section 2.4. However, since the model and therefore the correction was developed using these same preflight images, it is perhaps not surprising that the correction is successful in this case.

2.5.1. Verification Using In-Flight Geology Images

As a first independent verification of the correction we picked two examples of in-flight R7 images to correct and evaluate. A complex geology image from the Spirit rover was chosen based on it exhibiting high-contrast regions. Additionally, a similarly complex and high-contrast image from the Opportunity rover was chosen to further ensure that the algorithm is general enough to correct images from both rovers. Spectra from both images before correction and after correction were extracted and results are presented in section 3. The regions chosen for spectral extraction were regions exhibiting particular high contrast to their surroundings, that is, regions that are very bright and surrounded by a very dark area as well as very dark regions surrounded by a very bright area. The effect of the correction was inspected by producing images for each rover showing the ratio between the uncorrected and corrected R7 image, the ratio between the R2 filter image and all adjacent filters R3–R7, and the corrected version of the R7 filter image. All in-flight images from both rovers used for this test were I/F calibrated using the standard pipeline for calibration of MER images (Bell et al., 2006) and the resulting images are presented in section 3.

2.5.2. Verification Using In-Flight Calibration Target Images

The MER panoramic cameras rely on an external calibration target (cal target) for converting in-flight images of the martian terrain from radiance to radiance factor I/F (Bell et al., 2003). Radiance factor is π times the bidirectional reflectance of the surface and can be easily compared to laboratory reflectance data. The calibration target contains seven different surfaces with well-characterized reflectance properties. During the reflectance calibration process (Bell et al., 2006) radiance values are routinely extracted from these seven regions of a calibration target image. The observed radiances are then plotted against reflectances that are known from preflight characterization. Ideally, the observed radiance should be proportional to the reflectance of the surface, so when plotting observed radiances against known reflectances the data should fall on a straight line through the origin. The slope of this line is a measure of the solar irradiance and this quantity is used to convert images of the Martian terrain obtained at the same time from radiance to I/F.

Later in the mission dust deposition on the cal target must be taken into account and corrected for (Kinch et al., 2015), but early in the mission the cal targets were largely dust free and the observed radiances are generally observed to fall close to straight lines through the origin for all filters—except R7. We obtained data from cal target observations acquired through the first 5 sols of the mission. Sol 5 was chosen as a cutoff from a balance between wanting only very early images to minimize the dust influence but also desiring an adequate number of calibration target observations.

There are a total of 92 cal target observations from Spirit and 172 observations from Opportunity on the first 5 sols when counting all full-resolution images using any one of the 13 narrowband geology filters of the Pancam system. For each cal target observation we performed a fit of observed radiances to model reflectances using a simple linear model with two straight lines representing (a) regions of the cal target in direct sunlight and (b) regions in the shadow of the shadow post. The two lines were constrained to intercept the y (radiance) axis at the same point, but the offset from zero was a free parameter of the fit. Ideally, this offset should be zero as the radiance emitted by the calibration target is proportional to the reflectance of the surface. The recorded offset value was back converted from radiance to data numbers (DNs) by multiplying with the exposure time and dividing by standard preflight-determined conversion factors (Bell et al., 2003). This was done in order to make a better direct comparison of the offset value between filters. The offset in DN can be directly compared to the maximum value of the 12-bit digitization (0–4,095 DN) to get an idea of the size of the offset relative to the average signal level in the image, which will generally be set by the auto-exposure algorithm to be a significant fraction of the maximum value. The results of this analysis are shown in section 3.2.3.

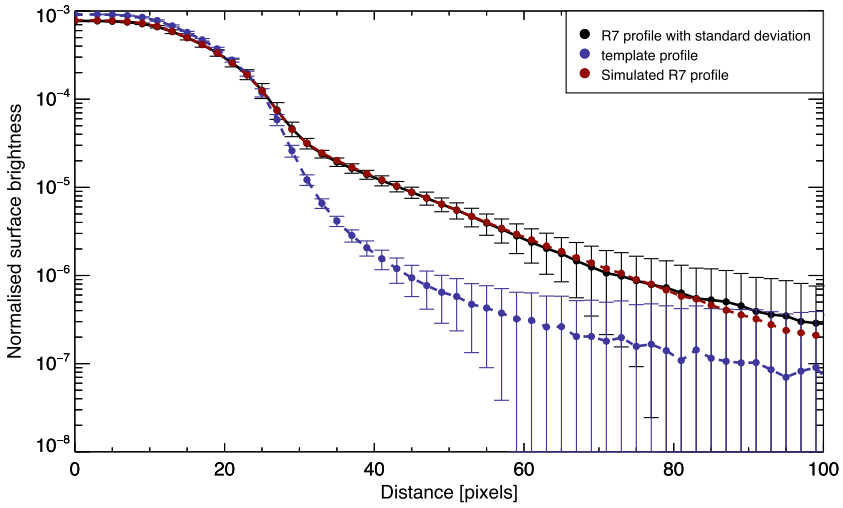

Figure 6. This figure shows a logarithmic plot of the fit used to determine the parameter values for the optimized model. The dark blue curve is the template profile, the black curve is the R7 profile to be simulated, and the dark red curve is the best fit, simulated profile. The uncertainties for the R7 profile is the standard deviation in each bin.

3. Results

This section is divided into two parts: Section 3.1 presents and discusses the results achieved during the development of the algorithm, that is, the simulation of the R7 anomaly based on the stray-light images from preflight calibration of the Pancam for the Spirit rover. Section 3.2 treats the performance of the algorithm when applied to the inverse problem, that is, the actual correction of the R7 filter images. This include both the preflight, stray-light images as well as the in-flight geology and cal target images.

3.1. Simulation

Here we present the results for the development of the algorithm. As described in section 2.3, we used a fitting routine to determine the best fit values for the model parameters. The result of the fit can be seen in Figure 6. The template profile is shown in blue, the R7 profile is shown in black, and the fitted, simulated model is shown as the dark red curve. The fit follows the measured R7 profile almost perfectly, indicating a very good fit. The fit was run with the parameter C held at $C = 33.0$, equivalent to a bulk thickness of the CCD of ~ 396 μm. The fit resulted in parameter values of $A = 96.2$, $B = 0.0388$, and $D = -0.211$. A is a scaling factor determining the overall strength of the effect; B is the absorption coefficient of the radiation at the longest wavelengths, equivalent to a penetration depth of $\sim$309 μm. D is an overall scaling factor ensuring the normalization, that is, that the sum of all pixels in the image is unchanged by the transformation.

The main goal of this investigation was to reproduce the observed R7 filter effect in order to be able to remove it from the affected images. With the fit in Figure 6 we have shown that the effect can be simulated to a good accuracy, which gives us confidence that we have a mathematical description of the effect that is accurate enough to be used for development of a useful correction. Second, since our mathematical description is based directly on a specific physical hypothesis for the origin of the effect, the good fit makes us believe that our physical hypothesis is, in fact, correct and that the instrumental artifact arises from photons close to 1,100 nm that penetrate the CCD, scatter from the backside, pass through the CCD chip again, and are registered on the way out in a different pixel from the one through which they originally entered the CCD. Nonetheless, we point out that the main requirement for the correction is that our model describes the effect well mathematically as demonstrated by Figure 6. The correction may be perfectly useful even if our physical hypothesis is not entirely correct.

3.2. Correction

In the following we look at the correction of the preflight images and the in-flight images in separate sections. The preflight, stray-light images from Spirit were used to develop the algorithm, and the results achieved from the correction of these images is presented in section 3.2.1. Section 3.2.2 presents the results achieved based on correction of images independent of the development of the algorithm, that is, in-flight, geology images from both Spirit and Opportunity, as well as cal target images from both rovers (3.2.3).

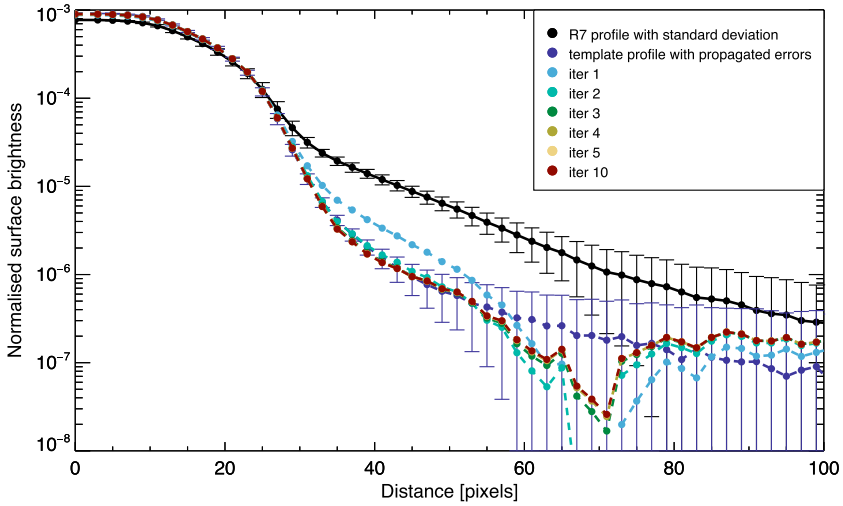

Figure 7. Convergence of the output toward a good correction over 10 iterations. A convergence of the output is reached already after three to four iterations. After four iterations the corrected profile does not change much with more iterations, and it matches the template profile within the errors.

3.2.1. Performance of the Correction on Stray-Light Image Data From the Spirit Rover (Preflight)

The results of the correction of the Spirit stray-light images is presented in Figure 7. The plot shows the convergence of the output after each iteration. The black curve is the light distribution in the uncorrected R7 image; the dark blue curve indicates how the light should in fact be distributed as expected from the template profile. The rest of the colored curves represent the first five, and the tenth iterations, showing that the R7 profile converges toward the correct distribution. It was observed by eye that the profile converged after approximately four iterations. The test value of the fourth iteration (equation (10)) were of the magnitude 10^{-14}, making this our choice for the cut-off value to be used for all other corrections, determining when the iterative process should be stopped. The number of iterations needed for correcting an image could vary; therefore, the cutoff value described in Section 2.4 is introduced to ensure that a sufficient convergence is reached.

3.2.2. Performance of the Correction on Some Examples of Image Data From Spirit and Opportunity (In-Flight)

In this section we present the results of the correction of two geology images, one from the Spirit rover, and one from the Opportunity rover. Both images are independent of the development of the correction.

In Figure 8 we show the uncorrected and corrected images side by side, comparing them for each rover. The figure also shows the ratio between the uncorrected and corrected image for each rover.

Both images were selected because they both contain a large number of high-contrast areas, created by clusters of small and large rocks on the Martian surface for the Spirit image, and by protrusions in the cliff walls in the case of the Opportunity image. Looking at some of the darkest shadows it is marginally noticeable by eye that they become more distinct in the corrected version of the images, and the edge between the darkest pixels and the brighter pixels outside the shadowed area becomes sharper. The same tendency is observable for the brightest areas. The ratio between the corrected and uncorrected R7 image shows that no artifacts develop during the correction, and no systematic effects are observed to distort the images. The primary effect of the correction is a sharpening of the corrected image relative to the uncorrected image. For this reason the ratio of the (sharper) corrected image to the (less sharp) uncorrected image shows a reasonable representation of the terrain.

In Figures 9 and 10 we show examples of spectra extracted from the Spirit and Opportunity images. The I/F values in the uncorrected and corrected R7 filter are represented by unfilled and filled symbols, respectively. We observe that the correction makes the darkest shadows darker, and the brightest regions brighter, as expected, and creates much flatter and unsurprising spectra, removing dramatic upturns and downturns in R7. As expected, the areas that do not exhibit high contrast do not change much.

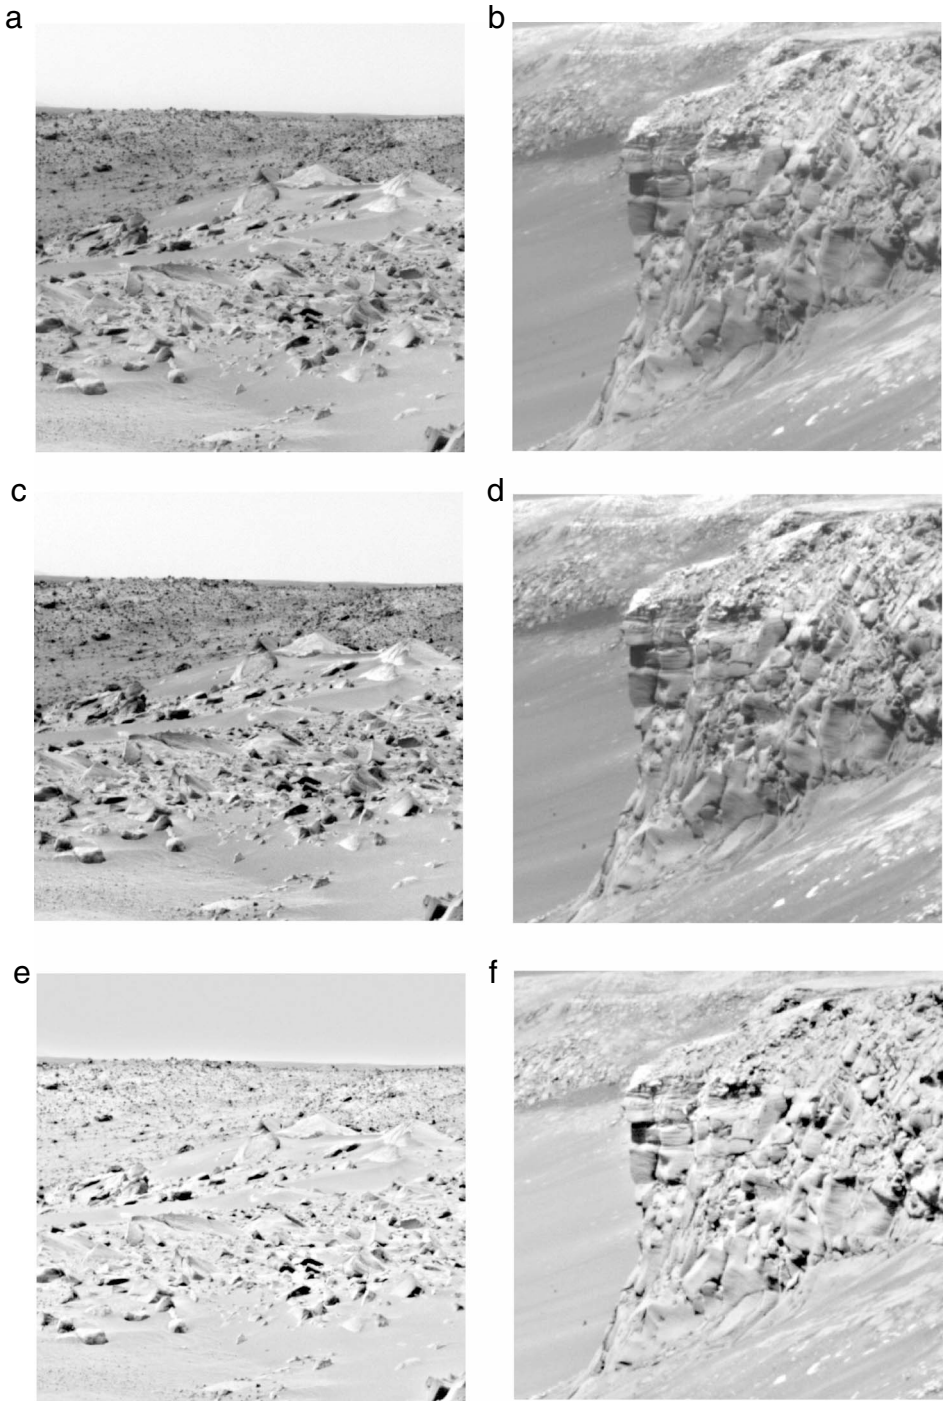

Figure 8. An example of how the correction can be used to improve the contrast in an R7 image. (a, c, and e) The improvement to a selected Spirit R7 image. The image is a radiance calibrated image obtained by the Spirit rover on Sol 77. Image ID: 2P133203792RAD2224P2572R7C5. From top, down: Uncorrected R7 image, corrected R7 image, the corrected image divided by the uncorrected one. The uncorrected and corrected images have been stretched to the same scale, which is the overall minimum and overall maximum of the two images. Min: 8.80 mW/(m^2*nm*sr); Max: 34.2 mW/(m^2*nm*sr). The ratio image was stretched according to its minimum and maximum value: Min: 0.771; Max: 1.06. (b, d, and f) The improvement to a selected Opportunity R7 image. The image is a radiance calibrated R7 image obtained by the Opportunity rover on Sol 954. From top, down: Uncorrected R7 image, corrected R7 image, and the corrected image divided by the uncorrected one. Again, the uncorrected and corrected images have been stretched between the overall minimum and maximum values of the images. Min: 0.0 mW/(m^2*nm*sr); Max: 26.7 mW/(m^2*nm*sr). The ratio image was again stretched to its minimum and maximum values. Min: 0.736; Max: 1.09. Image ID: 1P212872229RAD76EVP2586R7C1.

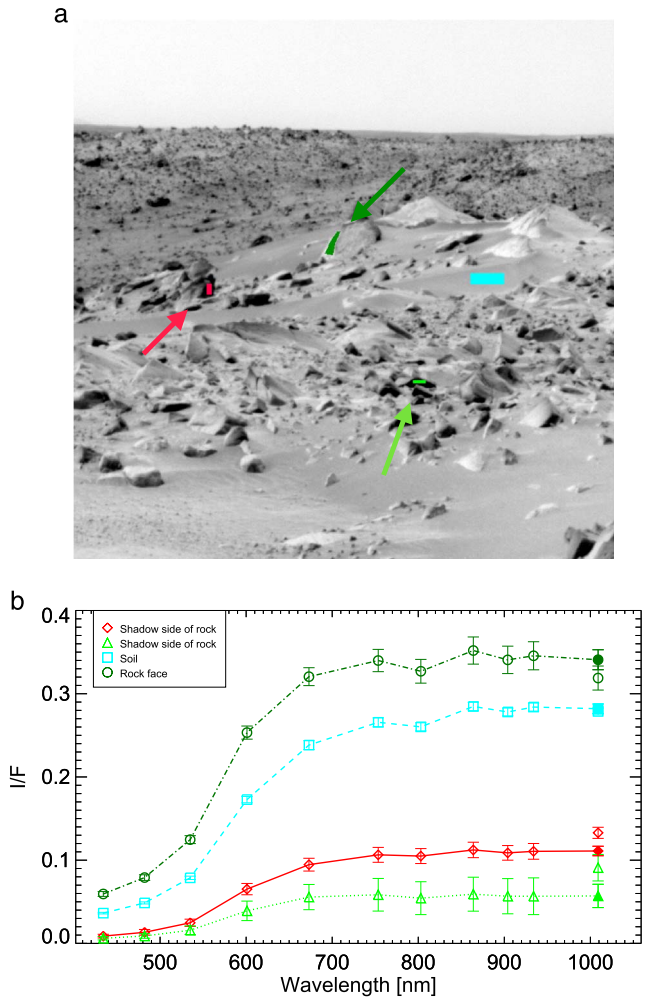

Figure 9. Image showing the ROIs used for extraction of I/F spectra from the Spirit image. (b) The extracted I/F spectra. For R7 both uncorrected (open symbols) and corrected (filled symbols) values are shown. No scaling was performed between the right and left eye values. Values at overlapping wavelengths (L7R1 and L2R2) are presented as an average of the two eyes. The error bars show the standard deviation between individual pixels within each region of interest.

In Figure 11 we show ratio images for the Spirit rover (left column) and the Opportunity rover (right column). Row 1 shows the ratio between the shortest wavelength filter R2 image at $\lambda_{\text{eff}} = 754$ nm and the R6 filter image at $\lambda_{\text{eff}} = 934$ nm, Row 2 shows the ratio between the R2 filter image and the uncorrected R7 filter image, and Row 3 shows the ratio between the R2 filter image and the corrected R7 filter image. The R2/R7 image is dramatically different from the R2/R6 image, almost with bright and dark regions inverted. The R2/R7 image shows a fair representation of the terrain for the same reason that the ratio image of the corrected R7 divided by uncorrected R7 did (Figure 8, bottom row), namely, that the R7 image is significantly less sharp and has lower contrast than the other images. The R2/R7 corrected is almost identical to the R2/R6 image and dominated by more subtle wavelength-dependent differences.

3.2.3. Performance of the Algorithm on In-Flight Calibration Target Images (Spirit and Opportunity)

As described in section 2.5.1 above, the panoramic camera cal target is routinely imaged in all filters of the camera system and these images are used in reflectance calibration of images of the martian terrain. Here we show how the radiance offset in linear fits on the calibration target radiance-reflectance plots demonstrate the artifact in the R7 filter on the close to dust-free calibration targets of the early mission. We also show how the correction improves the situation significantly.

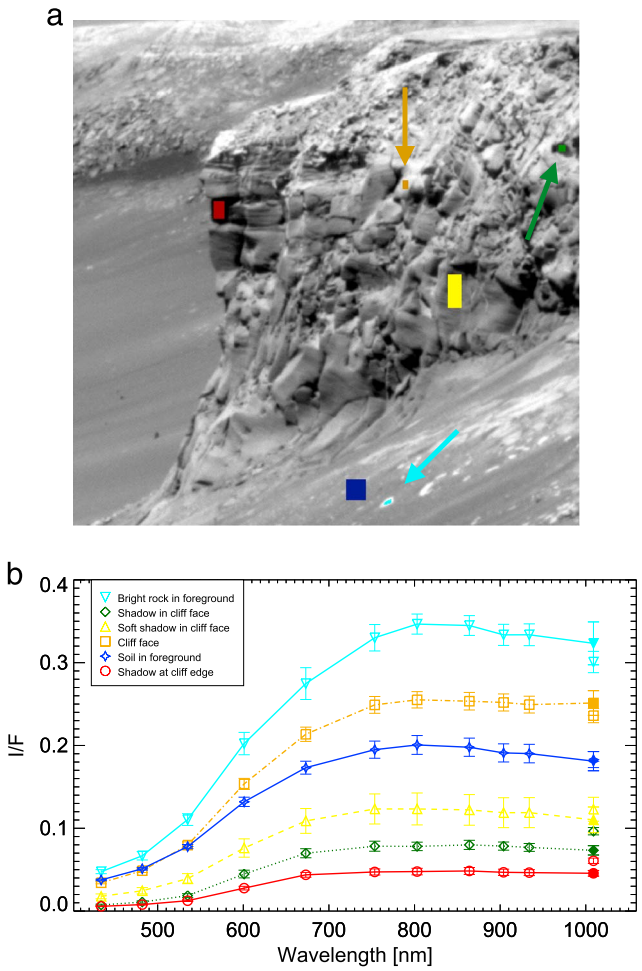

Figure 10. Image showing the ROIs used for extraction of I/F spectra from the Opportunity image. (b) The extracted I/F spectra. For R7 both uncorrected (open symbols) and corrected (filled symbols) values are shown. No scaling was performed between the right and left eye values. Values at overlapping wavelengths (L7R1 and L2R2) are presented as an average of the two eyes. The error bars show the standard deviation between individual pixels within each region of interest.

Figure 12 demonstrates the approach using a typical calibration target image in the R7 filter acquired on sol 3 of Spirit's mission. Ten regions of interest were marked by hand on the calibration target image including the four corner color chips and three grayscale rings plus shadowed regions of the three grayscale rings. Mean radiance values were extracted from each region of interest and plotted against reflectance factor values known from preflight optical characterization of the cal target. The reflectance factor, or $R*$, is defined as $R* = (\mathrm{I/F})/\cos(i)$, where i is the angle of incidence. For this image (ID: 2P126802681ESF0200P2110R7M1.IMG) the recorded offset is 225 DN before correction, 51 DN after correction. This shows that the correction, which has been implemented, is to a very high degree able to reduce the anomalous effect.

Figure 13 shows all the observed offsets in all filters for all calibration target observations for the first 5 sols of the mission on both rovers. The offsets are back converted from radiance to DNs as described in section 2.5.1 for better comparison between filters.

The figure shows very clearly on both rovers that offsets in almost all other filters scatter close to zero (mean offset within 60 DN of zero in all cases), but the offsets in the R7 filter are very clearly above zero with mean values of 107 DN on Opportunity and 220 DN on Spirit. In particular the R7 data are obviously different from the other right-eye longwave filters (within 30 DN of zero for all filters > 700 nm).

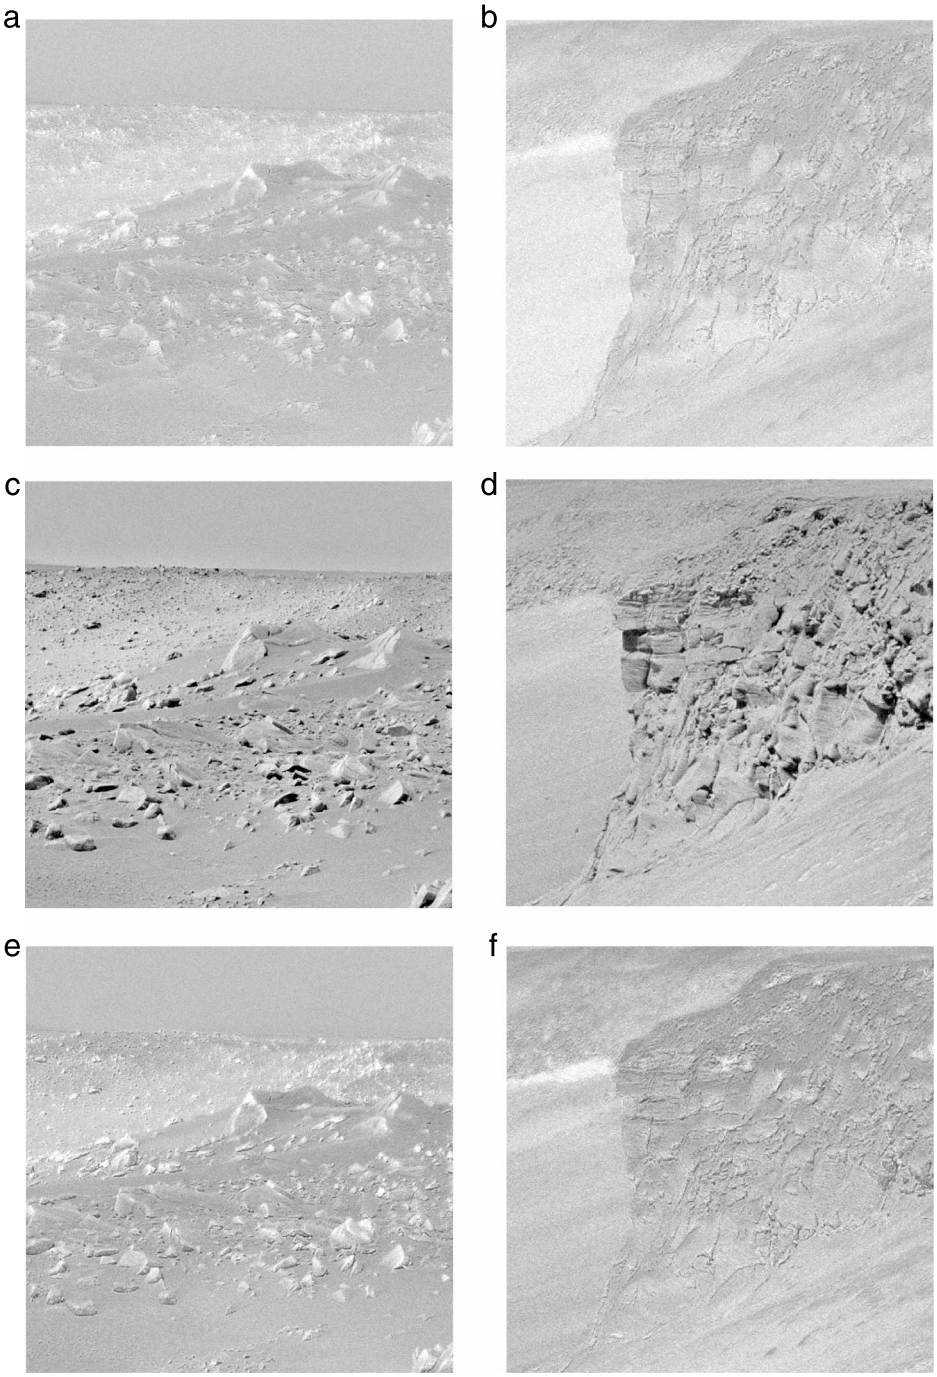

Figure 11. (a, c, and e) Spirit image. (b, d, and f) Opportunity image. From top to bottom: R2 (753 nm) divided by R6 (934 nm), R2 divided by the uncorrected R7 (1,009 nm), and R2 divided by the corrected R7. All images are radiance calibrated. Spirit color stretch (all three images): Min: 1.06 , max: 1.70. Spirit sequence: P2572 of Spirit Sol 77. Opportunity color stretch (all three images): Min: 1.04 , max: 1.80. Opportunity sequence: P2586 of Opportunity Sol 954.

In red are shown the same analysis performed on corrected R7-images. On Opportunity the correction brings the mean offset from 107 DN to −29 DN, bringing the R7-offsets in line with the observations in the other longwave filters. On Spirit the correction brings the mean offset down from 220 DN to 54 DN. This value is still higher than all other filters, but much less dramatically than before. In section 4.2.3 we further discuss the observed difference between the two rovers.

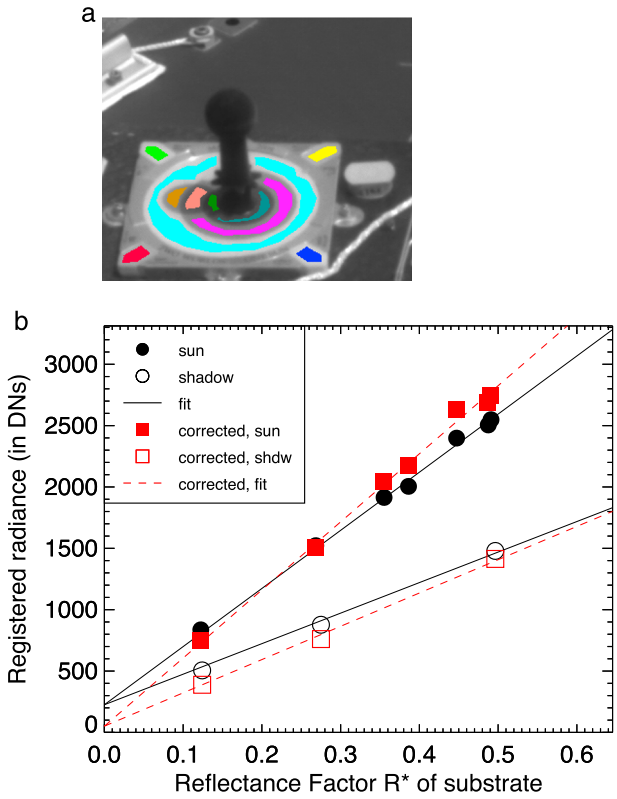

Figure 12. (a) Example calibration target image and (b) radiance-reflectance plot with demonstration of determination of the radiance offset. This image has image ID: 2P126802681ESF0200P2110R7M1.IMG acquired on Spirits sol 5 through the R7 (1009 nm) filter. In black circles (filled and open) are shown mean radiance values extracted from 10 regions on the uncorrected calibration target image. Filled circles are sunlit regions and open circles are shadowed regions. Solid black lines show the fit allowing two different slopes representing different solar irradiances in sunlit and shadowed regions. In red squares (filled and open) and red dotted lines are shown data and fit for the corrected image. The radiances are back converted from radiance to data numbers (DNs) by multiplying with the exposure time and dividing by standard preflight-determined conversion factors; that is, even though given in DN, these are not raw image data, the images have been dark subtracted, flatfielded, etc.

4. Discussion

In the following section we discuss the results from development of the correction, as well as the results from the implementation of the correction on in-flight images from the Spirit and Opportunity rovers.

4.1. Simulating the Effect—The Accuracy of the Model

Of the four parameter values used in the model, two of them, A and D, are scaling parameters that are hard to directly compare to independent information. The two others, however, B and C, have very clear physical interpretations. C, as already described, is the bulk thickness of the CCD, and was fixed at 33, equivalent to 396 μm. This value was based on the 400 μm minimum thickness reported by the manufacturer (rounded to the nearest whole number of pixels). Thus, the fact that the model fits the data using this specific, physically reasonable value for C, may be taken as further evidence in favor of our hypothesis. The physical interpretation of the parameter B in the model is the minimum absorption coefficient in silicon in the wavelength range covered by the model. Converting the parameter B to a penetration depth, for easier interpretation, gives us a value of $1/0.0388$ pixels$^{-1} * 12$ μm pixel^{-1} $=$ 309 μm, where 12 μm is the physical size of a pixel (Bell et al., 2003). Consulting table 3.1 on page 172 of Janesick (2001; see Figure 4) we find listed values for the penetration depth in silicon of 250 μm at 1,060 nm, 370 μm at 1,080 nm, and 582 μm at 1,100 nm. Thus, we would perhaps have expected a somewhat higher value than we get. However, (1) the tabulated values are for pure silicon, whereas we do not have access to measurements of the exact penetration depths in the bulk material of this specific (doped-silicon) CCD chip, and (2) our model neglects loss processes (impurities, thermal loss, etc.) that do not lead to the generation of signal in the CCD and these processes may well

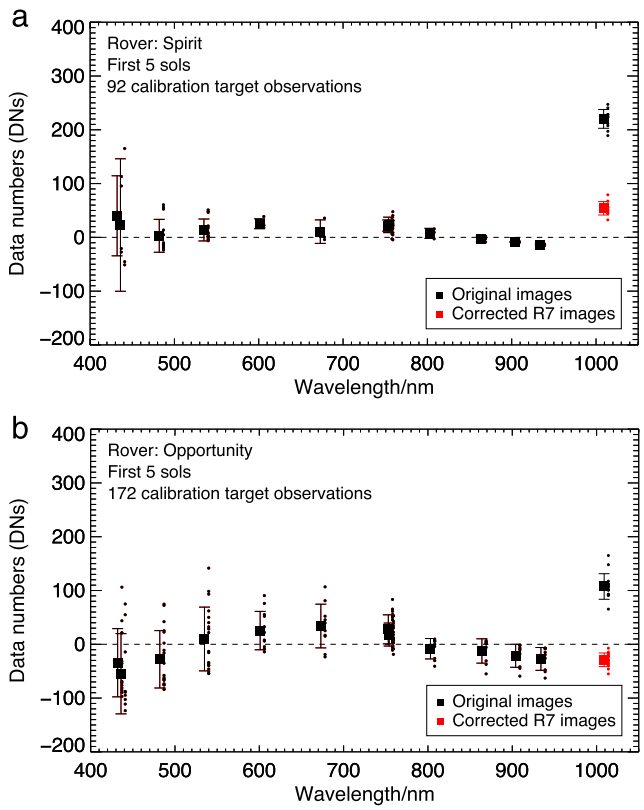

Figure 13. Offsets in radiance-reflectance plots for all calibration targets during the first sols on Spirit (a) and Opportunity (b). As described in section 2.5.1 the offsets are found by fitting two straight lines with a common intercept to the radiance-reflectance plot and recording this y axis intercept. The offsets are back converted from radiance to data numbers (DNs) by multiplying with the exposure time and dividing by standard preflight-determined conversion factors. This was done in order to make a better direct comparison of the offset value between filters. Each small filled circle is an offset found from a single cal target image. For each filter, the mean offset is shown as a large filled square with error bars showing ±1 standard deviation. The black data points are original images. Red data points are R7 images corrected using our correction algorithm as described in section 2.4.

dominate close to 1,100 nm. Thus the fit value we get probably reflects the highest penetration depth that contributes significantly to the signal, and this may not be all the way at 1,100 nm. Given these uncertainties, we consider the fit value of 309 μm to be fully consistent with our hypothesis.

4.2. Corrected Images

4.2.1. Preflight Images

The correction of the stray-light image from Spirit was done as a first check of the iterative method. The effect is expected to be most pronounced in areas of high contrast, and since the stray-light images are indeed high in contrast, correcting these images was a good first test. Furthermore, in correcting these images we were able to directly compare the output to a "true" value, since we were able to infer how the light should be distributed before it was affected by the R7 filter effect from a template profile calculated using the R2–R6 images. By comparing directly to this template, we could ensure that no inherent issues were present that would introduce other artifacts in the image. As discussed in section 3.1 we made sure that no artifacts were created due to the correction by first correcting a simple stray-light image of which we had a good representation of the true light distribution in the form of the template profile. Using this, we found the correction to be efficient as the result converged after only four iterations. A test value cutoff of 10^{-14} was chosen based on the initial correction, and ensures that this very simplistic image, which we would expect to be most simple to correct, sets a strict limit for an acceptable correction of the more complex geology images. It was important to choose a cutoff value, since iterating a large number of times will eventually result in a too noisy image.

4.2.2. In-Flight Images

From the correction of in-flight images, we saw an improvement in contrast in spectra extracted from both Spirit and Opportunity collected images. Generally, the corrected spectral values do not exhibit the dramatic changes from the R6 filter at $\lambda_{\text{eff}} = 934$ nm to the R7 filter at $\lambda_{\text{eff}} = 1{,}009$ nm that were observed for the uncorrected images. We observe an increase in contrast, as expected. When taking the ratio of the corrected R7 image divided by the uncorrected R7 image, the ratio ranges from 0.771 to 1.06 for the Spirit example. For Opportunity the ratio ranges from 0.736 to 1.09. In other words, in these specific examples, some (dark) pixels get $\sim$ 25% darker after correction and some (bright) pixels get 5–10% brighter after correction. Inverting these numbers, this indicates that the R7 artifact inflates the recorded radiances from some dark shadows by up to $\sim$ 33%.

When examining the images directly, we observed that taking the ratio between different infrared filters, there is a noticeable difference when including the uncorrected R7 filter image. We interpret this difference as due to the instrumental artifact in that the distribution of light is not expected to change drastically from the R6 filter to the R7 filter, especially not when the images obtained with the filters R2–R6 are so similar. The observed difference, that is the instrumental artifact, was efficiently removed after correction, resulting in an R7 image much more similar to the R2–R6 images.

4.2.3. Cal Target Radiance-Reflectance Offsets

The Pancam external calibration targets contain a number of surfaces that span relatively few pixels in Pancam images but have dramatically different reflectance properties in close proximity to each other. As such we would expect images of these targets to be exactly the type of images affected by the R7 artifact as photons from high-reflectance cal target areas pass through the CCD, scatter from the backside and are redistributed to nearby dark regions. The net effect would be to reduce observed radiances from bright regions and increase observed radiances from dark regions. This is what was observed in preflight cal target images (Figure 20 in Bell et al., 2003) and at the time hypothesized to be due to partial transparency of cal target materials at 1,009 nm. However, as we have shown, the R7 effect is present also in images of the Martian surface, proving that the effect is an instrumental effect unrelated to the cal target. When plotting observed radiances against known reflectances, the low-reflectance points will be offset upward in radiance while the high-reflectance points will be offset downward (see Figure 12). Thus, the straight-line radiance-reflectance fit will intercept the y axis above zero, as observed in Figures 12 and 13.

Figure 13 demonstrates that the correction performs very satisfyingly on the Opportunity early-mission calibration target observations, so that after correction no significant difference in offset is observed between R7 and the other infrared filters. On Spirit, before correction, the observed offsets are more than double what they are on Opportunity (220 vs. 107) and the correction appears not fully successful in removing the offset—although the situation is much improved after correction. The observed 54 DN offset on Spirit represents less than 1.5% of the full well level (4095) and we point out that there is some nonvanishing systematic uncertainty inherent in the calibration target reflectance model—particularly in the longest- and shortest-wavelength filters that have wider bands and lower S/N. This can also be seen from the filter-to-filter variation in mean offsets in Figure 13. It is of course also quite possible that small differences in the filter, optics, and particularly the CCD chip between Spirit and Opportunity make the R7 effect have different strength between the two rovers. Given that the correction algorithm was developed based on Spirit preflight images one would a priori expect it to perform better on Spirit data, though.

5. Conclusions

In this section we will summarize the results presented in the preceding sections and discuss implications for researchers working with archived Pancam data and our plans for future work.

5.1. Summary

We have demonstrated the existence of an instrumental artifact in preflight and in-flight Pancam images taken through the R7 (1,009 nm) narrowband filter and shown that the effect is an inherent instrumental effect, rather than an effect related to properties of cal target materials. We have described the effect, which manifests itself as effectively a very long, very faint tail in the point-spread function in images acquired through this filter. We have proposed a physical explanation, namely that the effect arises from photons close to the silicon band gap that penetrate the CCD chip, scatter from the backside, penetrate the CCD again, and are registered on the way back in a different pixel from the one through which they originally entered

the CCD. We have presented an algorithm for the correction of such images in order to (largely) remove the artifact and have presented preliminary verification that this correction algorithm (1) largely does remove the artifact and (2) does not introduce other artifacts.

5.2. Implications

The artifact in R7 images is in most cases fairly weak but can be significant in certain observations, particularly at sharp edges with high contrast. In some dark shadows we observed the signal to be artificially inflated by up to ~ 33% and analysis of early-mission cal target images indicated that the reduced contrast due to the artifact is equivalent to >100 DN (full well = 4095 DN) for a hypothetical perfectly dark pixel. Accordingly, it is possible that scientific conclusions relying on such high-contrast R7 images may be affected. For scientists currently working on archived or newly acquired Pancam R7 data, we recommend caution in drawing any firm conclusions. Additionally, previously published analyses based on R7 filter images are in need of reevaluation, in order to make sure that the conclusions drawn are left unchanged by the correction.

5.3. Future Work

For further verification that our correction algorithm is sound, we are working to implement the correction on a larger data set, particularly on a number of specific R7 observations that have previously been used as the basis of analysis in published work. In section 1 we briefly mentioned some examples of such work. Correction of the data used for those and other publications will give us a better understanding of the implication of this instrumental artifact on higher-order conclusions derived from Pancam observations. The results from this investigation will make a natural follow-up publication to the present work. It is our hope and belief that this work will demonstrate beyond doubt that the correction algorithm is stable, useful, and can be sensibly implemented on the full data set of acquired R7 images. Provided that the correction algorithm can be thus fully verified we will implement it in a recalibration of the full R7 data set for archiving of corrected R7 images.

Acknowledgments

This work was funded by the Danish Council for Independent Research in the Natural Sciences, grant 4002–00292, by the Carlsberg Foundation, grant CF16–0981, and by the NASA Mars Exploration Program and Mars Data Analysis Program. The data forming the foundation of this work are the radiance- and reflectance-calibrated images from the Pancams on the two Mars Exploration Rovers. These images are available through NASA's Planetary Data System at the geosciences node hosted by Washington University in St. Louis (http://pds-geosciences.wustl.edu/). Preflight calibration images and derived data products, such as corrected images as well as code used in the analysis, are available by direct communication with the first author (simonejj@nbi.ku.dk). We would be happy to work with any interested parties to provide corrected R7 images for analysis. When the algorithm is fully verified, work will be put into archiving corrected R7 images.

References

Arvidson, R. E., Ruff, S. W., Morris, R. V., Ming, D. W., Crumpler, L. S., Yen, A. S., et al. (2008). Spirit Mars Rover Mission to the Columbia Hills, Gusev Crater: Mission overview and selected results from the Cumberland Ridge to Home Plate. *Journal of Geophysical Research*, *113*, E12S33. https://doi.org/10.1029/2008JE003183

Arvidson, R. E., Squyres, S. W., Anderson, R. C., Bell, J. F. III, Blaney, D., Brückner, J., et al. (2006). Overview of the Spirit Mars Exploration Rover Mission to Gusev Crater: Landing site to Backstay Rock in the Columbia Hills. *Journal of Geophysical Research*, *111*, E02S01. https://doi.org/10.1029/2005JE002499

Bell, J. F. III, Godber, A., McNair, S., Caplinger, M. A., Maki, J. N., Lemmon, M. T., et al. (2017). The Mars Science Laboratory Curiosity rover Mastcam instruments: Preflight and in-flight calibration, validation, and data archiving. *Earth and Space Science*, *4*, 396–452. https://doi.org/10.1002/2016EA000219

Bell, J. F. III, Joseph, J., Sohl-Dickstein, J. N., Arneson, H. M., Johnson, M. J., Lemmon, M. T., & Savransky, D. (2006). In-flight calibration and performance of the Mars Exploration Rover Panoramic Camera (Pancam) instruments. *Journal of Geophysical Research*, *111*, E02S03. https://doi.org/10.1029/2005JE002444

Bell, J. F. III, Squyres, S. W., Arvidson, R. E., Arneson, H. M., Bass, D., Blaney, D., et al. (2004). Pancam multispectral imaging results from the Spirit Rover at Gusev Crater. *Science*, *305*, 800–806.

Bell, J. F. III, Squyres, S. W., Arvidson, R. E., Arneson, H. M., Bass, D., Calvin, W., et al. (2004). Pancam multispectral imaging results from the Opportunity Rover at Meridiani Planum. *Science*, *306*, 1703–1709.

Bell, J. F. III, Squyres, S. W., Herkenhoff, K. E., Maki, J. N., Arneson, H. M., Brown, D., et al. (2003). Mars Exploration Rover Athena Panoramic Camera (Pancam) investigation. *Journal of Geophysical Research*, *108*, 8063. https://doi.org/10.1029/2003JE002070

Janesick, J. R. (2001). *Scientific charge-coupled devices*. Bellingham, Washington: SPIE - The International Society for Optical Engineering.

Kinch, M. K., Bell, J. F. III, Goetz, W., Johnson, J. R., Joseph, J., Madsen, M. B., & Sohl-Dickstein, J. (2015). Dust deposition on the decks of the Mars Exploration Rovers: 10years of dust dynamics on the Panoramic Camera calibration targets. *Earth and Space Science*, *2*, 144–172. https://doi.org/10.1002/2014EA000073

Markwardt, C. B. (2008). Non-linear least squares fitting in IDL with MPFIT. In D. Bohlender, P. Dowler, & D. Durand (Eds.), *Astronomical data analysis software and systems XVIII* (Vol. 411, pp. 251–254). Quebec, Canada: Astronomical Society of the Pacific.

Rice, M. S., Bell, J. F. III, Cloutis, E. A., Wang, A., Ruff, S. W., Craig, M. A., et al. (2010). Silica-rich deposits and hydrated minerals at Gusev Crater, Mars: Vis-NIR spectral characterization and regional mapping. *Icarus*, *205*(2), 375–395.

Squyres, S. W., Arvidson, R. E., Bell, J. F., Brückner, J., Cabrol, N. A., Calvin, W., et al. (2004a). The Spirit Rover's Athena science investigation at Gusev Crater, Mars. *Science*, *305*, 1025–1035.

Squyres, S. W., Arvidson, R. E., Bell, J. F., Brückner, J., Cabrol, N. A., Calvin, W., et al. (2004b). The Opportunity Rover's Athena science investigation at Meridiani Planum, Mars. *Science*, *306*, 1698–1703.

Squyres, S. W., Arvidson, R. E., Blaney, D. L., Clark, B. C., Crumpler, L., Farrand, W. H., et al. (2006). Rocks of the Columbia Hills. *Journal of Geophysical Research*, *111*, E02S11. https://doi.org/10.1029/2005JE002562

Wang, A., Bell, J. F. III, Li, R., Johnson, J. R., Farrand, W. H., Cloutis, E. A., et al. (2008). Light-toned salty soils and coexisting Si-rich species discovered by the Mars Exploration Rover Spirit in Columbia Hills. *Journal of Geophysical Research*, *113*, E12S40. https://doi.org/10.1029/2008JE003126