



Design Through Operation of an Image-Based Velocity Estimation System for Mars Landing

ANDREW JOHNSON*, REG WILLSON, YANG CHENG, JAY GOGUEN, CHRIS LEGER, MIGUEL SANMARTIN AND LARRY MATTHIES

Jet Propulsion Laboratory, California Institute of Technology, CA, USA

Received November 7, 2005; Accepted December 6, 2006

First online version published in January, 2007

Abstract. During the Mars Exploration Rover (MER) landings, the Descent Image Motion Estimation System (DIMES) was used for horizontal velocity estimation. The DIMES algorithm combined measurements from a descent camera, a radar altimeter, and an inertial measurement unit. To deal with large changes in scale and orientation between descent images, the algorithm used altitude and attitude measurements to rectify images to a level ground plane. Feature selection and tracking were employed in the rectified images to compute the horizontal motion between images. Differences of consecutive motion estimates were then compared to inertial measurements to verify correct feature tracking. DIMES combined sensor data from multiple sources in a novel way to create a low-cost, robust, and computationally efficient velocity estimation solution, and DIMES was the first robotics vision system used to control a spacecraft during planetary landing. This paper presents the design and implementation of the DIMES algorithm, the assessment of the algorithm performance using a high fidelity Monte Carlo simulation, validation of performance using field test data and the detailed results from the two landings on Mars.

DIMES was used successfully during both MER landings. In the case of Spirit, had DIMES not been used onboard, the total velocity would have been at the limits of the airbag capability. Fortunately, DIMES computed the actual steady state horizontal velocity and it was used by the thruster firing logic to reduce the total velocity prior to landing. For Opportunity, DIMES computed the correct velocity, and the velocity was small enough that the lander performed no action to remove it.

Keywords: velocity estimation, feature tracking, computer vision, robotics, Mars lander, Mars Exploration Rover, DIMES

1. Introduction

The purpose of the Mars Exploration Rover (MER) Entry Descent and Landing (EDL) system is to deliver the rover and its science payload safely onto the surface of Mars. During entry, atmospheric drag slows the lander down to the conditions where a supersonic parachute can be deployed. Just before impact, retro rockets are fired to slow the lander down even further and airbags that fully encompass the lander are inflated. If the total velocity at impact is less than a threshold then the airbags will cushion the lander as it bounces along the surface of Mars before rolling to rest. If the velocity is too large then tests have shown that the airbag can rip and tear. The EDL system is designed to keep the lander velocity below the threshold of the airbag capability.

Early in EDL development it was determined that wind shear could cause the Rocket Assisted Deceleration (RAD) retro rockets to fire at an angle that would induce horizontal velocity. To eliminate this RAD-induced horizontal velocity and its adverse effect on airbag performance, a system called the Transverse Impulse Rocket System (TIRS) was added to reorient the lander system so that the retro rockets are always fired close to vertical. Later, in the fall of 2001, the MER EDL design team recognized that the current airbag design might not accommodate the additional velocity at impact induced by constant “steady state” winds that are common at the surface of Mars. This *steady state* velocity contribution made it possible that a strong and constant wind during landing could cause the airbags to tear on impact.

The EDL designers understood that, given a steady state horizontal velocity measurement, the TIRS system could also be used to reorient the lander system so that the retro rockets fired in a direction that would compensate for an additional wind-induced horizontal velocity component as well as the RAD-induced velocity. Since the TIRS system was well along in development, the difficulty was in obtaining a horizontal velocity measurement. At this late point in EDL development, adding a traditional radar-based sensor for measuring horizontal velocity was not feasible; a radar interface to the rover avionics was not in place, and accommodation of a multiple antenna radar would be difficult.

The EDL development team realized that due to a descope of the sun camera on the rover, the MER avionics had an unused camera interface and surplus camera electronics modules. Consequently, if there was a way to use a camera to measure horizontal velocity then the impact on the design of the MER avionics would be minimal and it was conceivable that a horizontal velocity measurement system based on images could be developed before launch.

The missing component was an algorithm and flight software that could take in images and other sensor data and output a horizontal velocity measurement. The design, development, testing and flight system implementation of this algorithm are the focus of this paper. This algorithm and the associated descent camera became known as the Descent Image Motion Estimation System

(DIMES). Because it combines multiple sensor measurements with images to estimate horizontal velocity and this velocity is used to control a descent trajectory, DIMES is truly a robotic vision system for planetary landing.

1.1. Design Constraints

DIMES was developed for a flight project that was already well along its development. Consequently, the DIMES algorithm had to be developed in the context of this relatively mature system. As described below, this placed numerous constraints on the design of the DIMES algorithm.

EDL is a very dynamic environment. During descent, the vertical velocity was about 75 m/s, with up to 30 m/s wind-induced horizontal velocity. The MER EDL system was essentially a double pendulum composed of a parachute, backshell containing the retro rockets, and the lander with the rover inside of it. This double pendulum induced pitch rates up to 60°/s in addition to coning and wrist oscillations. Figure 1 shows the dynamics and nominal timeline for the DIMES system. The DIMES sensors were all located on the lander. A Litton LN200 Inertial Measurement Unit (IMU), containing three gyros and three accelerometers, was inside the rover. A Radar Altimeter System, based on a radar altimeter built by Honeywell, and the DIMES camera were attached to a bracket located on the corner of the lander most often oriented towards the surface.

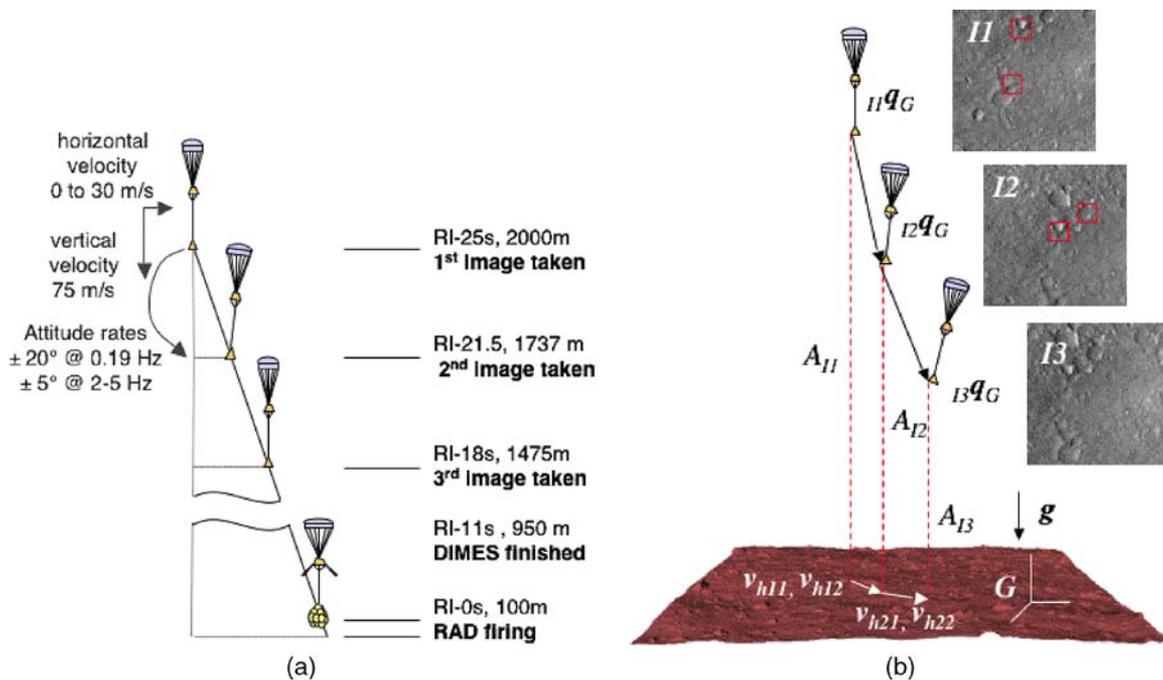


Figure 1. (a) Descent timeline and (b) associate sensor measurements. DIMES takes three images during descent and uses these images along with measurements of surface relative attitude (q) and altitude (A) to estimate horizontal velocity.

It was possible to add a descent camera to the MER EDL system because camera electronics (Maki et al., 2003) had already been developed for a recently de-scoped rover sun sensing camera. In addition, there were surplus optics from the rover cameras available with 16°, 45° and 120° fields of view (Maki et al., 2003). Due to time and cost constraints, the 45° field-of-view MER Navcam optics had to be used with its built-in band pass filter removed to increase optical throughput. In addition, it was not possible to make anything but minor changes to component values in the electronics because any redesign and testing would be too time consuming and costly; DIMES had to be made to work with the hardware that was available.

Since the DIMES camera was created by modifying a camera designed for surface operations it was not an ideal camera for the highly dynamic descent environment. To minimize motion blur, images were taken at the shortest exposure time (5 ms), but due to the frame transfer operation of the Charge Coupled Device (CCD), the images also contained an unwanted ramp of intensity caused by additional exposure during the 5 ms of frame transfer and 5 ms of fast flush. This ramp had to be removed before feature tracking.

The DIMES camera CCD was 1024×1024 pixels, but, to minimize readout time, on-chip binning was used to reduce the image to 256×1024 . Even with the reduced size, image readout still took 3.75 s. Given the long time between images, the 45° field of view was barely wide enough to ensure adequate overlap area to guarantee the same feature would appear in consecutive images.

The high dynamics, long time between images, and limited computing power drove the algorithm design to use onboard attitude and altitude measurements for image rectification instead of something more akin to structure from motion (Azarbayejani and Pentland, 1995; Oliensis, 2002).

The MER landing sites, Gusev Crater and Meridiani Planum, were chosen based on landing safety and science return. These sites are generally smooth and featureless, which results in low image contrast (2% to 3% of the image signal) when combined with the camera performance. As described in greater detail later, low image contrast drove our choice of feature selection and matching algorithms.

Given the short development time and the novelty of the approach for spacecraft landing, there was a lot of concern that DIMES could compute an incorrect velocity resulting in failure of a very expensive and high profile Mars mission. Consequently, the probability of generating an incorrect velocity estimate was required to be less than 0.1%. To compensate for this strict requirement, DIMES was allowed to not report a velocity 10% of the time. DIMES achieved this level of robustness by tracking multiple features and applying checks on feature cor-

relation and consistency of velocities across two image pairs.

The DIMES algorithm had to be implemented in C under VxWorks on a 20 Mhz RAD6000 flight processor. DIMES was allocated 40% of the processor over 20 s, which provided only 160 million instructions to do the job. This severe limitation on available computation drove many of the optimizations in the algorithm.

Given these numerous constraints, the algorithm that was ultimately designed worked as follows. The descent camera takes three images during descent at roughly 2000 m, 1700 m, and 1400 m above the surface. The DIMES algorithm starts by tracking two features between the first and second images and two features between the second and third images. To enable use of a 2D correlator, templates and windows for tracking are rotated and scaled using onboard measurements of lander attitude and altitude relative to the surface. Lander attitude is generated by propagating an inertial star-referenced attitude from prior to atmospheric entry down to the surface using attitude rate data supplied IMU. Lander altitude above the surface is measured by a wide-beam first-return Honeywell radar altimeter. The feature tracks provide estimates of the average velocity between images. If a valid velocity is computed it is propagated using IMU data down to thruster firing, which occurs at approximately 100m altitude. Figure 1 shows the descent time line and the measurements available for processing from the EDL system.

1.2. Related Work

The DIMES algorithm was created by combining algorithmic components from multiple ongoing research tasks focused on vision based approaches to safe and precise landing (Cheng et al., 2001; Johnson and Matthies, 1999; Roumeliotis et al., 2002). In Johnson and Matthies (1999), it was shown that image-correlation and structure from motion (Azarbayejani and Pentland, 1995; Oliensis, 2002) when combined with an altimeter measurement can provide full six degree of freedom motion estimates for landing on comets and asteroids. In Cheng et al. (2001), it was shown that the homography (Hartley and Zisserman, 2003) is the appropriate transformation for representing image motion during planetary landing because the scenes are typically flat and far from the camera. Finally in Roumeliotis et al. (2002), it was demonstrated that image feature tracking and inertial measurements are complimentary and that when combined in a single filter, highly accurate, robust and real-time measurements of vehicle motion are possible. During development we took concepts and in some cases code from all of these approaches to develop the DIMES algorithm.

A platform that has proven particularly useful in closed loop demonstration of safe and precise landing

algorithms has been an autonomous helicopter developed at JPL. The approaches used to estimate the motion of a helicopter for landing and autonomous flight can be applied to the problem of planetary landing. Describing work in this area is relevant to the development of DIMES. There are numerous examples of autonomous helicopter systems using vision for both state estimation and control. In Bosse et al. (1997), optical flow-based motion estimates are combined in an Extended Kalman filter along with IMU, GPS, and sonar altimeter measurements to provide a navigation solution for an autonomous helicopter. The use of optical flow however is restrictive since it is reliable only in domains where the motion between images is expected to be small. Amidi et al. (1999) present a visual odometer which estimates the position and velocity of a helicopter by visually locking on to and tracking ground features. Attitude information is provided by a set of gyroscopes while position and velocity are estimated based upon template matching from sequences of stereo vision data. In this approach, it is assumed that the field of view changes slowly while the helicopter hovers above the same area. New templates are acquired only when the previous ones disappear from the scene. In Corke (2004), stereo vision for height and velocity estimation is combined with IMU measurements in an EKF/CF (Complimentary filter). The goal is to minimize the use of GPS and instead rely on velocity estimates from vision for control. Various approaches for vision-based control for autonomous landing are described in Saripalli et al. (2003) and Shakernia et al. (2002). Both approaches describe and demonstrate a vision-based approach for locating a known target and then tracking it while navigating to and landing on the target. However, in these two approaches, the target area is known *a priori* to be flat and safe. Recently, JPL has demonstrated autonomous landing in unstructured terrain using visual inputs to estimate helicopter motion and scene structure (Johnson et al., 2005a).

In addition to MER, there have been some recent space science missions that are using imagery for spacecraft control. As in the case of DIMES, the resulting systems are highly tailored to operational scenario and spacecraft capabilities. In 2000, the NASA Near Earth Asteroid Rendezvous Mission (Williams, 2003) used imagery to orbit and touchdown on the surface of Eros, but all operations were manual. In July 2005, the NASA Deep Impact mission (Bank et al., 2001) successfully impacted a comet at high velocity with a penetrator spacecraft while another spacecraft imaged the impact site as it passed by. The targeting required closed loop image-based control using autonomous centroiding of the entire comet nucleus. In 2006, the Japanese MUSES-C mission will attempt to return a sample from an asteroid. The terminal control for this mission is performed by placing a known

marker on the surface of the asteroid (Kubota et al., 1999).

1.3. Paper Overview

This paper is organized as follows. In Section 2, the DIMES algorithm is described including many of the optimizations required to fit the algorithm into the processing budget. In Section 3, we describe validation of the performance of the algorithm using Monte Carlo simulation with a high fidelity descent image simulator and field testing with an engineering model camera and IMU. In Section 4, we then describe how DIMES performed for the two Mars landings in January 2004 including a comparison of a flight descent image to one produced by our simulator. In Section 5 we conclude the paper with a short description of future research directions. Portions of this paper have appeared in various peer reviewed conference papers (Cheng et al., 2005; Johnson et al., 2005b; Willson et al., 2005a).

2. Algorithm

In the DIMES system, we assume the lander state (attitude and altitude) is correctly measured and the terrain surface is flat and level. DIMES takes three descent images as input; for each descent image, DIMES also requires elements of the lander state and the vector to the time of image exposure. The descent images are transformed (scaled and rotated) into the same image coordinate system using the lander state information. Two features are matched for each transformed image pair using image correlation. These four matched features are then used to estimate the horizontal velocity for each image pair; the difference between these two velocities is also compared to the acceleration provided by the IMU to ensure a correct answer.

2.1. Two Image Algorithm

With very limited computing power and a very short timeline, the DIMES algorithm was fit into the flight system using four main strategies:

- Minimize data to be processed: It was clear that the slow onboard computer would not be able to accomplish the processing in time if each descent image was processed in its entirety. Therefore, feature search windows were made just big enough for tracking.
- Consolidate processes: To apply 2-D correlation between two descent images, radial distortion must be removed and each image must be rectified to the local level frame. By taking advantage of the moderate

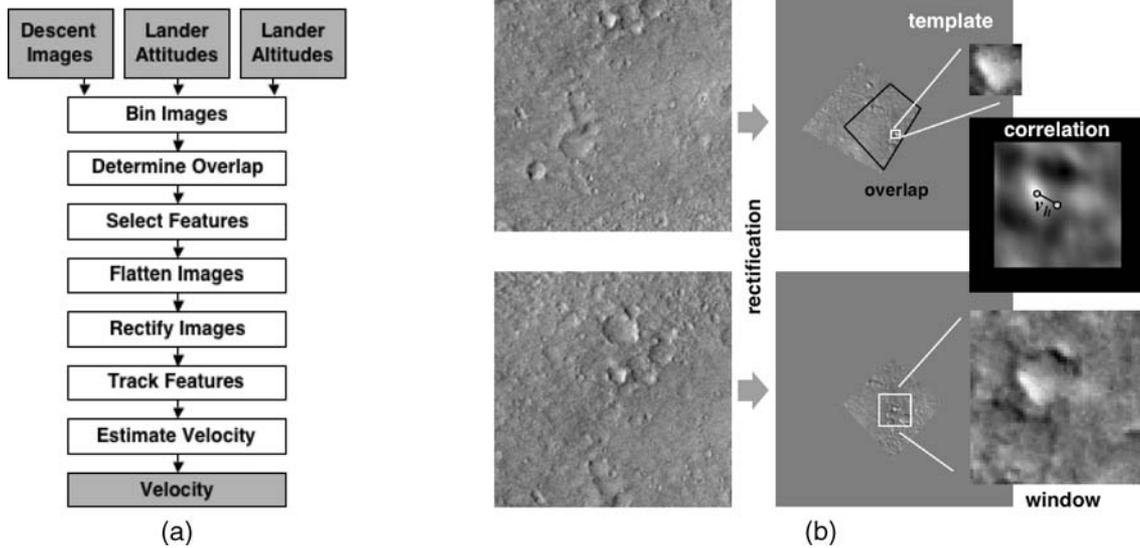


Figure 2. Two image single template horizontal velocity estimation algorithm. (a) Block diagram. (b) Algorithm process.

field of view and very small radial distortion, it is possible to consolidate these two processes into a single homography-based rectification.

- Simplify the algorithm: Some time consuming computations, such as the Harris operator and image rectification, were simplified.
- Use a multi-resolution approach: Multi-resolution was used for correlation matching and feature selection. Both implementations yielded a large time savings.

A block diagram and schematic that illustrate how velocity is estimated from two images is given in Fig. 2.

2.1.1. Image Binning. The descent camera was programmed to use 4×1 on-chip row binning to reduce the time of image readout. Consequently, the images read by DIMES are 256×1024 pixels. When DIMES obtains an image, it further bins it in software from 256×1024 to 256×256 images using 1×4 summing of the 12 bit pixels. The purpose of software binning is to make square pixels and reduce computation.

2.1.2. Zero Phase Masking. When viewing the ground from above, two potentially troublesome effects occur near the point on the ground opposite the direction of the sun: the shadow of the observer and the zero-phase brightening around the shadow (the opposition effect) (Hapke, 1986). The shadow moves with the observer, so it must be avoided to prevent its being tracked and used to compute an erroneous velocity. The opposition effect causes a non-linear increase in the image intensities near the shadow, so it can also hinder feature tracking. Fortunately, the location and angular size of the shadow and opposition effect can be com-

puted and masked out of the data to remove them from consideration.

The camera is placed above the surface using the altitude measurement and then the camera is oriented using the surface relative attitude measurement. The sun ray vector is then projected from the camera to compute its intersection with the ground. Finally the ground intersection is projected into the image to determine the pixel location of the shadow and surrounding opposition effect. A user-defined parameter (4.3°), based on the worst case attitude error and a sensitivity analysis of the algorithm performance in the presence of a model of opposition effect brightening, is then used to mask out a circular region around the shadow. This procedure is repeated for each image.

2.1.3. Setting the Search Window Size. The search window size in pixels depends on five factors: altitude (H), the bounds on the absolute and relative attitude measurement errors (e_a, e_r), the maximum expected horizontal velocity (v_m), the time interval between two descent images (dt), and the template size w_t . The largest contributing factor is the horizontal displacement (Fig. 3(a))

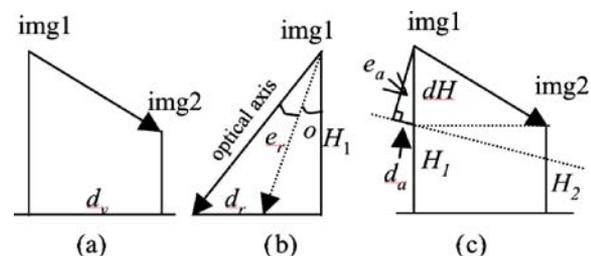


Figure 3. Geometry for setting the search window size.

between two images due to the velocity:

$$d_v = v_m dt.$$

The other two horizontal displacements are relatively small. The relative attitude measurement error between two images causes some horizontal displacement, which can be computed as the displacement of the intersection between the optical axis and ground as

$$d_r = H(\tan(o + e_r) - \tan(o))$$

where o is the off nadir angle (Fig. 3(b)). Finally, the absolute attitude error can be visualized as tilting the imaginary ground surface during rectification (Fig. 3(c)), which causes the vertical velocity to bleed into the horizontal velocity. This effect can be computed as

$$d_a = (H_1 - H_2) \sin(e_a)$$

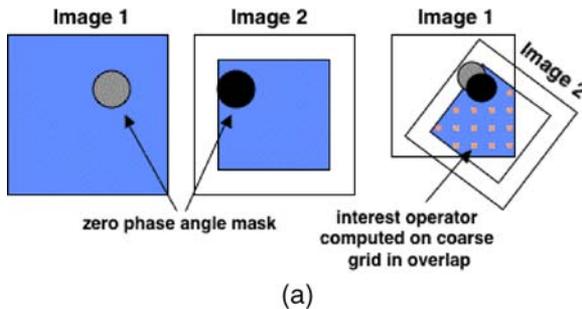
Since these effects are independent, the total horizontal displacement is computed by taking the RSS of d_v , d_r and d_e .

$$d_i = \sqrt{d_v^2 + d_r^2 + d_a^2}$$

The half window size in pixels (w_h) is then the difference in pixels between the template center and the reprojection (into the image) of the template point on the ground shifted horizontally by d_i . The total search window size is then $2w_h + w_t$.

2.1.4. Feature Selection. Features selected for matching must satisfy four criteria (Fig. 4(a)):

- They must be inside the image overlap.
- They must be away from the opposition effects in both images.
- They must have enough intensity variation to be trackable.



- The two features selected in each image pair must be far apart to ensure independent estimation.

To meet these criteria, the Harris interest operator is applied (Harris and Stevens, 1988) to the first image inside the overlap region between the images and outside the opposition effect masks. To reduce computation, feature selection is applied on a coarse grid (64×64) in the first image. Overlap is determined as follows. Because feature selection does not require high geometric accuracy, radial lens distortion can be ignored, and, since the ground is flat, the relation between the two images can be approximated by a 3×3 homography matrix h . If $(r0, c0)$ is a pixel in the first image and $(r1, c1)$ is a pixel in the second image then

$$[r1, c1, 1]^T = h \cdot [r0, c0, 1]^T.$$

h depends on the change in position and orientation between the images, the parameters of the plane describing the scene and the intrinsic parameters of the camera. From the onboard measurements, the change in orientation (provided by the IMU) and the change in vertical position (from altimeter measurements) are known. The valid assumption that the ground plane is flat and level provides the plane parameters. A camera model describing camera intrinsic parameters is available from pre-flight camera calibration (Gennery, 2001). The final unknown in the horizontal translation between images. Since there are obviously no measurements for this quantity, it is assumed to be zero and instead, the overlap region is shrunk by the maximum expected horizontal motion between images to account for this assumption.

h is computed through with a linear estimation procedure. Four 3D points on the ground plane are projected into each image using the geometric information provided by the assumptions in the paragraph above. This results in four pairs of corresponding pixels $(r0_i, c0_i)$ and $(r1_i, c1_i)$. We then solve the linear equation $Am = b$ for

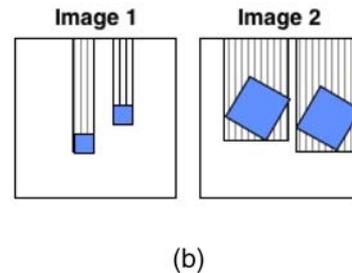


Figure 4. (a) Features have to be selected inside the overlap between two descent images. The dark circles indicate the opposition effects masks. The shaded areas in the left two images are the potential area for feature selection. After overlapping two images, only a small area is left for feature selection (blue area in right figure). (b) The frame transfer smear removal is only applied to a strip that covers the templates or search windows.

m where

$$A = \begin{bmatrix} A_0 \\ A_1 \\ A_2 \\ A_3 \end{bmatrix},$$

$$A_i = \begin{bmatrix} r0_i & c0_i & 1 & 0 & 0 & 0 & -r0_i r1_i & -r1_i c0_i \\ 0 & 0 & 0 & r1_i & c1_i & 1 & -r0_i c1_i & -c0_i c1_i \end{bmatrix},$$

$$b = \begin{bmatrix} b_0 \\ b_1 \\ b_2 \\ b_3 \end{bmatrix}, \quad b_i = \begin{bmatrix} r1_i \\ c1_i \end{bmatrix}.$$

h is then determined from m

$$h = \begin{bmatrix} m_0 & m_1 & m_2 \\ m_3 & m_4 & m_5 \\ m_6 & m_7 & 1 \end{bmatrix}.$$

Given h , the zero-phase mask for the second image is transformed into the first image and expanded by the size of the search window to allow for the unknown translation between the two images. Then each pixel from the 64×64 coarse grid in the first image is checked to see if it is outside of the zero-phase masks from both images. If it passes that test, it is transformed into the second image using the homography h ; if the transformed pixel is inside the second image bounds (shrunk by the search window size), then it is feasible for tracking and the Harris interest operator is applied to intensities in the first image.

The pixel with highest interest value is selected as the first feature. The second feature is the one with highest interest value outside the search window of the first feature.

To refine the selection made on the coarse grid, the Harris interest operator is applied to a 9 by 9 region centered at each selected feature. The final feature location is set to the pixel with the highest interest value in the region.

2.1.5. Image Flattening. Before rectification, the templates and window intensities are corrected for image-to-image variations so that image correlation has the best chance for success on the low contrast Martian terrain.

The DIMES camera did not have a shutter, so the CCD collected charge during the transfer of the frame off the CCD after exposure and during the flushing of the charge from the CCD before exposure. Because the exposure time is the same duration as the frame transfer and fast flush times, a ramp of intensity is visible in the descent images. Using the assumption that the camera is not moving

during exposure, this intensity ramp can be removed using the image data alone and the formula

$$P'[r][c] = f * P[r][c] + (1 - f) * P'[r - 1][c]$$

where P is the pixel intensity, $f = f_i/e_i/r$, f_i is frame transfer plus fast flush time (5 ms + 5 ms = 10 ms), e_i is the exposure time (5 ms), and r is number of rows of the image (256). In order to reduce the computation, only a strip of image, which covers the search window or template is processed (Fig. 4(b)).

After frame transfer ramp removal, gain correction is applied to each template and window (not the whole image) to reduce the intensity differences between images. The term ‘‘gain correction’’ as used here includes subtraction of the dark current and multiplication by a scale factor to correct for CCD pixel-to-pixel gain variation and optical transfer (vignetting). Gain correction cannot eliminate photometric differences between the images (e.g. zero phase angle brightening, lander parachute shadow).

Gain correction applies a scale and offset to each pixel. To reduce non-volatile memory requirements, the scale and offset images are represented as biquadratic polynomials whose coefficients are stored instead of entire images. Prior to DIMES image acquisition, the scale and offset images are precomputed from the stored coefficients. Since the order of the polynomials is low, they cannot account for pixel-to-pixel CCD variation. However, testing with representative CCD pixel-to-pixel variation has shown that not accounting for it with the scale and offset images does not affect DIMES performance. Furthermore, algorithm testing has revealed that the offset images produced by radiometric calibration of the flight cameras have little effect on DIMES results because the dark current is small for these short exposures, cold temperatures, and bright scenes. The velocity error and number of valid solutions are very similar regardless of whether or not the offset image is used. Consequently, the offset image was not computed or used during the MER landings (turned off by parameter).

2.1.6. Image Rectification. The templates and corresponding search windows for the two features are rectified to a camera frame parallel to the ground and at the altitude when the first image was taken. Again, a local homography is used for this procedure. The local homography serves two purposes. First, it provides a simple relationship between the rectified plane and the image plane. Second, it compensates for the local radial distortion. For the DIMES camera, the distortion is less than 0.03% of the image height (Smith et al., 2001), which means the pixel displacement caused by radial distortion at the corners of a 256 by 256 image is less than 0.054 pixel. Therefore radial distortion was ignored and a homography transform was used directly on the raw images.

The local homography transform is obtained by a numerical solution. The pixel coordinates of four corners of the template or search window are projected from the descent image into the local level frame using the intrinsic camera parameters encoded in a camera model and the camera attitude and altitude. The four pairs of image and local level pixels points are then used to determine the 8 coefficients of the homography transform as described in Section 2.1.4.

In order to reduce computation during rectification, the homography transform is implemented as a rolling sum. First it is decomposed into three components:

$$x_1 = \frac{A}{C} \quad y_1 = \frac{B}{C} \quad \text{where} \quad \begin{aligned} A &= a_1x + a_2y + a_3 \\ B &= a_4x + a_5y + a_6 \\ C &= a_7x + a_8y + 1 \end{aligned}$$

If the current pixel's transform has been calculated, then the right adjacent pixel's transform can be reduced to three additions and two divisions as:

$$\begin{aligned} A_{i+1} &= A_i + a_1, & x_{1i+1} &= \frac{A_{i+1}}{C_{i+1}} \\ B_{i+1} &= B_i + a_4, & & \\ C_{i+1} &= C_i + a_7, & y_{1i+1} &= \frac{B_{i+1}}{C_{i+1}} \end{aligned}$$

The pixel's intensity is then calculated using bilinear interpolation.

2.1.7. Feature Tracking and Velocity Estimation. The result of rectification is a pair of templates from the first image and their corresponding search windows from the second image mapped into a single view of the surface. The only difference between templates and windows is a horizontal shift induced by the unknown horizontal motion. A classical pseudo-normalized correlation algorithm (Moravec, 1977) is used to determine this horizontal motion. The pixel with highest correlation is the template's most likely location in the window. A two level block average image pyramid is used to speed up the process. First, correlation is performed using 2×2 binned templates and windows. This result seeds correlation at the finer level between the original template and a 5×5 sub-window around the coarse correlation peak.

Biquadratic interpolation of correlation scores is applied to obtain subpixel match locations. The biquadratic form is also used for calculation of peak width and peak ratio, which are used for match validation.

To determine the correlation peak width, the biquadratic used for subpixel interpolation is analyzed. Fitting a biquadratic to a 3×3 neighborhood of correlation scores yields a biquadratic surface:

$$C = ax^2 + by^2 + cxy + dx + ey + f$$

The subpixel maximum of the fitted surface is then found by partial differential, which leads to the following solution for the peak pixel (x_p, y_p) :

$$\begin{aligned} x_p &= (-2bd + ce)/(4ab - c^2) \\ y_p &= (-2ae + cd)/(4ab - c^2) \end{aligned}$$

To compute the maximum width of the correlation peak a change of variables to coordinates (x', y') around the peak is applied to the biquadratic form.

$$C = ax'^2 + cx'y' + by'^2 + f'$$

Let $x' = r \cos t$, $y' = r \sin t$, then

$$C = ar^2 \cos^2 t + cr^2 \cos t \sin t + br^2 \sin^2 t + f'$$

Given a constant C , the curve's maximum is at

$$t_0 = \frac{1}{2} \arctan\left(\frac{c}{b-a}\right)$$

Therefore, the maximum peak width w can be calculated as

$$w = [(c - f')/(a \cos^2 t_0 + c \cos t_0 \sin t_0 + b \sin^2 t_0)]^{1/2}$$

As described below, this width is used to indicate poor correlations.

After image correlation, the center of the template and its peak correlation pixel are projected onto the ground. The horizontal velocity is then

$$\begin{aligned} v_x &= (X_C - X_T)/dt \\ v_y &= (Y_C - Y_T)/dt \end{aligned}$$

where (X_C, Y_C) and (X_T, Y_T) are the template center and peak ground points, respectively, and dt is the time interval between them.

2.1.8. Match Verification. To ensure a correct match, multiple checks are made on the intermediate and final results. The Harris interest operator value for a feature is the template contrast metric. A low contrast metric indicates little or no contrast and consequently a template that is susceptible to image noise. The correlation peak value is another metric that indicates how similar the template is to its best match location in the search window.

Two other metrics are also used. The first is the correlation peak width, which is a very good indicator of the local frequency content of the window. When low frequency dominates in a template or along a certain direction, such as along an edge, the peak width will be large and the match is not reliable. The correlation peak ratio is the ratio of the highest correlation peak to the

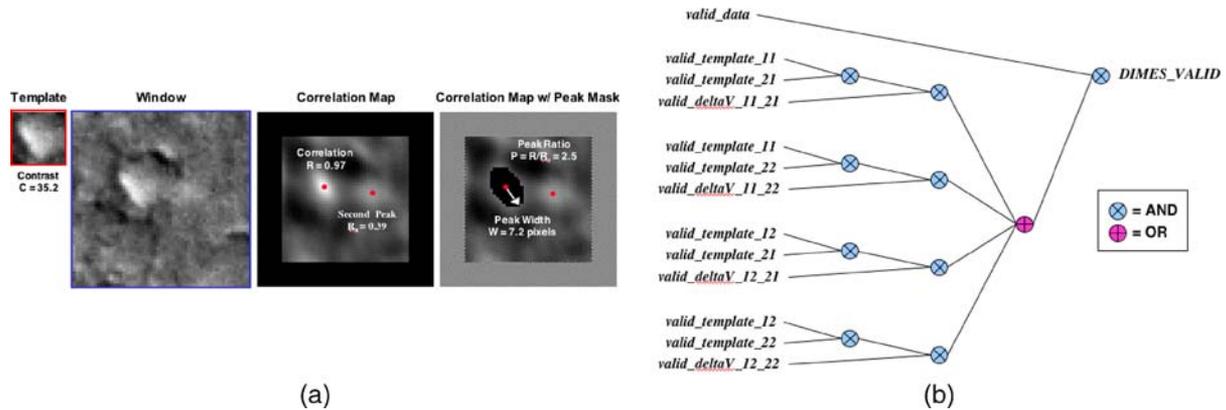


Figure 5. (a) Image correlation metrics and (b) validity checking logic. Valid_template_ij indicates that the j th template from the i th image passed all of its correlation checks. valid_deltaV_1a_2b indicates that the deltaV computed between the a th template from the first image and the b th template from the second image agree with the imu deltaV. valid_data indicates that all of the sensors are working nominally.

second highest correlation peak. This metric is useful for detecting repetitive terrain or multiple features of similar appearance, such as craters. Figure 5 depicts these correlation metrics.

Furthermore, the algorithm will not produce any result when it encounters any one of three additional anomalies. The first anomaly is that the matched point is on the border of the window, where subpixel interpolation is not feasible. The second anomaly is that the peak is too flat (i.e., where $(4ab - c^2)$ is close to zero). Finally, if the subpixel correction is greater than 1.5 pixels, the matched result will be rejected.

2.2. Robust 3-Image Velocity Estimation

The DIMES algorithm uses three descent images. Two templates are tracked between the first and second image and two templates are tracked between the second and third images. This results in four velocity measurements which makes the DIMES algorithm extremely robust in the presence of off nominal effects like dust on the lens, bad pixels in the CCD and the appearance of the heat shield in the field of view. One of the templates from each pair can fail correlation and DIMES can still compute a velocity. Also, as described below, using two image pairs allows for a mechanism to check the image velocity measurements using the completely independent measurements from the IMU.

Although the IMU does not have enough accuracy to measure horizontal velocity, it is very good at measuring changes in velocity over short periods of time. This fact is used to extract from the IMU data a measurement of delta velocity between the first and second image pair. By taking the difference of the velocity computed from a template in the first image pair and a velocity computed from the second image pair, a image-based delta velocity can

also be generated. The image-based and IMU delta velocities should be close to each other. If they are not then one of the templates used to compute the image-based delta velocity has been tracked incorrectly. If a combination of templates from the first and second image pair generate a delta velocity that matches the IMU (within a velocity threshold established through Monte Carlo simulation), then the DIMES algorithm reports a velocity. Otherwise it reports that velocity estimation was unsuccessful. The logic diagram to determine when the DIMES algorithm reports a valid velocity is given in Fig. 5(b).

3. Pre-Flight Testing

Validation of the performance of DIMES was critical to prove that DIMES would “do no harm” during EDL. Because the entire DIMES flight system could not be tested completely in a realistic flight like environment, the validation tests were broken into three categories. Monte Carlo simulation provided velocity accuracy statistics. Field testing proved that the camera and algorithm would produce reasonable velocity estimates when imaging Mars-like terrain at representative altitudes. Finally, flight system testing proved that the flight software worked on the flight system and that the DIMES velocity answer was available in time. The next sections describe the simulation and field testing approaches used to validate DIMES performance.

3.1. Monte Carlo Simulation

Camera simulators can generate images for developing algorithms, tuning parameters and evaluating the performance of machine vision systems where it is too expensive or not possible to acquire actual test images. Our camera simulator is called MOC2DIMES because

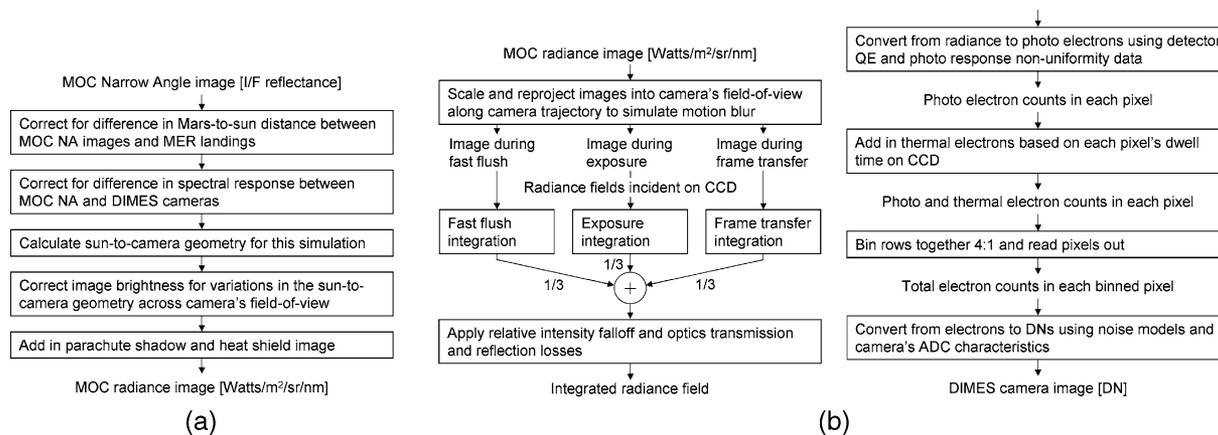


Figure 6. MOC2DIMES (a) scene and (b) camera modeling flow. DN are Data Numbers or greyscale values.

it uses the appearance of the Mars surface in Mars Orbital Camera (MOC) images to generate simulated DIMES descent images. MOC2DIMES uses as inputs images from orbit of candidate Mars landing sites, lander descent profiles and operating conditions, fixed camera parameters and calibration data, and produces geometrically and radiometrically accurate descent camera images. The images generated include a full range of descent dynamics, planetary terrain effects, and non-ideal camera behaviors. MOC2DIMES enabled realistic testing and tuning of the DIMES from early proof-of-concept through development and final landing site selection. This section first describes the MOC2DIMES camera simulator and then describes how it was used to quantify the performance of DIMES using the actual appearance of the MER landing sites. In Section 4.3 it is shown that simulated images are comparable to those collected in flight thereby validating the MOC2DIMES image simulator.

3.1.1. MOC2DIMES Descent Image Simulator. The DIMES camera has a 1024×1024 pixel frame transfer CCD, a 45° FOV and a spectral response from 400–1100 nm (Maki et al., 2003). The DIMES images are taken between 2000 m and 1500 m altitude. In contrast, the Mars Global Surveyor Mars Orbiter Camera Narrow Angle (MOC-NA) instrument is a push-broom line-array camera that takes images of Mars from a spacecraft orbiting Mars at 400 km. MOC-NA has a 0.4° FOV across a 2048 element line array and it has a spectral sensitivity from 500–900 nm (Malin et al., 1992).

Surprisingly, MOC image resolution and photometric conditions are well matched to the MER descent imaging scenario. Specifically, there exist MOC images of the MER landing with incidence and phase angles less than 60° , emission angle less than 30° , spatial resolution ~ 3 m/pixel, and cross-track widths around 3 km. This

compares well to the expected DIMES image parameters of 1–5 m/pixel resolution and 18° – 33° incidence angle. Each MOC image is at least as wide as a DIMES image projected on the ground, so a DIMES image triplet could be rendered from a single MOC image. Furthermore, during the period leading up to the landing, the MOC took images of the MER landing sites, so the images used in MOC2DIMES sampled the appearance of the actual landing sites.

MOC2DIMES projects MOC-NA orbital images into the DIMES camera field of view. As shown on Fig. 6, this modeling process has two stages: scene modeling and camera modeling. Scene modeling (Fig. 6(a)) takes raw MOC-NA reflectance images and converts them into the radiance field the DIMES camera would see while looking down at the terrain. Camera modeling (Fig. 6(b)) converts the radiance field into realistic 12-bit/pixel DIMES camera images.

The specific imaging effects that were modeled in MOC2DIMES are listed below. Further details on each of these effects are described in Willson et al. (2005a).

Scene appearance effects derived through photometric analysis

- Mars terrain albedo and shading from MOC-NA images
- Reflectance as a function of emission and phase angles including the opposition effect (Hapke, 1986)
- Solar illumination intensity for Martian day of year and time of day at landing
- Atmospheric extinction due to dust
- Lander parachute shadow

Descent dynamic effects from multi-body landing simulations (Raiszadeh and Queen, 2004)

- Image scene scale due to lander altitude
- Image orientation due to lander attitude

- Image motion blur due to angular rates and velocities during exposure
- Position and appearance of heat shield (possibly in image due to release shortly before image acquisition)

Camera optical effects determined from calibration data

- Perspective projection with radial lens distortion (Gennery, 2001)
- Optical intensity fall off (vignetting)
- Dust on the lens (Willson et al., 2005b)

Camera electronics effects derived from analysis and calibration data

- Detector sensitivity for conversion from radiance to photo electrons
- CCD frame transfer and fast flush during an open shutter
- CCD pixel random non-uniformity including dead pixels
- Dark current electrons
- Shot noise electrons
- Analog to digital conversion noise

3.1.2. Simulation Results. Simulation inputs were set based on the MER EDL environment. Multi-body aerodynamic simulations of the lander system were used to generate EDL trajectories (Raiszadeh and Queen, 2004). Flight system testing was used to determine the typical altitude of the first image exposure (2000 m) and the time between the second and third images (3.75 s). This information was used to look up the true camera state in each EDL trajectory. Representative values were:

- Altitudes: 2000 m, 1725 m and 1450 m
- Attitude: off nadir angle $< 30^\circ$ and roll $< 45^\circ$
- Angular rates: $< 60^\circ/\text{s}$
- Velocity: vertical ~ 72 m/s and horizontal < 30 m/s
- Position of heatshield relative to lander: 140–200 m axial separation, 50–75 m lateral separation.

Each MOC image is bigger than a set of descent images, so multiple simulation test cases can be created from each by placing the EDL trajectory at different locations across the MOC image. The constraints for placing the trajectories are that each descent image must stay within the horizontal bounds of the MOC image and that the trajectories should evenly sample the appearance of the MOC image so that DIMES performance statistics are not biased. Figure 7 shows an example of image states from three trajectories and associated fields of view on a MOC image; this figure conveys the large attitude and altitude

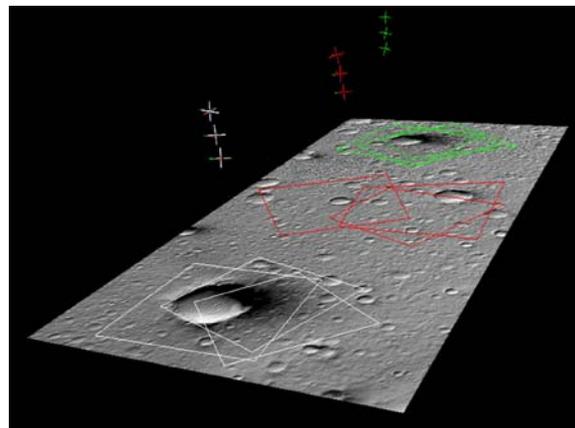


Figure 7. Representative EDL trajectories (+) and DIMES camera fields-of-view (polygonal boxes), projected within a MOC image.

changes between images. In addition to the descent images, the DIMES algorithm also requires estimates of lander attitude, altitude, and biased horizontal velocity. Sensor models were used to generate these measurements, including noise, from the true EDL trajectories that were used to generate the images. Attitude errors, including absolute biases ($\sigma = 1^\circ$), between image random errors due to timing accuracy ($\sigma = 0.05^\circ$), and alignment accuracy ($\sigma = 0.1^\circ$), were based on EDL requirements. Altitude errors were based on the measured performance of the radar altimeter ($\sigma = 0.3\%$ of altitude).

Each MER landing site was imaged by MOC multiple times for landing site selection and hazard assessment. The DIMES simulation used the subset of these images that had photometry similar to that expected during landing. MOC images and associated coverage for each of the landing sites are given in Figs. 8(a) and 9(a).

Terrain appearance has a strong influence on DIMES performance; the number of valid velocities will decrease as the terrain becomes more bland. Both of the MER landing sites had wide variability in appearance. To capture this variability in performance each landing site was manually segmented into “appearance classes” based on orbital images (Figs. 8(b) and 9(b)). Each landing site had three appearance classes which covered different fractions of each landing ellipse.

The Monte Carlo simulation of DIMES performance proceeded as follows for each landing site. First, the camera parameters specific to the flight camera for the landing site (MER-A or MER-B) were input into MOC2DIMES. Next, the MOC images for each appearance class were selected. Then for each MOC image from the appearance class, a set of descent profiles specific to the landing site (the trajectories for Gusev Crater (MER-A) had stronger winds and consequently greater horizontal velocities and attitude excursions) were selected and placed to cover the image. Descent images and sensor measurements were

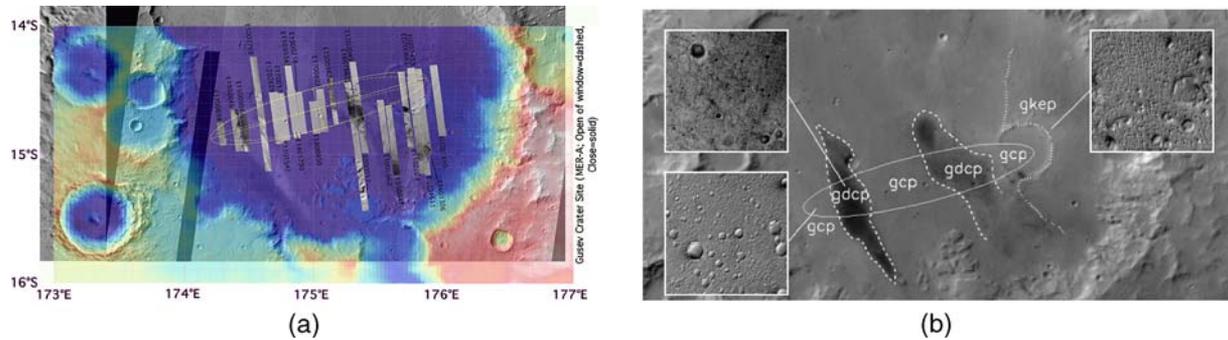


Figure 8. Gusev Crater MOC coverage (52%). (Image courtesy Tim Parker/JPL). Gusev Crater appearance class coverage. gcp: Gusev cratered plains—higher albedo, smooth plains with few craters, bright crater rims, and low contrast overall (ellipse fraction = 59%). gdcp: Gusev dark cratered plains—lower albedo, mostly due to linear dark dust devil tracks, cratered plains area (ellipse fraction = 41%). gkep: Gusev knobby etched plains—knobs or mesas of positive relief dominate this area surrounding crater at east end of ellipse (ellipse fraction = 0%).

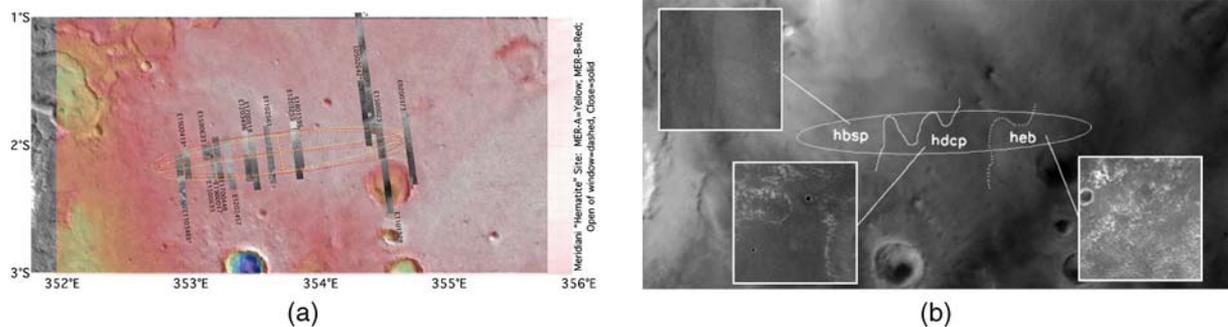


Figure 9. Meridiani Planum MOC coverage (31%). (Image courtesy Tim Parker/JPL). Meridiani Planum Appearance Class Coverage. hbcp: Hematite bright smooth plains—higher albedo, smooth plains with few craters, low contrast (ellipse fraction = 34%). hdcp: Hematite dark cratered plains—lower albedo, cratered plains (ellipse fraction = 40%). heb: Hematite ejecta blanket—ejecta apron surrounding large crater near east end of ellipse (ellipse fraction = 26%).

generated for each trajectory and fed into the DIMES flight software. The velocity estimation results were then added to the results for the current appearance class. This process was repeated for each appearance class. The final result for each landing site was then created by taking an ellipse fraction weighted sum of the results from each appearance class.

Table 1 shows the simulated performance of DIMES at Gusev Crater and Meridiani Planum. The DIMES software parameters were fixed for each landing site, but tuned to Gusev where there is more scene contrast and also stronger winds. These conservative parameters are the reason for the decrease in the valid velocity percentage at Meridiani Planum. If the parameters had been tuned to Meridiani, the valid velocity percentage would have been greater than 90%. The velocity accuracies are for an altitude of approximately 1600 m.

MOC2DIMES was also used to test velocity estimation performance sensitivity to a number of off nominal conditions including high motion blur, excessive dark current and spikes, dust particles on the lens, specular reflections

and blooming off of the heatshield and energetic particle hits.

3.2. Field Testing

Monte Carlo simulation was used to estimate the performance of DIMES at each landing site under realistic EDL dynamics using images of actual Martian landscapes. In Monte Carlo simulation thousands of test cases can be run, so it is important for assessing algorithm performance. However, it cannot replace taking pictures with a real camera at altitude over Mars like terrain. For this field testing is essential. In the summer and fall of 2002, the DIMES team performed a series of field tests in the Mojave Desert. These tests proved that the DIMES algorithm could provide accurate velocity estimates using real images taken at altitude and attitude rates typical of EDL over terrain that was representative of the landing sites.

The field test system consisted of an engineering model (EM) camera and IMU mounted on a 2-axis gimbal

Table 1. (a) Gusev Crater and (b) Meridiani Planum MOC2DIMES performance.

Appearance class	Number of cases	Number of valid cases	% valid	Velocity error (m/s) (mean + 3 sigma)	% landing ellipse	Weighted % valid	Weighted velocity error (m/s) (mean + 3 sigma)
(a)							
gcp	186	182	98	3.90	59	58	2.30
gdcp	112	112	100	3.48	41	41	1.43
gkep	78	77	99	3.29	0	0	0.00
Gusev	376	371				99	3.73
(b)							
hbsp	44	22	50	4.06	34	17	1.38
hdcp	179	125	70	3.92	40	28	1.57
heb	31	31	100	3.03	26	26	0.79
Meridiani	254	178				71	3.74

platform, which was then mounted to a 3-axis stabilized platform attached to the front of a helicopter. Ground support equipment was developed, including a data acquisition and controls system to command the 2-axis gimbal, and to log field test data. A GPS receiver was used to measure and log helicopter position and time data. Figure 10 shows the integrated field test system.

The field test system collected sensor data and ground truth data needed for DIMES validation. After the field test, triplets of images with associated measurements needed by the DIMES flight software were created and a performance analysis using the actual DIMES flight software was conducted. No onboard computation of velocity was performed during the field test.

Field testing was used to validate aspects of the DIMES algorithm that were not covered by Monte Carlo simulation. Specifically, field testing verified the following:

- Performance of the flight-like camera hardware
- Algorithm performance with images containing topography and the associated photometric and viewing effects.
- Algorithm performance with images of increasing resolution and scale.

For field testing to validate flight performance, the field test sensors, dynamics, measurements and environment must be as flight-like as possible. Below we describe the field test design that achieved this goal.

3.2.1. Site Selection. An important aspect of field testing was to replicate Mars like terrain for imaging. The Mars science community was polled for possible Mars analog sites close to Southern California that could be used for testing DIMES. The requirements were that the terrain be free of vegetation in an area large enough to

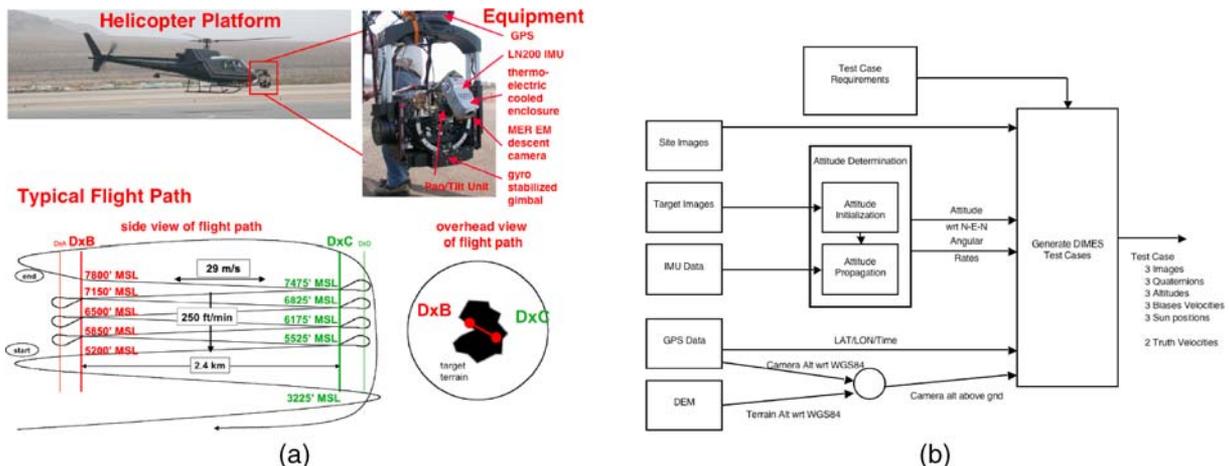


Figure 10. (a) Field test equipment and flight path. (b) Field test data flow.

fit a DIMES image field of view at 2000 m. Also the surface slope and roughness, brightness and native contrast should match that of the MER landing sites. After consensus was reached, the selected sites were Pisgah Lava Flow, Kelso Sand Dunes and Ivanpah Dry Lake Bed. Although none of these is an ideal MER landing site analog, they test DIMES performance over a range of accessible and representative Earth terrains. Further details on the rationale for selecting these sites can be found in Johnson et al. (2005b).

3.2.2. Field Test Sensors. DIMES uses a descent camera, an IMU and a radar altimeter. For the field tests we wanted these sensors to be as close to flight-like as possible. The IMU used in the DIMES field tests was the same model used by the flight system (Litton LN200 IMU). A radar altimeter was not available for the DIMES field tests, so altitude measurements from a flight-like sensor were not available. However, an altitude measurement was derived by taking the difference between GPS altitude and the altitude of the terrain below the helicopter available from digital elevation maps of the test sites. For imaging a MER engineering model descent camera was used. To replicate the camera performance expected on Mars the field test camera's neutral density filter was changed and the camera head was cooled.

3.2.3. Field Test Dynamics. The large vertical velocities close to the ground that are present during EDL cannot be safely achieved by helicopters or planes. The only way to obtain these rates is to drop a system by parachute which is expensive and produces limited data. Before field testing, analysis was conducted that showed that the vertical velocity of the lander had much less of an effect on DIMES performance when compared to altitude and attitude rates. This analysis made it reasonable to use a manned helicopter as the data collection platform thereby simplifying and reducing the cost of the field test system.

A schematic of a typical flight path is shown in Fig. 10(a). The IMU and camera are attached to a pan/tilt unit that is placed inside a gyro-stabilized gimbal on the front of a manned helicopter. The helicopter takes off and flies to the test site while climbing to an altitude of 1000 m. While constantly gaining in altitude, the helicopter flies back and forth over the terrain in a zigzag pattern along a fixed horizontal line; the typical velocities are 30 m/s horizontal and 1 to 2 m/s vertical. During this time, the camera operator points the gimbal to avoid imaging undesirable terrain. The run ends when the helicopter reaches a height of 2000 m. During each one of these runs the pan/tilt can be activated to obtain attitude rates up to $60^\circ/\text{s}$ and off nadir angles up to 45° . Except for vertical velocity, the field test dynamics covered the range dynamics expected during Mars landing.

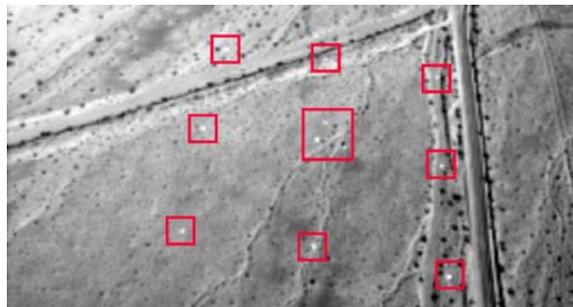


Figure 11. GPS surveyed surface targets for attitude determination.

3.2.4. Generating Flight Like Inputs. After the field test, the collected sensor data were processed to produce altitude, attitude and sun direction for each image. Ground truth position is also available for each image from GPS. After the measurements are determined, DIMES test cases are built from a series of three not necessarily consecutive images. Test cases can be generated from images taken close in time (short time triples), from images taken when the helicopter was flying in the same direction but at different altitudes (same direction triples), and from images taken when the helicopter was flying in opposite directions at different altitudes (opposite direction triples). Each test case consists of three images, the attitude quaternions, altitudes, biased IMU horizontal velocities and the sun vector for each image, and two truth velocity vectors. Figure 10(b) shows the flow of field test data into measurements and the details on the generation of each measurement are given below.

The DIMES algorithm needs a ground relative attitude estimate for each image. To form this attitude reference, eleven white square targets (1 m edges) were placed on the ground (see Fig. 11, for highlighted targets). A detailed GPS survey determined the latitude, longitude and height of each of the targets. The area containing all of the targets was imaged during a target flyover at the beginning and end of each flight; typically about 20 images (<2 min of data collection), were acquired for each fly-over of the targets. Test site data collection occurred during the ~ 40 min between flyovers.

A “inertial” coordinate frame $CF(T_0)$ was established by choosing a time T_0 . Using the Earth's rotation rate, Earth measurements taken at time T can be mapped to $CF(T_0)$ using the transformation $\Phi(T_0 - T)$. Mapping measurements to the inertial frame was important since the gyros only measured changes relative to the inertial coordinate frame—an inertially fixed sensor will see the Earth rotating beneath it. (The Earth's motion about the sun was ignored here.)

Given the helicopter GPS position data in Earth rotating coordinates (vector form), and the transformation $\Phi(T_0 - T)$, we computed the location of the targets relative to the helicopter camera for each time when an

image was taken. This produced a set of unit vectors from the camera position to target location in the reference frame $CF(T_0)$. Finally we extracted the target centroids from the images and computed the location of the targets in camera coordinates which gave corresponding vectors in both camera and $CF(T_0)$ coordinates.

Next, all the target vectors for each image were mapped to a common frame using the gyro data. With the GPS based position vectors in the $CF(T_0)$ frame, we computed an initial attitude estimate at time T_0 by performing a QUEST solution (Shuster and Oh, 1981) on the pairs of vectors from all of the targets in all of the images. This process was repeated for the end of run fly-over giving two attitude sequences in $CF(T_0)$. The gyros biases were assumed to have no drift during each flyover. Assumed constant gyro biases were then estimated by minimizing image centroid to target position match errors between the beginning and end flyovers. Bias estimation improved the attitude estimate and allowed the establishment of an initial attitude estimate using all the target data.

For each of the pictures taken (i.e., the pictures where the targets were not present), the attitude was estimated by propagating the initial attitude using the bias compensated gyro data to the time of the image exposure. The position was directly determined by GPS. Finally all the data was mapped back to the surface fixed frame using the mapping $\Phi(T_0 - T)^{-1} = \Phi(T - T_0)$.

3.2.5. Field Test Results. Field testing occurred in October 2002. A day of flying was spent at each of the test sites and three runs were performed per day at 10am, noon and 2pm.

The first test day was at Pisgah Lava Flow. In Fig. 12 a DIMES result is shown; the first image pair is shown on the left, and the second is shown on the right. The bottom row shows the original un-rectified images with selected templates as red squares and tracked locations as green squares. The top row shows the result of rectifying the images using the image attitude and altitude (note that these images are shown only rectified for display purposes; the DIMES algorithm does not actually rectify the entire image). The correlation window is shown as a blue square. The brightening on the left of the image in the bottom right is due to the opposition effect where the photometric phase angle goes to zero. Excessive noise in the image data, due to a faulty cable, corrupted the data from the first run of the day making it unusable. The noon and 2pm runs were acquired successfully.

On the second day of testing the Kelso Sand Dunes were imaged. A typical DIMES result with repetitive dunes is shown in Fig. 13. Image noise prevented the use of the data from the first run of the day. The noon and 2pm runs were acquired successfully.

On the third and final day of testing, Ivanpah Dry Lake Bed was imaged. A typical DIMES result is shown in Fig. 14. Once again excessive noise in the image data, prevented the use of the data from the first run and second runs of the day. The 2pm run was acquired successfully.

For each of the five successfully acquired runs, image triples were generated resulting in a total of 1913 test cases. During error analysis it was noticed that some of the triples generated much larger errors than others. To investigate this issue, the triples were segmented into 3 categories. The first category contained images that were

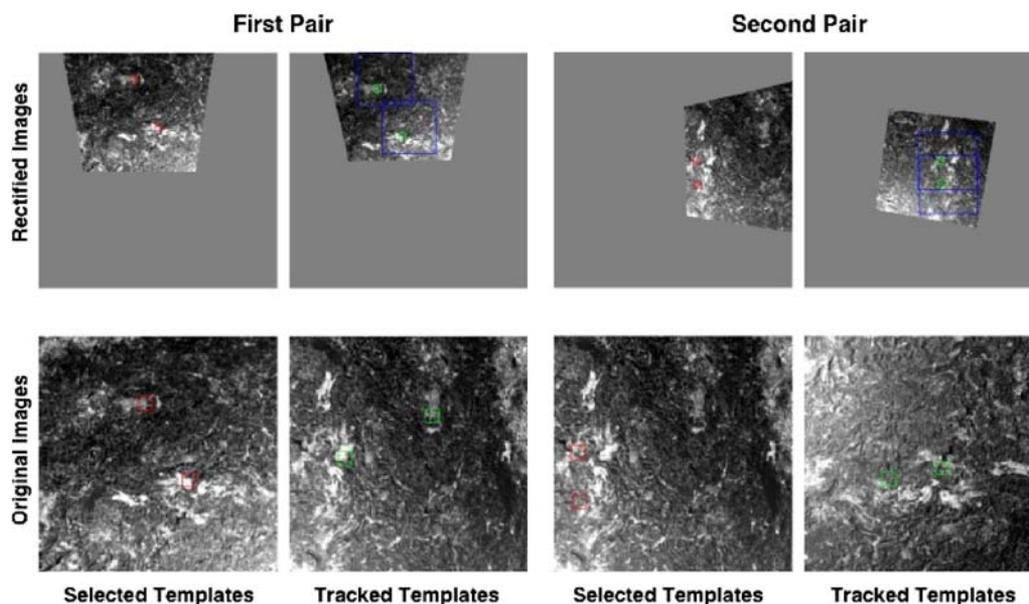


Figure 12. Example DIMES result from Pisgah Lava Flow.

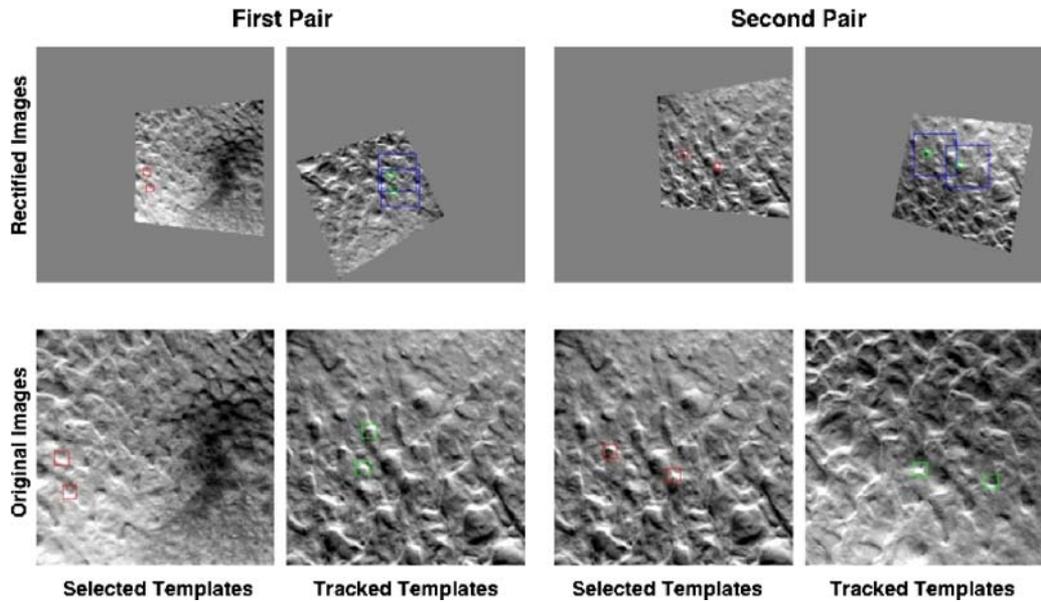


Figure 13. Example DIMES result from Kelso Sand Dunes.

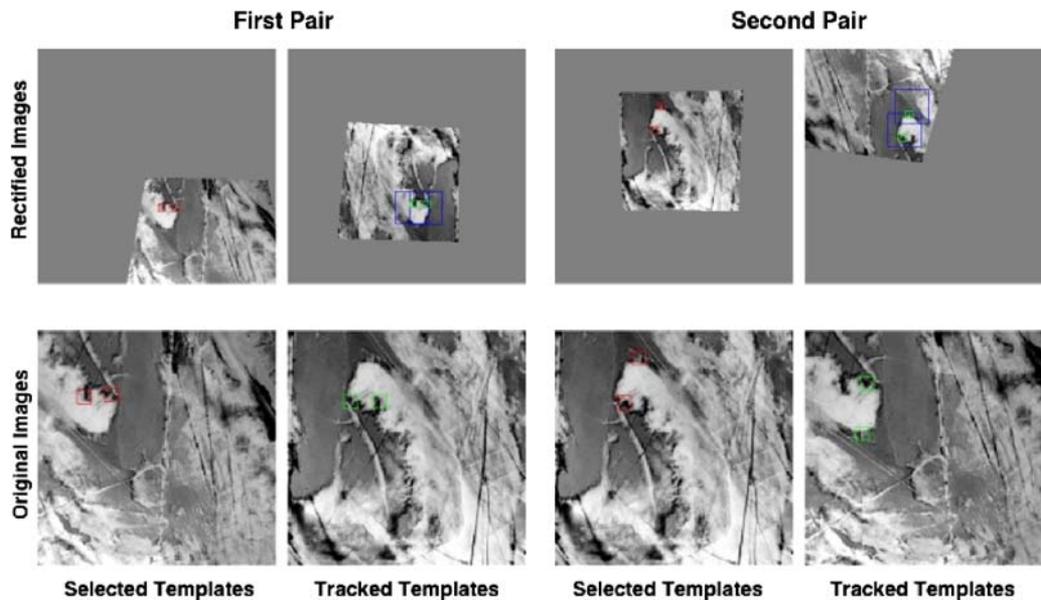


Figure 14. Example DIMES result Ivanpah Dry Lake Bed.

taken close in time and consequently had similar altitudes and were taken when the helicopter was flying in a single direction. These triples are given the label *short time*. To investigate altitude error dependencies a second category was created that contained images where the helicopter flew in the same direction, but the images were separated in altitude like the landing images. The altitude separation forces the images to be separated in time as well. This category was labeled *same direction*. The final category contained images that were separated in altitude (and time) and where the direction of the helicopter travel

switched at least once between images. This category was labeled *opposite direction*.

Velocity errors for all three categories and all five successful test runs are plotted in Fig. 15. For the short time and same direction test cases, the velocity errors are all within the DIMES requirement of 5 m/s. The opposite direction triples have a greater spread and in some cases do not meet the requirement. Also each run is made of one or more clusters. After further analysis, the problem with the opposite direction velocities was attributed to an unmodeled attitude bias across the direction of travel.

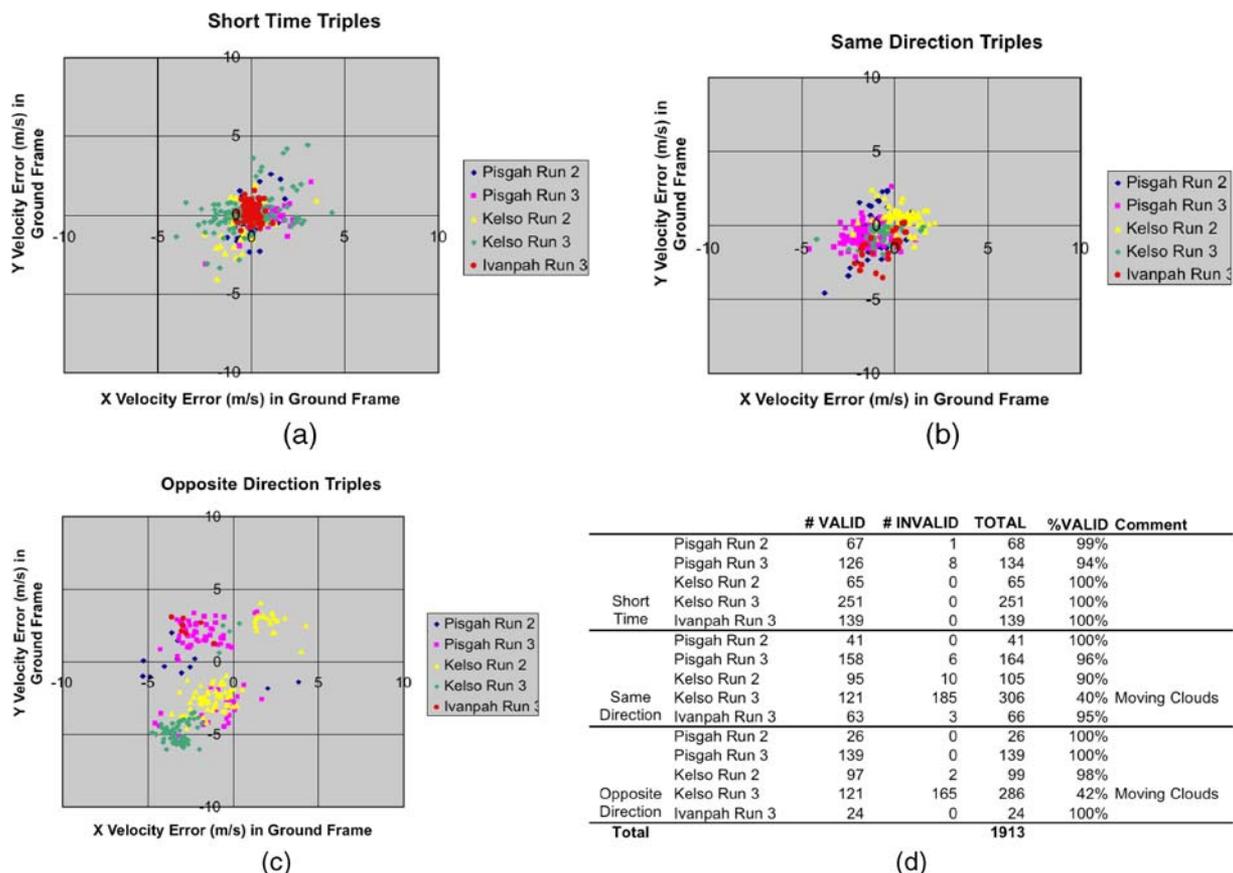


Figure 15. Valid velocity errors for three image triple types: (a) short time, (b) same direction and (c) opposite direction. (d) Field test valid velocity results for three image triple types.

This bias was most likely due to the long time between images and temperature variations in the IMU which negated the constant bias assumption used to determine attitude. Since the temperature of the IMU varies only a little and the images are taken very close together in time this type of bias or effect will not occur in flight. After eliminating the opposite direction triples, all of the remaining triples satisfy the DIMES velocity error requirement of 5 m/s making the field test successful in validating DIMES performance.

The number of test cases and valid velocity results are given in Fig. 15. Except when for the Kelso Run 3, where the shadows of clouds were moving across the terrain, most test cases were valid.

4. Flight Results

The combination of successful Monte Carlo MOC2DIMES testing, field testing and flight system testing indicated that DIMES met all of its requirements and it was sufficiently robust to be used during landing. Consequently, DIMES was enabled for

both landings. During each landing, DIMES computed a valid velocity that was used by the EDL system.

4.1. MER-A/Spirit

On January 4th, 2004, the MER-A (Spirit) spacecraft landed in Gusev Crater, a 160 km diameter crater just south of the Martian Equator. During its decent, it took three images of the Martian surface at altitudes of 1983 m, 1706 m, and 1433 m. These images showed a terrain covered with craters, hollows, and dust devil streaks. During EDL, DIMES computed a velocity of (4.1, 9.7) m/s at 1570 m altitude (all velocities are in the local level (North, East) coordinates). After propagation down to RAD fire, this steady state velocity was (-1.2, 10.7) m/s. Due to a wind gust just before RAD fire, the angle between the lander and backshell was significant at the time of RAD fire, and would have caused an additional horizontal velocity of (-5.7, 11.7) m/s making the without-TIRS total velocity (-6.9, 22.5) m/s. This velocity was significant enough to cause TIRS to fire in its stronger of two modes resulting in a reduced bridle cut velocity of (0.0, 11.0) m/s. The total

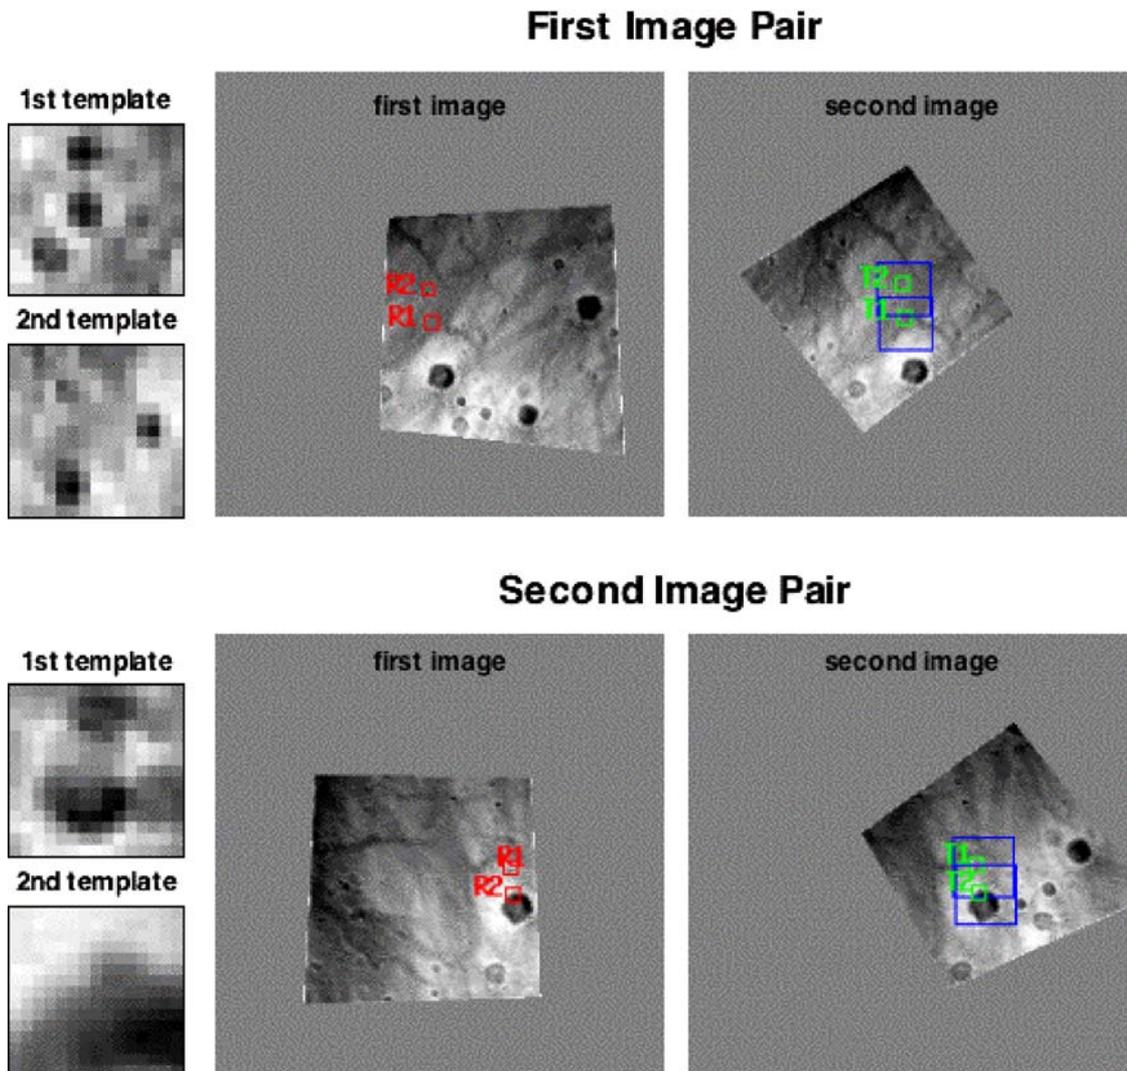


Figure 16. MER-B DIMES results at Gusev Crater.

processing time was 13.75 seconds, which was well within the processing margin.

The RAD-induced velocity was less than the threshold to fire TIRS. Had DIMES not been available to measure the steady state velocity, the EDL system would not have fired TIRS and the total velocity would have been just on the threshold of airbag performance. Furthermore, the velocity would have been to the East toward the rockier terrain surrounding the so-called Bonneville Crater.

The DIMES software tracked all of the features correctly, but threw out the second feature of the second image pair because the peak width parameter for this feature was larger than the maximum allowable peak width parameter (Table 2). Post flight analysis showed that DIMES did compute the correct velocity. Figure 16 shows the images taken by the DIMES camera and the associated features that were tracked.

4.2. MER-B/Opportunity

On January 25th, 2004, the MER-B (Opportunity) spacecraft landed on Meridiani Planum, halfway around Mars from Spirit and just north of the Martian Equator. During its descent, it took three images of the Martian surface at altitudes of 1986 m, 1690 m and 1404 m. These images showed a very bland terrain spotted with a few craters. During EDL, DIMES computed a mostly northern velocity of (8.0, -0.3) m/s at 1547 m altitude. After propagation down to RAD fire, this steady state velocity was still (10.4, -2.8) m/s but in this case the RAD-induced velocity was (-7.7, 3.9) m/s which would effectively negate the effect of the steady state wind. Consequently, TIRS did not fire during the Opportunity landing. The bridle cut velocity was (9.0, -2) m/s, so the lander bounced north, ultimately ending up in the so called Eagle crater. The total processing time was 13.875 seconds.

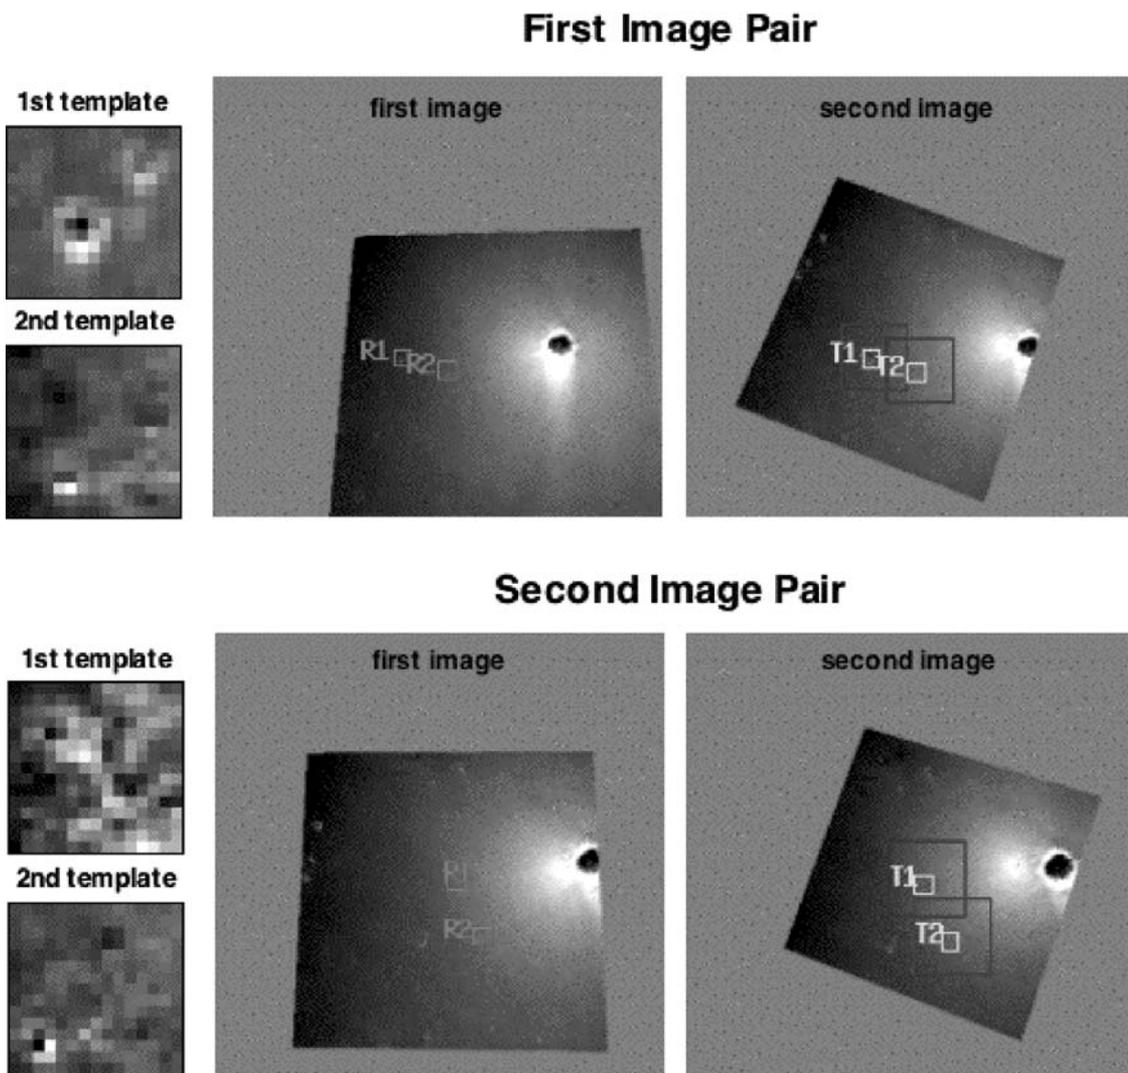


Figure 17. MER-B DIMES results at Meridiani Planum.

Table 2. In flight feature outcomes.

	MER-A				MER-B			
	Pair 1		Pair 2		Pair 1		Pair 2	
	1st	2nd	1st	2nd	1st	2nd	1st	2nd
Valid	Y	Y	Y	N	Y	Y	Y	N
Velocity (m/s)	4.2, 10.4	4.1, 10.6	4.1, 9.7	N/A	6.3, 1.2	6.2, 1.1	8.0, -0.3	N/A
Brightness (dn)	4710	3898	4310	4694	3242	3819	3697	4403
Contrast(dn/pix)	47	42	54	56	31	16	18	17
Correlation	0.88	0.91	0.97	0.997	0.82	0.81	0.76	0.68
Peak ratio	1.3	1.5	1.8	1.1	1.5	1.1	1.4	2.0
Peak width (pix)	2.4	3.3	4.3	39	2.1	4.4	3.3	2.8

The DIMES software tracked all of the features correctly, but threw out the second feature of the second image pair because the correlation for this feature was smaller than the minimum allowable correlation parameter (Table 2). Post flight analysis showed that DIMES did compute the correct velocity. Figure 17 shows the images taken by the DIMES camera and the associated features that were tracked for MER-B.

4.3. Comparison of MOC2DIMES Images to Flight Images

After both landings, the DIMES descent images, state measurements and algorithm results were sent back to Earth. In the case of MER-A, which landed in Gusev Crater, a MOC-NA image that covered the MER-A landing site had already been taken before landing. Using this MOC-NA image, the state measurements and the velocity computed on-board, it was possible create MOC2DIMES images that matched the viewing conditions of the MER-A descent images.

The left column of Fig. 18 shows a simulated raw MOC2DIMES image and the raw image taken during flight (both images scaled between 1200 and 200 DN). The right column of Fig. 18 shows these images after they have been binned to 256×256 and had the radiometric fall-off, fast flush and frame transfer removed. As the figures show, the images are similar in both the raw and flat-

tened versions, but the MOC2DIMES image is brighter and has less contrast than the MER-A image.

Figure 19, which plots a column from each of the raw images, helps explain these differences. Both column plots show the frame transfer ramp. The MOC2DIMES pixels have a higher DN values than the MER-A image while the MER-A pixels vary more indicating a higher level of scene contrast. By plotting a scatter plot of one column vs. the other (Fig. 19, right), it is easy to see that the data are highly correlated (correlation coefficient of 0.99) which indicates that over the entire column, the images are similar up to a constant scale factor indicated by the slope of the scatter plot (slope = 0.85). This 15% difference is quite small given that it encompasses all possible radiometric error sources including MOC image I/F calibration, camera radiometric calibration and scene photometric effects.

The loss of contrast in MOC2DIMES was expected. During field testing, it was discovered that the contrast of the field test imagery was significantly higher than that seen in MOC2DIMES. It was determined that this discrepancy was caused by MOC viewing the Martian surface through the entire dusty atmospheric column as opposed to just 2 km of atmosphere encountered during field testing and flight. Even with the decrease in contrast, the templates selected from the two images are similar (shown on the right in Fig. 18 as red squares) which, while not proving, at least indicates that the images simulated by MOC2DIMES are have similar appearance to

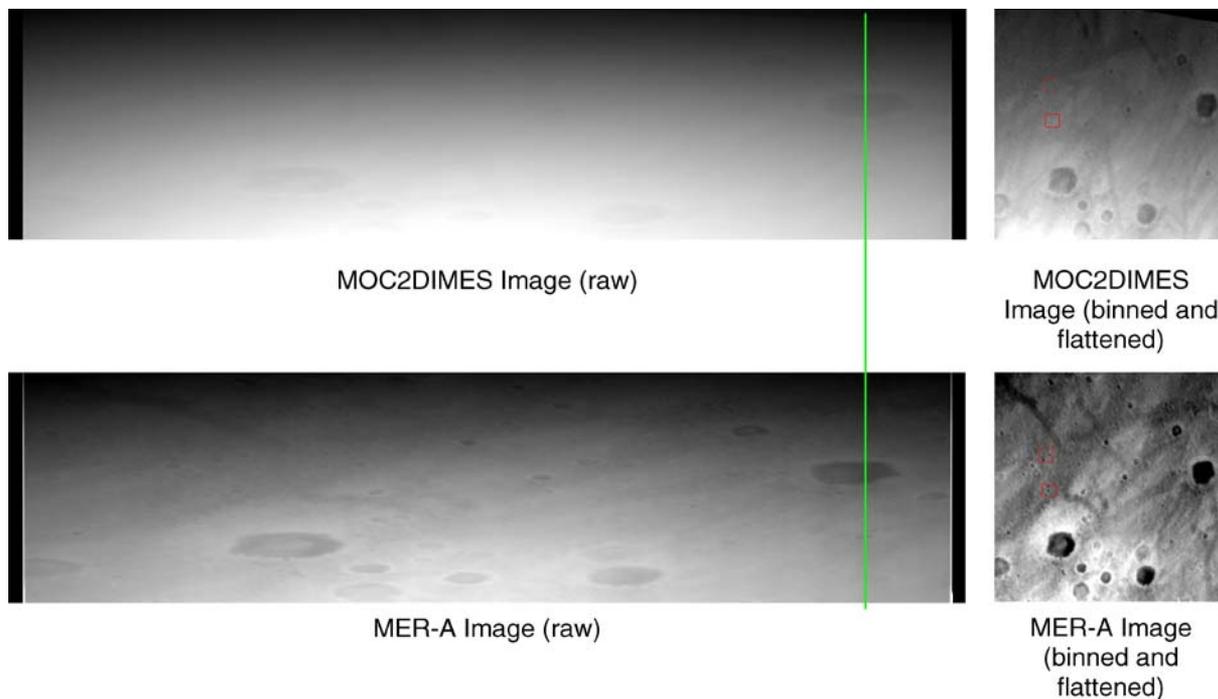


Figure 18. Comparison of MOC2DIMES image of MER-A landing site to actual MER-A descent image. Upper right corner of MOC2DIMES image has no texture because image field of view went outside MOC image. Green line indicates columns plotted in Fig. 19.

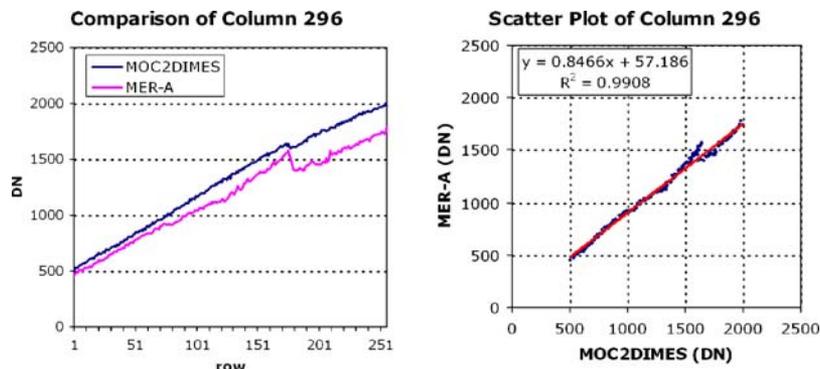


Figure 19. (left) Comparison of column 296 from MOC2DIMES and MER-A image shown in Fig. 18. (right) Scatter plot of columns 296 and associated best fit line and correlation coefficient.

those taken in flight. Further comparisons between flight and MOC2DIMES images are given in Willson et al. (2005a).

5. Conclusions

The Descent Image Motion Estimation System is the first passive image based system to estimate lander velocity during planetary descent. DIMES is composed of sensors and software including a descent imager, a radar altimeter, an inertial measurement unit and an algorithm for combining sensor measurements to estimate horizontal velocity. Although the sensors are not novel technology, the algorithm and flight software that combines them is novel. This algorithm combined radar, image and inertial data in a new way to create a low cost, robust and computationally efficient solution to the horizontal velocity estimation problem. Simulation and field testing both showed that the DIMES algorithm could estimate horizontal velocities that met the requirement of 5 m/s at 1600 m altitude (3 sigma) and this performance was confirmed when DIMES computed the correct velocity for both MER landings.

During development of the DIMES algorithm computational efficiency and robustness to tracking errors were the driving factors behind algorithm design. One way these issues were addressed was by utilizing all available sensor data. For example, the IMU and altimeter provided the information required for image rectification so that 2D correlation could be used for feature matching. If these sensor measurement were not used then, given the large attitude and scale changes between images, a more complicated feature matching procedure would have been required. By enabling 2D correlation these non image measurements made the algorithm much more efficient. The IMU also provided an independent check on the velocity estimate that was used to throw out bad feature matches. This independent check was critical especially

in the event that the DIMES algorithm locked onto and tracked the velocity with respect to the (accelerating) heat shield instead of the surface.

A downside of using imagery with inertial measurements in a very dynamic environment was that the timing knowledge between the IMU and imager had to be very accurate (2.5 ms). This timing knowledge was provided by other lander software modules but the accuracy had to be confirmed during system testing. In addition, errors in the attitude and altitude estimates would leak into the horizontal velocity estimate. This coupling of sensor errors had to be accounted for when generating a total velocity estimation error budget for DIMES. However, this extra analysis and precautions were necessary because the velocity estimation algorithm would not have been efficient enough to work in the allotted processing budget or robust enough to satisfy independent reviewers without the use of these additional sensor measurements.

Another effective approach utilized to improve robustness and efficiency was to apply domain knowledge to tailor the algorithm to the specific operational scenario. For example, the scene was known to be flat relative to the height of the camera, so the algorithm did not need to account for depth to features during velocity estimation. This insight also allowed the use of an efficient homography based warping for image rectification. Another example is the way the correlation search windows were set based on analysis of the dynamics and measurement errors between image captures. This made it possible to set the windows only as large as they actually needed to be which balanced tracking robustness and computational efficiency typical of small windows with the requirement to measure a velocity in the worst case horizontal wind conditions.

During development, it became clear that field testing and simulation were complimentary approaches to system testing. Simulation showed that the expected vertical velocity during image exposure had little impact on velocity estimation performance. This indicated that field

testing on a platform with a large vertical velocity was not necessary thereby reducing the cost and complexity of the field testing effort. The field testing indicated that the contrast in the simulated images was very low and therefore any analysis done with simulation would be conservative. Both testing approaches generated similar velocity estimation performance thereby cross confirming each other and increasing confidence in the system above that possible with just one test approach. Both testing approaches were useful and necessary.

Image simulation with MOC2DIMES was required during development because it provided velocity estimation performance using the appearance of the actual Mars terrain. It also provided a mechanism for investigating non-ideal imaging effects like energetic particle strikes on the CCD or dust particles on the lens. All of this simulation would have been invalid if the simulated images were not comparable to those collected in flight. Fortunately, post flight analysis showed that this was not the case indicating that realistic descent images can be simulated for algorithm development and testing in future flight programs.

The success of the DIMES has boosted confidence in using vision systems for future lander applications. DIMES can be used as is by future Mars lander missions (e.g., Mars Science Laboratory, Phoenix, Mars Sample Return). With minor modifications DIMES can also be used for velocity estimation during landing on any planetary or small body surface making it useful for proposed space science missions to return samples from comets and the moon. DIMES can also be applied to autonomous navigation of unmanned aerial vehicles for use in surveillance and other defense applications. Since a much faster onboard processor is expected to be available for future space flight missions, the performance of DIMES can be improved by using higher resolution images, using a full structure from motion solution to cope with terrain relief, and/or tightly integrating the IMU and vision system for higher rate, more accurate velocity estimation.

Two new capabilities desired the space science community are landing hazard detection and avoidance and pin-point landing. In landing hazard detection and avoidance, imagery of the landing site is collected during descent to look for large rocks, steep slopes, craters and crevasses. We are investigating multiple approaches including dense structure from motion, hazard shadow detection and crater recognition. In pin point landing, images are taken and from these the position of the lander with respect to a map is determined. We are currently investigating approaches based on landmark recognition and image correlation. Once proven to be robust and accurate, we expect our hazard detection and pin-point landing technologies to be adopted by future flight projects which will enable landing in

much more challenging and scientifically interesting terrain.

Acknowledgments

We thank Rob Manning and the entire Mars Exploration Rover project management for supporting the development of DIMES. We also thank the DIMES Review Board for the withering but useful reviews during the DIMES development and the DIMES field test team for providing the system and data for DIMES performance validation. In addition, we would like to thank other members of the DIMES team for their professional dedication and hard work. The work described in this publication was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract from the National Aeronautics and Space Administration.

References

- Amidi, O., Kanade, T., and Fujita, K. 1999. A visual odometer for autonomous helicopter flight. *Robotics and Autonomous Systems*, 28:185–193.
- Azarbayejani, A. and Pentland, A. 1995. Recursive estimation of motion, structure, and focal length. *IEEE Pattern Analysis and Machine Intelligence*, 17(6):562–575.
- Bank, T., Frazier, W., Blume, W., Kubitschek, D., Null, G., Mastrodemos, N., and Synnot, S. 2001. Deep impact: 19 gigajoules can make quite an impression. In *Proc. 24th Annual AAS Guidance and Control Conference*.
- Bosse, M., Karl, W., Castanon, D., and DiBitetto, P. 1997. A vision-augmented navigation system. In *Proc. IEEE Conf. Intelligent Transportation Systems*, pp. 1028–1033.
- Buratti, B. 1984. Voyager disk resolved photometry of the saturnian satellites. *Icarus* 59:392–405.
- Cheng, Y., Johnson, A., Matthies, L., and Wolf, A. 2001. Passive image-based hazard avoidance for spacecraft safe landing. In *Proc. 7th Int'l Symp. Artificial Intelligence, Robotics and Automation in Space (ISAIRAS 01)*.
- Cheng, Y., Johnson, A., and Matthies, L. 2005. MER-DIMES: A planetary landing application of computer vision. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2005)*.
- Corke, P. 2004. An inertial and visual sensing system for a small autonomous helicopter. *Journal of Robotic Systems*, 21:43–51.
- Gennery, D. 2001. Least-squares camera calibration including lens distortion and automatic editing of calibration points. In A. Gruen and T. Huang (eds.), *Calibration and Orientation of Cameras in Computer Vision*, Springer-Verlag, Berlin.
- Hapke, B. 1986. Bidirectional reflectance spectroscopy. IV. The extinction coefficient and the opposition effect. *Icarus*, 67:264–280.
- Harris, C. and Stevens, M. 1988. A combined corner and edge detector. In *Proc. 4th Alvey Vision Conf.*, pp. 147–151.
- Hartley, R. and Zisserman, A. 2003. *Multiple View Geometry in Computer Vision*, Cambridge University Press.
- Johnson, A. and Matthies, L. 1999. Precise image-based motion estimation for autonomous small body exploration. In *Proc. 5th Int'l Symp. Artificial Intelligence, Robotics and Automation in Space (ISAIRAS'99)*, pp. 627–634.
- Johnson, A., Montgomery, J., and Matthies, L. 2005a. Vision guided landing of an autonomous helicopter in hazardous terrain. In *Proc.*

- IEEE International Conference on Robotics and Automation, (ICRA 2005).*
- Johnson, A., Willson, R., Goguen, J., Alexander, J., and Meller, D. 2005b. Field testing of the mars exploration rovers descent image motion estimation system. In *Proc. IEEE International Conference on Robotics and Automation, (ICRA 2005).*
- Kubota, T., Sawai, S., Hashimoto, T., Kawaguchi, J., and Fujiwara, A. 1999. Autonomous landing system for MUSES-C sample return mission. In *Proc. 5th Int'l Symp. Artificial Intelligence, Robotics and Automation in Space (iSAIRAS'99).*
- Maki, J. 1999. The color of Mars: Spectrophotometric measurements at the pathfinder landing site. *Journal of Geophysical Research*, 104(E4).
- Maki, J., et al. 2003. Mars exploration rover engineering cameras. *Journal of Geophysical Research*, 108(E12).
- Malin, M., et al. 1992. The mars observer camera. *Journal of Geophysical Research*, 97(E5):7699–7718.
- Moravec, H. 1977. Towards automatic visual obstacle avoidance. In *Proc. of the 5th International Joint Conference on Artificial Intelligence.*
- Oliensis, J. 2002. Exact two image structure from motion. *IEEE Pattern Analysis and Machine Intelligence*, 24(12):1618–1633.
- Raiszadeh, B. and Queen, E. 2004. Mars exploration rover terminal descent mission modeling and simulation. In *Proc. AIAA Space Flight Mechanics Meeting*, AAS-04-271.
- Roumeliotis, S., Johnson, A., and Montgomery, J. 2002. Augmenting inertial navigation with image-based motion estimation. In *Proc. Int'l Conf. Robotics and Automation (ICRA 2002)*, pp. 4326–4333.
- San Martin, A. and Bailey, E. 2005. The MER transverse impulse rocket system: An active system to compensate for martian winds during landing. In *Proc. 28th AAS G&C Conf.*, AAS-05-086.
- Saripalli, S., Montgomery, J., and Sukhatme, G. 2003. Visually-guided landing of an unmanned aerial vehicle. *IEEE Transactions on Robotics and Automation*, 19(3):371–381.
- Shakernia, O., Vidal, R., Sharp, C.S., Ma, Y., and Sastry, S.S. 2002. Multiple view motion estimation and control for landing an unmanned aerial vehicle. In *Proc. Int'l Conf. Robotics and Automation (ICRA '02)*, pp. 2793–2798.
- Shuster, M. and Oh, S. 1981. Three-axis attitude determination from vector observations. *Journal of Guidance and Control*, 4(1):70–77.
- Smith, G.H., et al. 2001. Optical design for the Mars' 03 rover cameras. *SPIE*, 4441:118–131.
- Williams, B. 2003. Technical challenges and results for navigation of NEAR Shoemaker. *John Hopkins APL Technical Digest*, 23(1).
- Willson, R., Johnson, A., and Goguen, J. 2005a. MOC2DIMES: A camera simulator for the mars exploration Rover descent image motion estimation system. In *Proc. 8th Int'l. Symp. Artificial Intelligence, Robotics and Automation in Space.*
- Willson, R., Maimone, M., Johnson, A., and Scherr, L. 2005b. An optical model for image artifacts produced by dust particles on lenses. In *Proc. 8th Int'l. Symp. Artificial Intelligence, Robotics and Automation in Space.*