Proceedings of the 2005 IEEE
International Conference on Robotics and Automation
Barcelona, Spain, April 2005

Field Testing of the Mars Exploration Rovers Descent Image Motion Estimation System*

Andrew Johnson, Reg Willson, Jay Goguen, James Alexander, and David Meller

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109
firstname.lastname@jpl.nasa.gov

***Abstract* – The Mars Exploration Rover (MER) Descent Image Motion Estimation System (DIMES) is the first autonomous machine vision system used to safely land a robotics payload on another planet. DIMES consists of a descent camera and an algorithm for estimating horizontal velocity using image, inertial and altitude measurements. Before DIMES was accepted by MER for inclusion in the mission, its performance was validated through field testing using a manned helicopter to image three Mars analog test sites. Statistical analysis of the resulting 1900+ test cases showed that DIMES met its velocity estimation requirement. This paper describes the DIMES field test approach and associated results.**

Index Terms – Mars Landing, velocity estimation, feature tracking, MER

I. INTRODUCTION

Late in development, it was discovered that steady state winds during descent could impart a surface relative horizontal velocity to the Mars Exploration Rovers (MER) landing system. Models showed that this wind-induced velocity could cause the airbags, that cushion the landing of the rover, to rip and tear. Although the MER landing system could reduce horizontal velocity by pointing the final deceleration thrusters, the system did not have a sensor for measuring horizontal velocity. The project quickly realized that due to mass and volume constraints, they could not insert a traditional radar velocimeter into the mature landing system. However, it was determined that due to serendipitous events a down-looking camera would be relatively easy to accommodate. Development of the missing ingredient, an algorithm to process the imagery, altitude and inertial measurements to estimate horizontal velocity, was promptly started. The resulting system composed of the descent camera and horizontal velocity estimation software became known as the Descent Image Motion Estimation System (DIMES).

The DIMES camera has the same electronics as all of the MER cameras [5] and has the same optics as the rover navcams (45˚ field-of-view). As shown in Fig. 1, three images are taken during descent at roughly 2000m 1700m and 1400m above the surface. The DIMES algorithm was created from pieces of various approaches developed to solve safe and precise landing problems [1][4][6]. It 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 surface relative attitude and altitude. 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 by a Litton LN200 Inertial Measurement Unit (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.

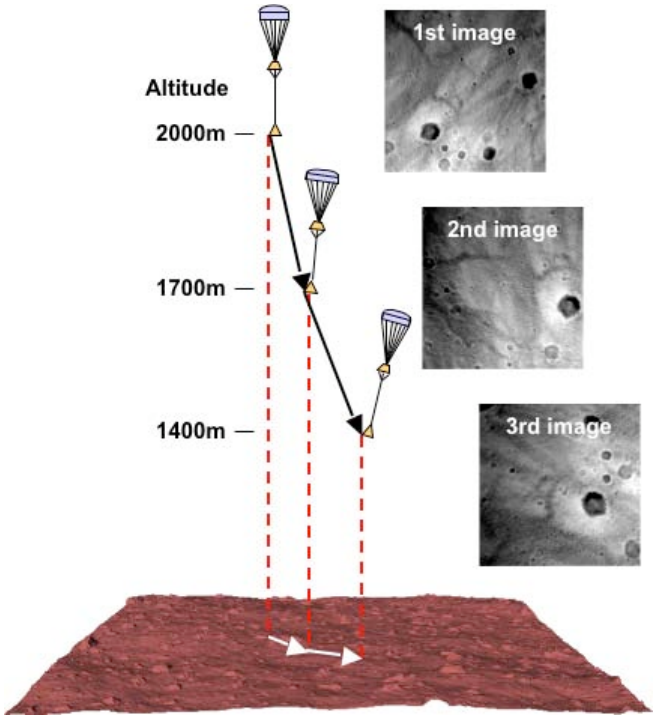

Fig. 1 DIMES descent imaging scenario.

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 correct horizontal velocity and the thruster firing took it out 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.

The development of DIMES started 19 months before launch, so validation of the performance was critical to prove that DIMES would "do no harm". 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

* The research described in this publication was carried out at the Jet Propulsion Laboratory, California Institute of Technology under contract from the National Aeronautics and Space Administration.

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 to effect thruster firing.

This paper describes the DIMES field tests. Reporting these tests is important for two reasons. Significant effort went into the design of the tests so that the dynamics and environment were as close as possible to the conditions expected during Mars landing. Also, unlike parachute drop tests or rocket sled tests, the DIMES field test approach using a manned helicopter provided large amounts of data that enabled a statistical analysis of the DIMES velocity estimation performance. The close emulation of Mars and the large amount of data make the DIMES field test approach worth investigating for other future autonomous landing validation activities.

The paper starts with a brief description of the DIMES algorithm. This is followed by a description of the field test requirements and design. Finally we show results from field tests over three different test sites in the California's Mojave Desert.

II. ALGORITHM

At the core of the DIMES algorithm is a procedure for computing horizontal velocity from a single feature visible in a pair of images. This procedure is applied to two locations in two pairs of images to obtain four independent estimates of the horizontal velocity. Below we briefly describe the single feature horizontal velocity estimation procedure and then we describe how it is applied multiple times to produce the final robust estimate of horizontal velocity. For additional details the reader should refer to [2].

A. Single template velocity estimation

First, to reduce computation, on-camera hardware (rows) and software (columns) binning are used to reduce each descent image from 1024x1024 pixels to 256x256 pixels. After binning, the region in the first image that can contain templates for tracking is computed. This region is inside overlap between the two images and outside the zero phase spot. The zero phase spot is the region on the surface containing the lander shadow and the opposition effect (brightening around the shadow). It is defined by the line between the sun and the lander and, because it moves with the lander, it can mislead correlation. However, the location of the spot in the image can be computed from known quantities (sun direction and lander attitude), which allows the algorithm to mask out the zero phase spot during template selection. To determine the image overlap, the projection of the corners of the field of view of the second image are projected into the first image using the lander attitude, altitude and the assumption that the lander has zero horizontal motion between images. The area inside the polygon defined by the projected corners and away from the zero phase spot is available for template selection. The DIMES images are taken 3.75 seconds apart and the maximum horizontal velocity expected during descent is 30 m/s so this area must be at least 110 m across to be large enough for correlation.

Next, the template selection area is searched for the point of highest contrast by computing, on a coarse grid of pixels, the Harris Interest Operator. The pixel with the highest contrast is selected as the template location. Selecting the template also automatically selects the window or correlation search area in the second image.

After the template and window have been selected they are corrected to reduce the intensity differences between the images. There are two major sources of intensity differences: frame transfer ramp and radiometric falloff. The frame transfer ramp is due to the fact that the camera does not have a shutter so it is exposing while the image CCD is being cleared and clocked out. Radiometric falloff (vignetting by the camera lens) causes an intensity fall-off from the center of the image. The frame transfer correction is computed directly from the image intensities using an algorithm that inverts the frame transfer process assuming that the image is not moving during exposure. The radiometric fall-off is removed by multiplying each template or window by corresponding pixels in an image of scale coefficients computed during radiometric calibration using an integrating sphere [5].

The next step is to rectify the template and window to remove perspective effects, scale changes and rotation between the images. Rectification uses a homography, computed from the camera projection model, and attitude and altitude measurements, to transform a descent image into an image that would be seen by a virtual camera pointed straight down. Each image is transformed to the same virtual camera. Because the horizontal displacement between images is not known, it is assumed to be zero. After rectification, any horizontal motion will show up as a shift in image intensities between the images. For efficiency, only window and template pixels are rectified.

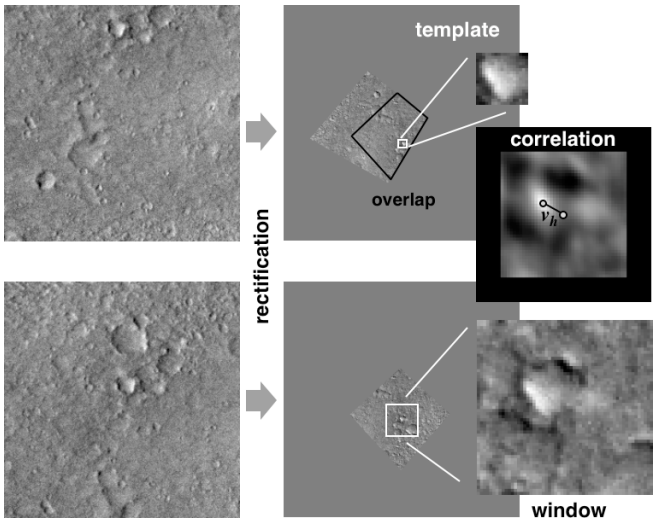

Fig. 2 Single template horizontal velocity estimation algorithm.

Pseudo-normalized correlation is used to find the best match between the template from the first image and the window from the second image. The pixel shift between the best correlation pixel and the center of the window corresponds to the horizontal motion between the images. The lander attitude, lander altitude and camera model are used to transform the pixel shift into the horizontal displacement of the lander in meters.

Once the correlation has been performed DIMES checks to make sure that correlation was successful using four metrics. Template contrast as output by the Harris Interest Operator is used to detect the images with little or no contrast or features. The height of the correlation peak is used to detect images with high noise or high frequency differences between images. The width of the correlation

peak is indicative of a poor correlation and large velocity errors. The ratio of the correlation peak to the second highest correlation peak is used to detect repetitive terrain and multiple features of similar appearance. If a template fail a test based on any of these metrics, it is considered invalid.

B. 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. First of all, 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 be also be generated. The image-based and IMU delta velocities should be close to each other. If they are not then the 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.

III. Field Test System

Monte Carlo simulation was used to estimate the performance of DIMES at each landing site under realistic EDL dynamics using imagery 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 you need field testing. In the summer and fall of 2002, the DIMES team did 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 engineering model (EM) camera and IMU mounted on a 2-axis gimbal 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 log helicopter position and time data. Fig. 3 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 detail the field test design that achieved this goal.

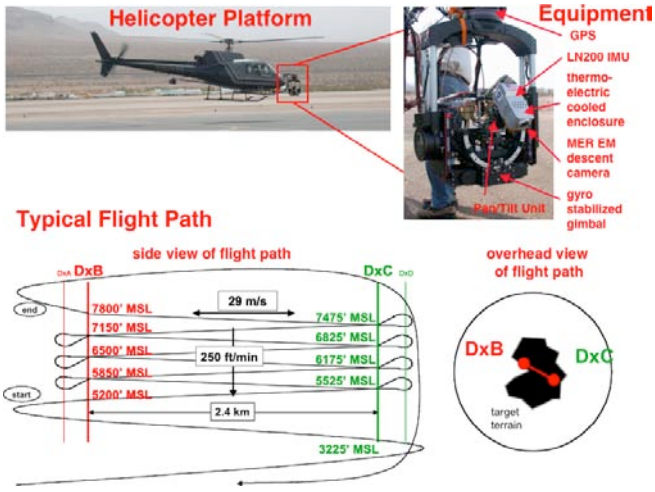

Fig. 3 Field test equipment and flight path

A. Site Selection

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 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.

Pisgah is a lava field in the Eastern Mojave Desert, consisting of numerous thin flows extending from the vent about 18km to the west and 8 km to the southeast. The lava field is predominantly "*pahoehoe*" lava (smooth, ropy surface) with some "*Aa*" lava (rough surface composed of broken lava blocks). The Pisgah lava flow shows extensive surface roughness, but on a scale of only a few meters. Pisgah terrain slopes correspond to the smoothest MER landing sites. Pisgah is the darkest of the test sites with very little contrast. Some representative images of Pisgah are shown in Fig. 7.

Kelso Dunes rise more than 200 meters above the desert floor. The dunes were created by southeast winds blowing finely-grained residual sand from the Mojave River sink,

which lies to the northwest. Rose quartz particles are responsible for the golden color. The Kelso Dunes, with peaks over 200m high, correspond to MER landing sites with greater terrain relief. Some representative images of Kelso are shown in Fig. 8.

Ivanpah Dry Lake Bed lies on the border of California and Nevada near the town of Primm. Large alluvial fans slope down to an expansive, almost perfectly flat *playa*, whose elevation is approximately 800 meters. Ivanpah Dry Lake Bed was the smoothest of the Field Test imaging sites, with slopes corresponding to landing sites with the smallest vertical relief. Some representative images of Kelso are shown in Fig. 9.

Quantitative comparisons of the field test sites to the MER landing sites were conducted for albedo, contract and slope. The *albedo* of the terrain effects the illumination within shadows. MER landing sites have albedos from 0.1-0.3, whereas the corresponding albedos for Pisgah, Kelso and Ivanpah are 0.1, 0.5 and 0.4, bracketing the MER sites. *Contrast* within an image plays a key role in DIMES performance. Contrast measured in the field test images at Pisgah was similar to contrast measured on descent images synthesized from Mars Orbital camera (MOC) images at most sites, however the contrast in the MOC images is a lower limit on the contrast that DIMES would see because the orbital MOC images view the surface through a dusty atmosphere. Topographic relief was quantified by the mean *slope* measured from Digital Terrain Models (DTM) for both MER sites and the field test sites over 30 m and 90 m baselines. On both scales, the field test sites had greater slopes (2˚ - 3˚) than the MER landing sites slopes (1˚ - 2˚).

B. Field test sensors

DIMES uses a descent camera, an IMU and a radar altimeter. These sensors should be as close to flight-like as possible. The same IMU product used by the flight system (Litton LN200 IMU) was used in the DIMES field tests. An engineering model radar altimeter (from Honeywell) was not available for the DIMES field tests, so altitude measurements from a flight like sensor were not available. In the Data Processing section we describe how altitude was computed for the tests. For imaging a MER engineering model descent camera was used. Since imaging during descent contained the greatest uncertainty, a great deal of effort went into making sure that the descent camera used in the field test would produce images that were comparable to those seen in flight.

1. **Field test camera hardware:** The DIMES camera used a 14.67mm focal length f/12 lens with a 45°x45° field-of-view. The lens includes a neutral density (ND) filter used to match the camera's sensitivity to the brightness anticipated at the different landing sites. The camera's detector was a 1024x2048 pixel frame transfer CCD (1024x1024 pixel imaging area) with 12 micron square pixels. The camera's images were digitized to 12 bits per pixel.

The field test data collection was done with an engineering model of the DIMES camera. To replicate the camera performance expected on Mars the field test camera's ND filter was changed and the camera head was cooled.

2. **Brightness Compensation:** MER A and B arrived at Mars when Mars is 1.5 AU from the sun. At this distance the sunlight on Mars was $(1/1.5)^2$ or 44% of the intensity experienced on earth. To compensate for the brighter conditions the field test camera used a darker ND filter. The ND filters used for the flight cameras were 1.44mm thick NG4 glass with a spectrally weighted optical density of 0.67 (22% transmission) . The field test filters were 1.35mm thick NG3 glass with an optical density of 0.97 (11% transmission). Fig. 4 shows the spectral transmission for the flight and field test filters.

3. **SNR Improvement and Stabilization:** Given the Martian terrain's low-contrast and the 3:1 range of brightness for the possible landing sites, the image Signal to Noise Ratio (SNR) was an important characteristic to control during the field test data collection Thermally generated dark current was the largest source of noise variation during the field testing. Dark current doubles every 6°C for the MER cameras (Fig. 4). During landing, the DIMES camera was expected to operate at -34°C. At this temperature the dark current signal represents less than 0.1% of the overall image noise signal for the most challenging (i.e. darkest) terrain. The minimum SNR for these images would be 102:1.

On early helicopter test flights the DIMES camera CCD temperatures were found to vary from 18 to 31°C over the course of a data collection run. At +23°C the dark current signal represented 15% of the overall image noise and the minimum SNR for the most challenging terrain drops to 87:1. To increase image SNR to EDL-like levels and to make the SNR consistent across all of the collected images the field test camera was cooled to 0°C for the final data collection flights. At this temperature the dark current signal represented only 2% of the overall image noise signal and the minimum SNR for the most challenging terrain rose to 100:1. Cooling was achieved using a compact, lightweight Thermo-Electric Cooler (TEC) plate on the back of the camera's detector head.

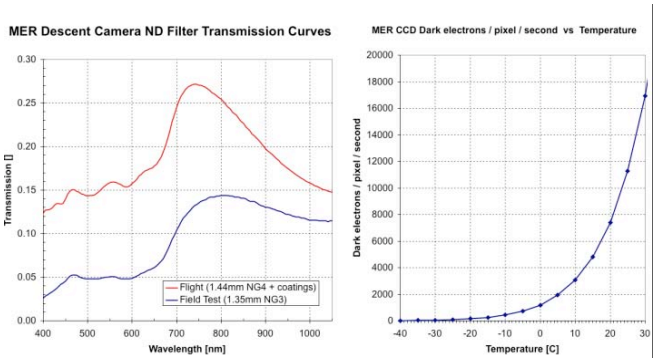

Fig. 4 Flight and field test neutral density transmission curves. (left) and CCD dark current as a function of temperature (right).

C. Field test dynamics

During landing, the DIMES images are acquired under the following conditions. The images are taken near 2000m, 1700m and 1400m above the surface. The landing system is descending at roughly 75 m/s and can be moving horizontally by as much as 30 m/s. During descent, the lander experiences a low frequency coning about the parachute at an angle of 30˚ from nadir. It also experiences a higher frequency nodding about its attach point to bridle connecting it to the backshell which can result in attitude rates up to 60˚/s.

The large descent rates close to the ground that are needed to test DIMES 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. As described below, this allowed for a significant simplification of the DIMES field test design.

A schematic of a typical flight path is shown in Fig. 3. Recall that 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 1000m. 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 2000m. During each one of these runs the pan/tilt can be activated to obtain attitude rates up to 60˚/s and off nadir angles up to 45˚. Except for vertical velocity, the field test dynamics covered the range dynamics expected during Mars landing.

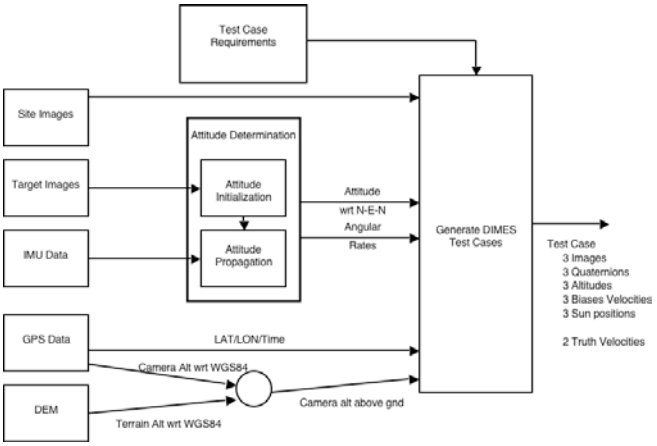

Fig. 5 Field test data flow.

D. 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 computed for each image. 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. Fig. 5 shows the flow of field test data into measurements and the details on the generation of each measurement are given below.

1. **Altitude measurement:** Helicopter altitude data was logged using an Ashtech GPS Z-Reference Station. This stand-alone unit logs data internally and was not integrated into the support electronics, save for power. A omni-directional GPS antenna was mounted above the helicopter cockpit for maximum reception. The data was post-processed at JPL using the GIPSY-OASIS II (GOA II) Software Package, an ultra high-precision GPS satellite orbit determination package providing sub-centimeter position accuracy on the ground [3]. The resulting altitude data is referenced to the World Geodetic System 1984 (WGS84) NEE.

Shuttle Radar Topography Mission (SRTM) DEMs of each imaging site, referenced to the WGS84 NEE, were obtained from the USGS. For each data point of the GPS helicopter position data, the elevation of the terrain directly below was determined by 2-D interpolation of the DEM. Taking the difference of helicopter GPS derived altitude and the elevation of the underlying terrain gave the altitude above ground. There are two sources of error in this measurement: the underlying SRTM DEM accuracy (6m) and DEM interpolation error (~1m). Combined the errors are less than the 1% of altitude errors expected from the flight sensor, but they are still significant.

2. **Attitude measurement:** To form an attitude reference, eleven white square targets (1m edges) were placed on the ground (see Fig. 6, for highlighted targets). A detailed GPS survey determined the lat, long, and height of each of the targets in WGS84 NEE. The 3 targets in the center were used to determine orientation. Only nine of the targets were directly used in the processing. 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 minutes of data collection), were acquired for each fly-over of the targets. Test site data collection occurred during the ~40 minutes 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 (v1, ...v9) 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 [7] on the ~ 9x20 = 180 pairs of vectors. 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 $\Phi(T_0 - T)^{-1} = \Phi(T - T_0)$.

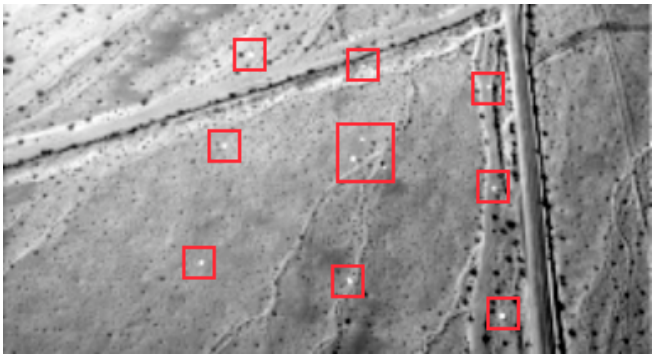

Fig. 6 GPS surveyed surface targets for attitude determination

3. **Biased IMU Horizontal velocity measurement:** Since the images making up each test case are not necessarily taken consecutively, it is not possible to generate a meaningful biased horizontal velocity from IMU propagation. The DIMES flight software still needs this measurement, so three biased horizontal velocities are constructed from the GPS positions of each image and a fixed virtual time interval associated with each image pair. The first and second image are assigned a velocity generated by dividing the change in GPS positions for the first image pair by the virtual time. The velocity for the third image is generated by dividing the change in GPS positions for the second image pair by the virtual time.

4. **Sun vector determination:** Sun position was determined using GPS time and position data to compute ephemeris. During the imaging portion of a flight (~30min. duration) sunlight incidence angle with respect to the local vertical direction of the terrain being imaged, remained essentially unchanged.

5. **Truth velocities determination**: The difference of GPS positions for each image divided by the same virtual time interval mentioned in Section D.III.3 result in 2 truth velocities for each test case.

IV. Field Test Results

Field testing occurred in October 2002. A day of flying was spend 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. 7 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 squared and tracked locations as green squares. The top row shows the result of rectifying the images using the image attitude and altitude. 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, prevented the use of the data from the first run of the day 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. 8. 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. 9. One 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.

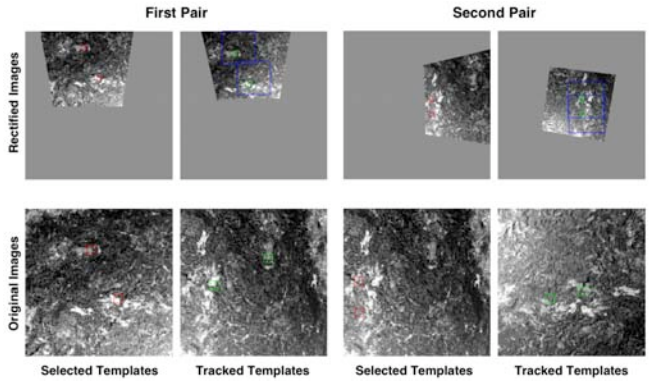

Fig. 7 Example DIMES result from Pisgah Lava Flow.

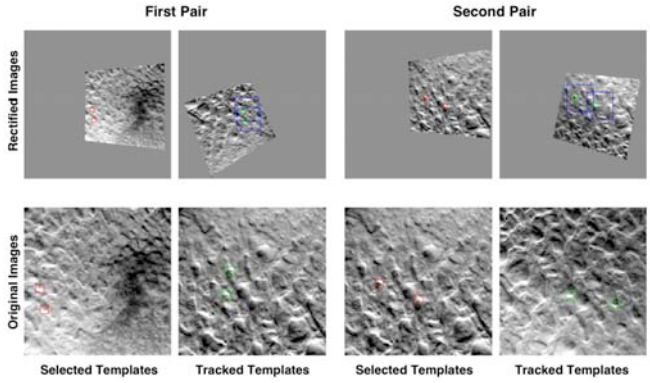

Fig. 8 Example DIMES result from Kelso Sand Dunes.

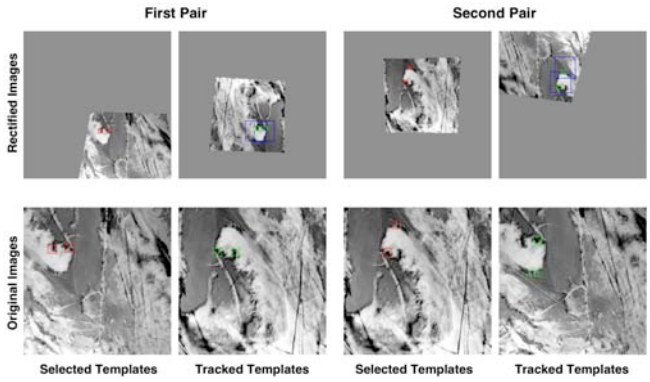

Fig. 9 Example DIMES result Ivanpah Dry Lake Bed.

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 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. 10. 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. 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. 11. Except when for the Kelso Run 3, where the shadows of clouds were moving across the terrain, most test cases were valid.

V. Conclusions and Lessons Learned

The Mars Exploration Rover Descent Image Motion Estimation System is the first passive image based system to estimate lander velocity during planetary descent. The DIMES algorithm combines sensor data from a descent imager, a radar altimeter and an inertial measurement unit in novel way to create an autonomous, low cost, robust and computationally efficient solution to the horizontal velocity estimation problem. DIMES performed successfully during both of the MER landings, and during the landing in Gusev Crater, the measurement provided by DIMES was used by the landing system to remove a possibly catastrophic horizontal velocity.

Field testing of DIMES proved very useful in validating the performance of the flight system. It showed that DIMES would estimate velocity correctly over three different types of terrain: the dark and locally rough terrain of Pisgah Lava Flow, the steep sloped and repetitive terrain of Kelso Sand Dunes and the flat and bland terrain of Ivanpah Dry Lake Bed.

Acknowledgment

We would like to thank the Ed Konefat and Konstatin Gromov for building the data collection system, Mike Armatys for providing GPS support, Mike Gradziel for designing and building the cooling system, Greg Laborde for test management and Jeff Mellstrom for over all management of DIMES.

References

[1] Y. Cheng, A. Johnson, L. Matthies and A. Wolf, "Passive Image-Based Hazard Avoidance for Spacecraft Safe Landing," P*roc. 6th Int'l Symp. Artificial Intelligence, Robotics and Automation in Space (iSAIRAS'01)*, June 2001.

[2] Y. Cheng, J. Goguen, A. Johnson, C. Leger, L. Matthies, M. San Martin, and R. Willson, "The Mars Exploration Rovers Descent Image Motion Estimation System," *IEEE Intelligent Systems* **19**(3), pp. 13-21, May/June 2004.

[3] http://gipsy.jpl.nasa.gov/orms/goa/

[4] A. Johnson and L. Matthies. "Precise Image-Based Motion Estimation for Autonomous Small Body Exploration." *Proc. 5th Int'l Symp. Artificial Intelligence, Robotics and Automation in Space (iSAIRAS'99)*, pp 627-634, June 1999.

[5] J. Maki et al. "Mars Exploration Rover Engineering Cameras." Jour Geophysical Research. **108**(E12), 2003.

[6] S. Roumeliotis, A. Johnson and J. Montgomery, "Augmenting Inertial Navigation with Image-Based Motion Estimation," *Proc. Int'l Conf. Robotics and Automation (ICRA 2002)*, pp. 4326-4333, April 2002.

[7] M. Shuster & S. Oh. "Three-Axis Attitude Determination from Vector Observations," *Jour. of Guidance & Control*, **4**(1), pp. 70-77, 1981.

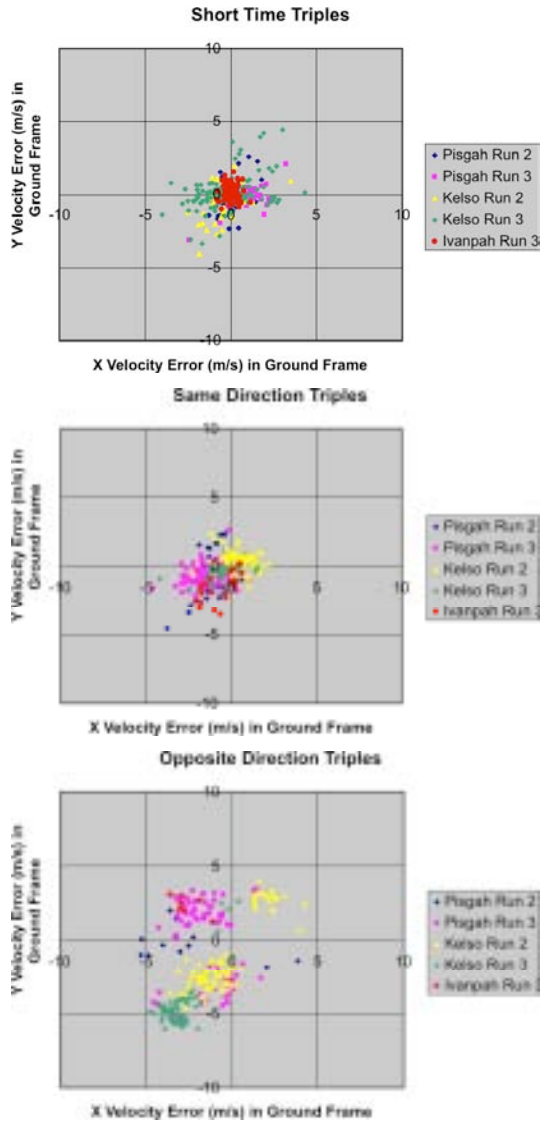

Fig. 10 Valid velocity errors for three image triple types : short time (top), same direction (middle) and opposite direction (bottom).

		# VALID	# INVALID	TOTAL	%VALID	Comment
Short Time	Pisgah Run 2	67	1	68	99%	
	Pisgah Run 3	126	8	134	94%	
	Kelso Run 2	65	0	65	100%	
	Kelso Run 3	251	0	251	100%	
	Ivanpah Run 3	139	0	139	100%	
Same Direction	Pisgah Run 2	41	0	41	100%	
	Pisgah Run 3	158	6	164	96%	
	Kelso Run 2	95	10	105	90%	
	Kelso Run 3	121	185	306	40%	Moving Clouds
	Ivanpah Run 3	63	3	66	95%	
Opposite Direction	Pisgah Run 2	26	0	26	100%	
	Pisgah Run 3	139	0	139	100%	
	Kelso Run 2	97	2	99	98%	
	Kelso Run 3	121	165	286	42%	Moving Clouds
	Ivanpah Run 3	24	0	24	100%	
Total				1913		

Fig. 11 Field test valid velocity results for three image triple types .