Editor: **Richard Doyle**
Jet Propulsion Lab
rdoyle@jpl.nasa.gov

The Mars Exploration Rovers Descent Image Motion Estimation System

Yang Cheng, Jay Goguen, Andrew Johnson, Chris Leger, Larry Matthies, Miguel San Martin, and Reg Willson, *Jet Propulsion Laboratory*

JPL's Descent Image Motion Estimation System is the first machine-vision system for estimating lander velocity during planetary descent. Composed of sensors and software, DIMES features a descent imager, a radar altimeter, an inertial-measurement unit, and an algorithm for combining sensor measurements to estimate horizontal velocity—the speed across the planet's surface the lander travels as it descends. Although the sensors are not novel technology, the algorithm and flight software that combines them are new. This algorithm combines radar, im-age, and inertial data in a novel way to create a low-cost, robust, and computationally efficient solution to the horizontal-velocity-estimation problem.

This article describes the DIMES algorithm, its testing, and its performance during both Mars Exploration Rover landings.

What led to DIMES

MER's Entry Descent and Landing system must deliver the rover and its science payload safely onto the martian surface. During entry, atmospheric drag slows down the lander far enough that a supersonic parachute can safely deploy. Just before impact, retro rockets fire to slow it down even further, at which point airbags inflate. The airbags fully encompass the lander, cushioning it as it bounces along the martian surface before rolling to rest. If the velocity at impact is too large, tests have shown that the airbags can rip and tear. The EDL system is designed to slow the lander velocity at impact sufficiently to let the airbags do their job.

In late fall 2000, the EDL design team determined that wind shear could cause the retro rockets (also known as the RAD 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, the Transverse Impulse Rocket System was added to reorient the lander system so that the retro rockets would always fire close to vertical. However, this system could not compensate for the effect of steady-state winds. This became a problem when, in the fall of 2001, the MER EDL design team learned that constant or steady-state winds were common on the martian surface.

EDL's designers soon realized that, given a steady-state horizontal-velocity measurement, the TIRS system could also serve for reorienting the lander system so that the retro rockets fired in a direction to cancel a wind-induced horizontal velocity. Because the TIRS system was well along in development, the difficulty was obtaining a horizontal velocity measurement. At this late point in EDL development, adding a traditional Doppler radar-based sensor for measuring horizontal velocity was not feasible; a radar interface to the rover avionics was not in place, and accommodating a multiple-antenna radar would be difficult.

Editor's Perspective

Spirit and Opportunity, the two Mars Exploration Rovers, achieved historic landings on the surface of Mars on 4 January and 25 January 2004. Everyone who tuned in to the real-time news coverage of these exciting events heard about the extreme engineering challenges of designing the critical Entry, Descent, and Landing sequences. Everyone who worked on the problem was relieved and amazed at how flawlessly the EDL phases unfolded.

Fortune does indeed favor the well prepared. Engineering success in deep-space exploration always comes down to anticipating contingencies and designing for them. When it came to EDL, some of the challenges were obvious: ensuring the parachute would open robustly, arranging the timing of the various events with great precision, and determining how well the airbag technology would scale up, for example. Some of the challenges were subtler—and one in particular was subtle enough that it wasn't fully appreciated until mission development was well underway.

This was the challenge of martian winds. How to detect and compensate for them? In the worst-case scenario, they could tip the vehicle over in the final stages of descent such that the powered thrust intended to eliminate downward velocity might actually drive the platform sideways and down into the surface beyond the safety envelope of the airbag cushions.

This article tells the story of how this late-understood challenge was addressed successfully—and, as it turned out, critically, for Spirit.

—Richard Doyle

Mars Exploration Rover Flight Performance

Earlier this year, Dimes operated successfully for the two Mars Exploration Rover landings. Each time, Dimes computed a valid velocity that was used by the EDL system.

MER-A—Spirit

On 4 January 2004, the MER-A or Spirit spacecraft landed in Gusev Crater, a 165 km diameter crater just south of the martian equator. During its descent, Spirit took three images of the martian surface at altitudes of 1,983, 1,706, and 1,433 meters. 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 1,570 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. Owing 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 combined steady-state and RAD-induced 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 (–11.0, 0) m/s, which caused the lander to bounce south.

Taken alone, 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
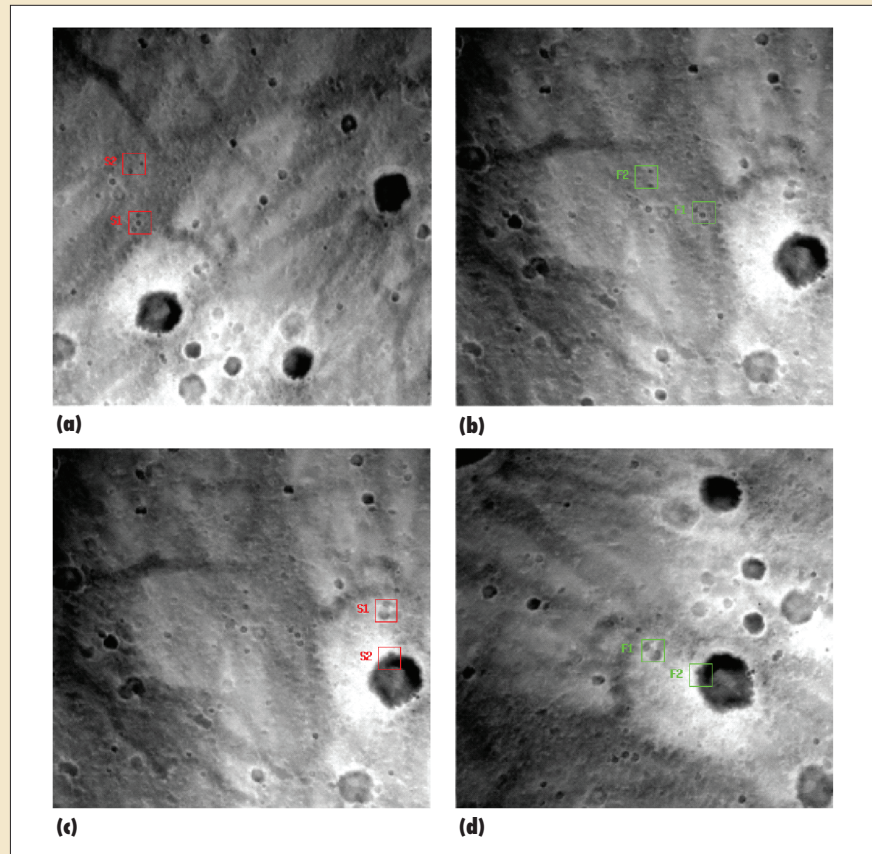

Figure A. MER-A descent image pairs of Gusev Crater with selected and tracked templates. First pair: (a) first descent image with selected templates shown as red squares and (b) location of selected templates shown as green squares on second descent image. Second pair: (c) second descent image with selected templates shown as red squares and (d) location of selected templates shown as green squares on third descent image.

performance. Furthermore, the velocity would have been to the east, toward the rockier terrain surrounding Bonneville Crater.

In late October 2001, Rob Manning, EDL development manager for MER, bumped into Miguel San Martin, MER's Attitude and Control System lead. San Martin was going on vacation for the next two weeks, but he still had MER EDL on his mind. In an inspired moment, he told Manning to give him two images and they would have a solution to the horizontal-velocity measurement problem. This chance encounter was the beginning of DIMES's development.

At this point, the EDL design team realized that, since a rover camera to measure the direction to the sun was recently removed from the MER design, the MER avionics had an unused camera interface and surplus camera electronics modules. So, if they could find a way to use a camera to measure horizontal velocity, the impact on MER avionics design would be minimal. Conceivably, a horizontal velocity measurement system based on images could be developed before launch.

On 1 November 2001, Manning emailed Andrew Johnson, a JPL machine vision researcher, and Andy Collins, a JPL advanced-imager developer, saying that "MER has a problem that I think might be solvable with a little help from you." This email laid out the concept for DIMES, a machine-vision system for estimating horizontal velocity during landing. Over the next two years, this JPL team designed, implemented, tested, and ver-

ified DIMES. Four reviews, three field tests, and 26 months later, DIMES performed successfully during Spirit's landing in Gusev Crater on 4 January 2004. As the "Mars Exploration Rover Flight Performance" sidebar explains, DIMES contributed to Opportunity's successful landing on Meridiani Planum three weeks later, on 25 January.

DIMES is the first machine-vision system to perform onboard analysis of images taken during planetary descent for landing a payload safely on the surface. Many scientists and engineers were skeptical that DIMES would work given the difficulty of making onboard algorithms robust enough to handle all possible anomalies that might occur dur-
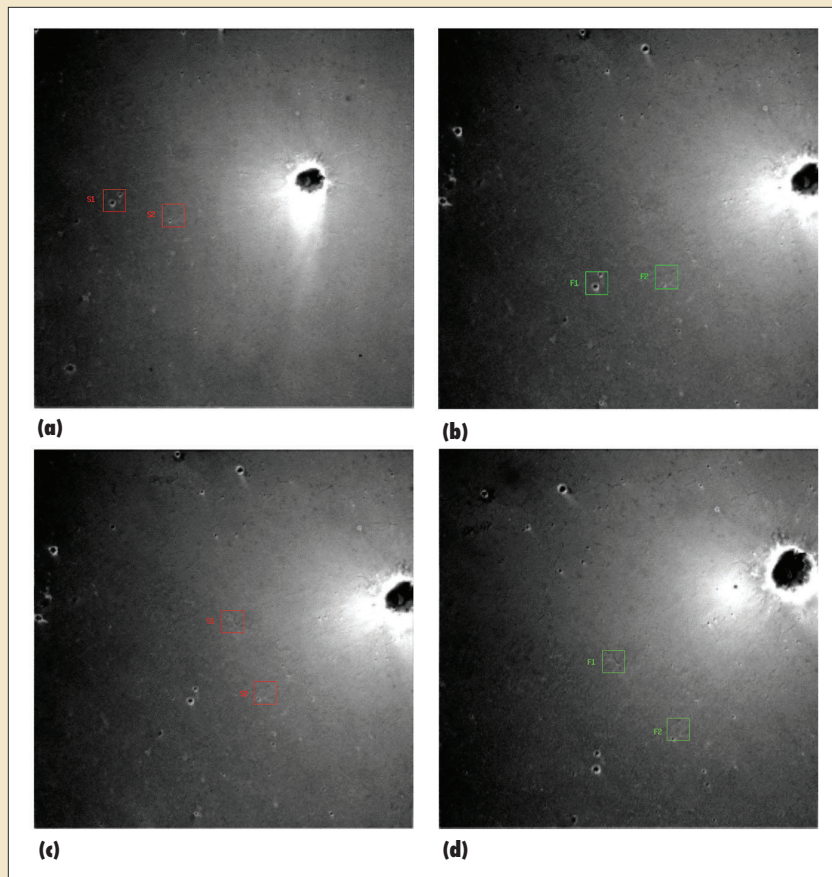
Figure B. MER-B descent image pairs of Meridiani Planum with selected and tracked templates. First pair: (a) first descent image with selected templates shown as red squares and (b) location of selected templates shown as green squares on second descent image. Second pair: (c) second descent image with selected templates shown as red squares and (d) location of selected templates shown as green squares on third descent image.

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. Postflight analysis showed that DIMES did compute the correct velocity.

Figure A shows the images taken by the DIMES camera and the associated features that were tracked.

MER-B—Opportunity

On 25 June 2004, the MER-B or Opportunity spacecraft landed on Meridiani Planum, halfway around Mars from Spirit and just north of the martian equator. During its descent, Opportunity took three images of the martian surface at altitudes of 1,986, 1,690, and 1,404 meters. 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 1,547 m altitude.

After propagation down to RAD fire, this steady-state velocity was still present (10.4, –2.8) m/s, but here 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. After RAD fire, 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 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. Postflight analysis showed that DIMES did compute the correct velocity.

Figure B shows the images taken by the DIMES camera and the associated features that were tracked.

ing planetary landing. The development, verification, and validation activities and DIMES's successful performance on both MER-A and MER-B have shown that it is possible to autonomously process and react to imagery during landing. This groundbreaking achievement has opened the door to future image-based safe and precise landing systems that will increase science return through landing-hazard avoidance in scientifically interesting terrain and pinpoint landing for sample return.

System design constraints

Given the novelty of using images to estimate horizontal velocity during descent and the extremely short development time, there was some concern that DIMES might produce an incorrect velocity that could lead to mission failure. Consequently, we designed DIMES to "do no harm"; the DIMES algorithm will produce no answer rather than one that cannot be trusted. DIMES handles this internal checking of the velocity estimate by comparing the velocities produced from two image pairs constructed from three images taken during descent.

The sequence of events for DIMES is as follows: At 2,000 meters above the surface, the EDL flight software tells DIMES to start taking images. Each image is available for processing 3.75 seconds after image exposure. Given a nominal descent speed of 75 m/s, DIMES will take images at altitudes of 2,000, 1,720, and 1,440 m. Because the firing of the RAD rockets (RAD fire) occurs around 120 m altitude, this leaves only about 17 seconds for velocity computation. DIMES images while the lander is on the parachute and can experience 60 degrees per second instantaneous angular rates and as much a 45-degree rotation about nadir and 20° off of nadir between images. Wind modeling of the MER landing sites showed that the worst-case steady state horizontal winds were less than 30 m/s.

Descent imaging of the surface of Mars must consider numerous imaging effects. First and foremost is the scene's native con-
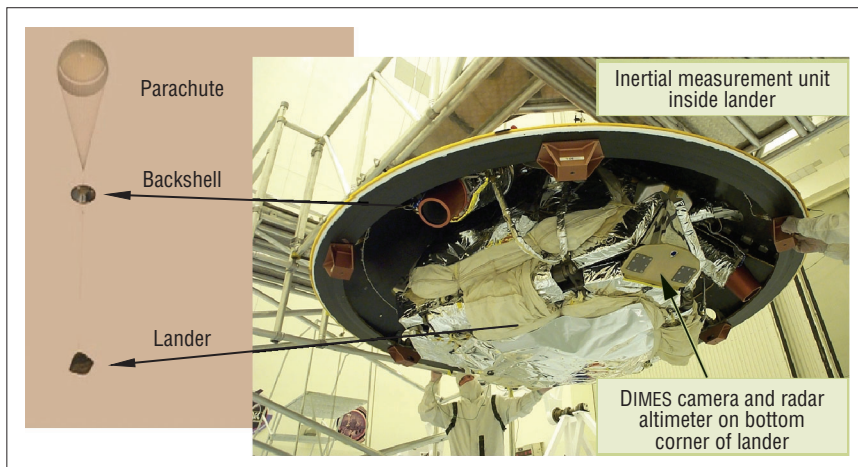
Figure 1. DIMES imager location.

trast caused by albedo variations (the fraction of light that is reflected by a body or surface) and shading. Shading depends on the terrain shape (smooth), time of day (2 p.m.), the landing site latitude (near the equator), and the season (summer). These conditions led to a situation where the sun was high in the sky, making the parachute shadow and the associated brightening around the shadow likely elements of the scene. Because the parachute shadow moves with the lander, it cannot be used for measuring horizontal velocity.

Some imaging effects are difficult to quantify. Dust in the atmosphere can reduce image contrast, while large dust particles sticking to the camera lens can cause circular dots fixed to the image. Cosmic rays can cause bright pixels in the image. The lander heat shield could still be falling when the images are taken. All of these effects could confuse an algorithm for measuring horizontal velocity and must be considered when designing the algorithm.

The DIMES camera had to be developed from existing MER camera electronics and optics. We selected the MER navigation camera optics for the DIMES camera because the 45-degree field of view had the best balance between pixel resolution and image overlap given the descent dynamics. Modifications were made to the optics and electronics to optimize the camera's performance for the low-contrast martian terrain. The DIMES camera takes images with 5 ms

exposures to minimize the effect of motion blur during imaging. As Figure 1 shows, it is mounted on the radar bracket, which is located on the corner of the lander that is nominally the lowest point on the lander during parachute descent.

The DIMES algorithm is implemented in flight software in C on a 20-MHz RAD6000 rover processor. Nominally 40 percent of the processor over 17 seconds is available for DIMES processing. The DIMES flight software communicates with other software modules to obtain measurements: images from the DIMES camera, ground-relative attitude from integration of angular rates from a Litton LN200 inertial measurement unit (IMU), and ground-relative altitude obtained from a Honeywell radar altimeter system (RAS).

Algorithm

DIMES takes as input three descent images. For each descent image, DIMES also requires elements of the lander state at the time of image exposure including the surface-relative attitude and altitude. Using the lander state, DIMES warps each image to the ground plane and then computes horizontal displacements between images using image correlation. Image correlation is applied to two locations in the first and second image and two locations in the second and third images. This approach produces four image-based horizontal-velocity estimates that are compared for consistency to each other and the acceleration computed from velocity differences from the IMU.

DIMES uses these velocity consistency checks along with image correlation metrics to decide if the computed velocity is correct. If DIMES determines that the velocity is correct, it sends a velocity correction to the EDL flight software module for inclusion in the TIRS firing solution. If it determines the velocity to be incorrect, DIMES reports that a velocity cannot be computed and TIRS proceeds without the input from DIMES.

At the core of the DIMES algorithm is a procedure for computing horizontal velocity from a single location visible in a pair of images. DIMES applies this procedure to two locations in two pairs of images to obtain four independent estimates of the horizontal velocity.
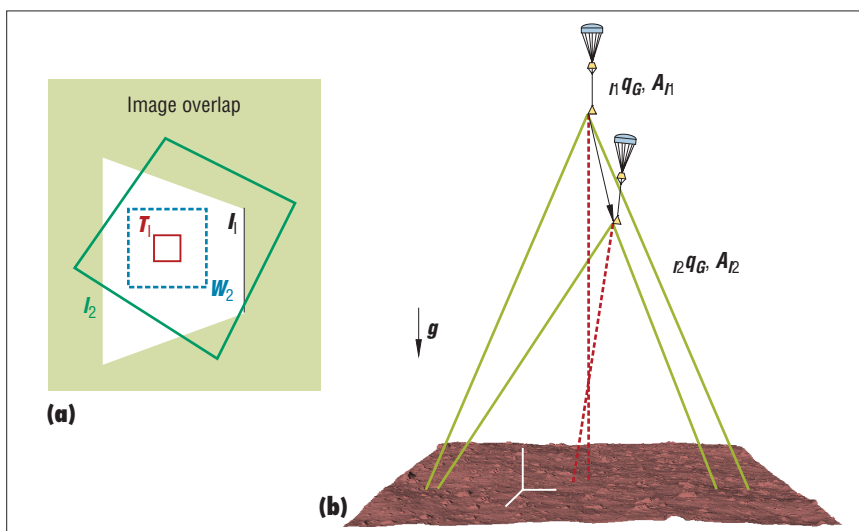

Figure 2. Computing horizontal velocity from two images. (a) DIMES uses the attitude ($_{I1}q_G$, $_{I2}q_G$) and altitude (A_1 and A_2) of the lander to determine the overlap between between two images (I_1, I_2). (b) Inside the overlap region, a template (T_1) is selected in I_1 and this defines a window (W_2) in which to search for the template.

Computing horizontal velocity with two images

Figure 2 shows the operations needed to

compute horizontal velocity from a pair of images. First, the DIMES algorithm downsamples the images from the $1,024 \times 256$ hardware binned images provided by the image acquisition flight software to 256×256 images. Next, it warps the images onto the ground plane, flattening them to remove radiometric and image frame transfer effects.[1] It selects a high-contrast location in the first image, which it them compares, through image correlation, to pixels over a window in the second image. The pixel with highest correlation corresponds to the horizontal shift between the images, which can be converted into a horizontal velocity.

Template selection. After binning, the algorithm computes the region in the first image in which to search for interesting templates to track. This region must be in the overlap between the two images and also not contain the lander parachute's shadow. To determine the image overlap, the algorithm computes the projection of the corners of the second image into the first image using the lander attitude measurements and altitude along with the assumption that the lander has zero horizontal motion between images. To account for perspective projection, a descent camera model is used. The projections of the four corners define a polygon in the first image inside which the interesting template locations can be identified.

The zero phase spot is the point on the surface that is on the line between the sun and the camera. This spot is brighter than the surrounding terrain and also contains the lander shadow. Because the shadow and brightening can mislead image correlation, the algorithm computes this spot's location in the first image from the known sun direction and the lander's attitude. Once it knows the location in the image, the algorithm eliminates this region from consideration for template selection.

Once it identifies the overlap area in the first image, the DIMES algorithm determines the location of highest contrast by computing, on a coarse grid of pixels, the Harris Interest Operator, which is the minimum variance of a 2D linear approximation of the autocorrelation function.[2]

To reduce computations, this operation occurs before image flattening and rectification. Ideally, template selection would occur after these operations, but testing has shown that template selection before image

processing still results in robust and accurate DIMES performance.

The algorithm selects the pixel with the highest contrast as the template location. Selecting the template also automatically selects the window or constrained correlation search area in the second image.

Correlation. The algorithm uses image correlation to locate the template from the first image in the window from the second image (see Figure 3). It places the template at each pixel location in the window and computes the pseudo-normalized correlation between template and window pixels.[3,4] The window pixel with the maximum correlation value (peak) is the template's most likely location in the window. The pixel shift between the best correlation pixel and the window's center corresponds to the horizontal motion between the images where the DIMES algorithm computes the size of the pixels in meters from the altitude of the first image and the angular extent of each pixel.

Correlation metrics. Once the correlation finishes, DIMES checks to make sure the identified correlation peak is the correct peak. The Harris Interest Operator calculates the total template contrast metric. This metric serves to detect the images with little or no contrast or features. Correlation is the metric corresponding to the maximum template and window correlation. This metric serves to detect images with high noise or high frequency differ-
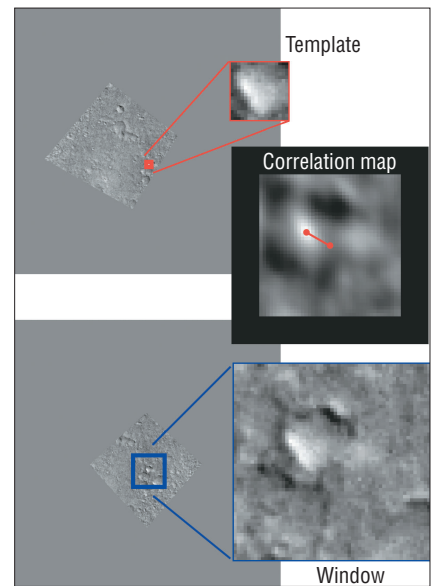

Figure 3. Image correlation.

ences between images. The DIMES algorithm measures the correlation peak width metric by fitting a biquadratic to correlation values around the peak. This metric detects low-frequency differences between the template and the searched window, which results in inaccurate motion estimation, poor correlation, and larger correlation errors. The correlation peak ratio is the ratio of the highest correlation peak to the second-highest correlation peak. This metric detects repetitive terrain and multiple features of similar appearance. Figure 4 depicts these correlation metrics.
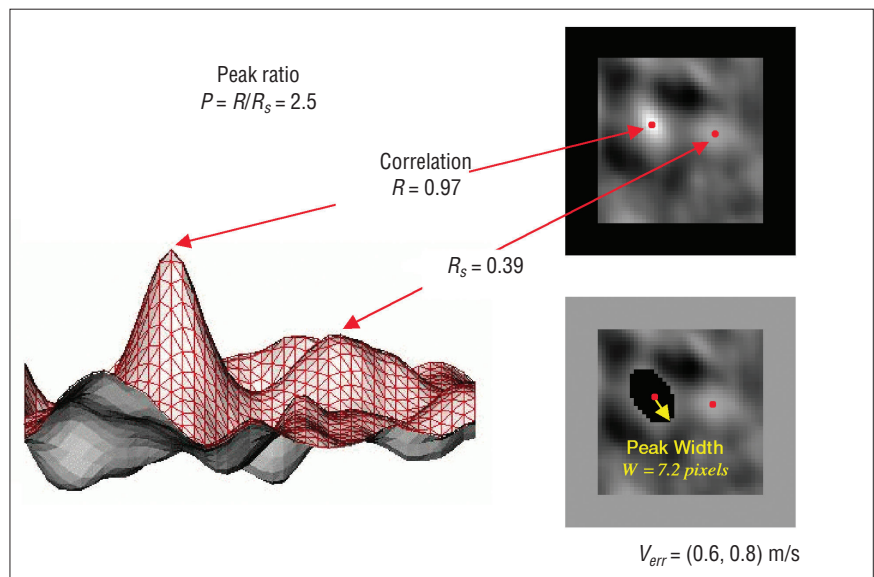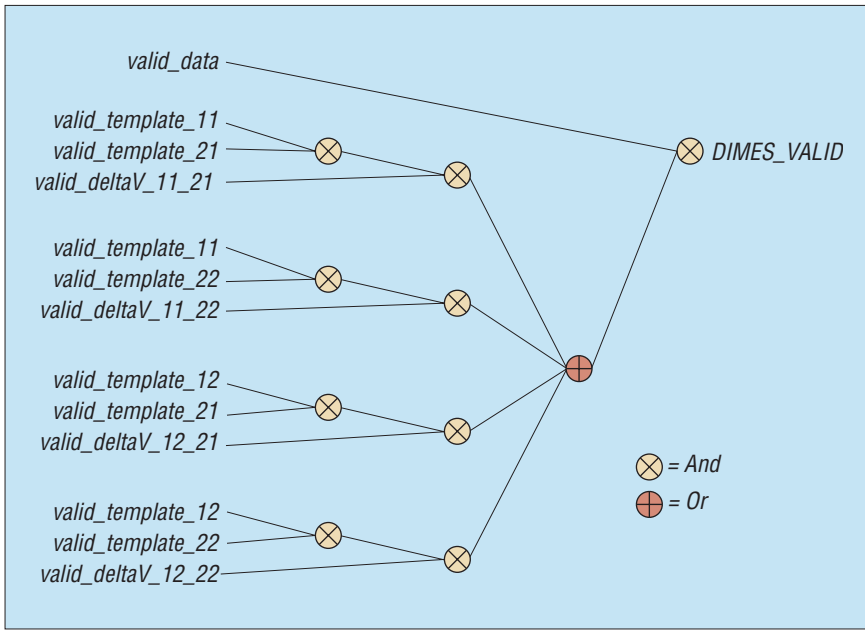

Peak ratio
$P = R/R_s = 2.5$

Correlation
$R = 0.97$

$R_s = 0.39$

Peak Width
$W = 7.2$ *pixels*

$V_{err} = (0.6, 0.8)$ m/s

Figure 4. Correlation performance metrics.

Figure 5. DIMES **validity checking logic.**

DIMES three-image algorithm validity checking

Built into DIMES are a sequence of checks on data validity, template correlation, and velocity consistency designed to ensure that DIMES does not report a false horizontal velocity. Figure 5 shows the overall validity logic. For DIMES to report a successful velocity computation, all input data (images and states) must be valid and within bounds, at least one template from each image set must have a valid correlation, and a pair of velocities (one from each image set) must be consistent with the delta velocity computed from the IMU propagation.

Software implementation

The DIMES algorithm was coded for flight in C under the VxWorks operating system. The DIMES flight software module comprises two tasks: the DIMES data task and the DIMES computation task. DDT's main responsibility is to ensure that images are acquired on time, while the DCT implements the DIMES algorithm, handles the DIMES command methods (setting and dumping parameters, and dumping images), and generates telemetry.

The DDT and DCT interface to a number of other flight software modules. Figure 6 shows the major interfaces used while DIMES is running. The numbers indicate the ordering of events, with the caveat that the figure shows only a single image acquisition. Steps 5 to 12 repeat three times before the velocity correction (step 13) and finished message (step 14) are sent.

Once the lander passes a triggering altitude (2,000 m), the EDL flight software module starts the DDT processing. The DDT's main responsibility is to ensure that images and state measurements are acquired on time and passed to the DCT.

We implemented the DIMES algorithm in the DCT as an asynchronous dataflow machine. The algorithm breaks down into operations that depend on different pieces of input data, so that the DCT will perform whatever operations are possible given the states and images available so far. This approach eliminates all dependencies on the arrival order of the data, guaranteeing that the DCT will make optimal use of available CPU time. The dataflow machine updates whenever a piece of data (an image or state measurement) arrives from the DDT.

When the DIMES algorithm has completed, a signal goes to the DDT and a status variable and valid flag are set so that
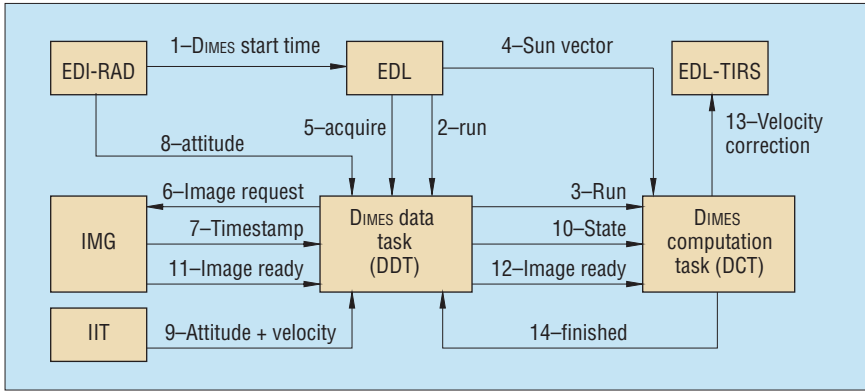
DIMES three-image algorithm

The DIMES algorithm uses two sets of images to determine the pre-RAD horizontal velocity. It requires three images to perform, with images 1 and 2 used as the first pair and 2 and 3 used for the second pair. For each pair, the two-image algorithm operates twice to determine the horizontal velocity for two templates, resulting in a total of four velocities.

An interesting advantage of the three-image algorithm is its use of the IMU to perform a check on feature tracking that is completely independent of the image data. DIMES computes the differences between each velocity in the first and second image sets, resulting in four delta velocities. It compares each of these delta velocities to the delta velocity computed from the velocity that the IMU software module has been propagating since before entry. If at least one image delta velocity is consistent with the IMU delta velocity, DIMES is deemed to have measured a correct surface relative velocity. For DIMES to succeed, at least one velocity from each image set must be correct. Consequently, one template from each image set can result in an incorrect correlation, without compromising DIMES. Allowing for the possibility of template correlation failures gives DIMES a robustness to nonideal imaging conditions.

If DIMES successfully computes a velocity, it passes an IMU velocity correction to TIRS. This velocity correction is the average of all valid velocities from the second image set.

Figure 6. DIMES **tasks and interfaces.**

TIRS, which polls DIMES every Real Time Interrupt (RTI), knows that a velocity correction is available. Figure 7 shows the DIMES algorithm's dataflow graph. Circles represent input and output data, rectangles indicate operations, and arrows show the flow of information between inputs, operators, and outputs. Each arrow has a *ready flag* associated with it, indicating whether the data is available (in the case of input data) or the processing step has completed (in the case of operators). When all of the inputs to an operation have their ready flags set, the operation can be performed.

Preflight validation

Validation of the performance of DIMES was critical to prove that DIMES would do no harm during EDL. Because we could not completely test the entire DIMES flight system in a realistic flight-like environment, we broke the validation tests 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.
- 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 RAD fire.

Monte Carlo performance

We used Monte Carlo simulation to estimate the performance of DIMES at each landing site under realistic EDL dynamics using imagery of actual martian landscapes.

To estimate DIMES performance during landing, we performed thousands of Monte Carlo runs. During each run, we simulated sensor measurements and input them into the DIMES flight software to produce a velocity estimate. We compared this estimate to the true horizontal velocity to generate a velocity measurement error. Using multiple Monte Carlo runs, we generated the velocity error statistics.

Multibody aerodynamic simulations of the EDL lander system served to generate multiple lander trajectories during EDL. These trajectories give the camera's position, attitude, angular rates, and velocities during image exposure. Using this state information, we then used a high-fidelity model of the DIMES camera to generate descent images
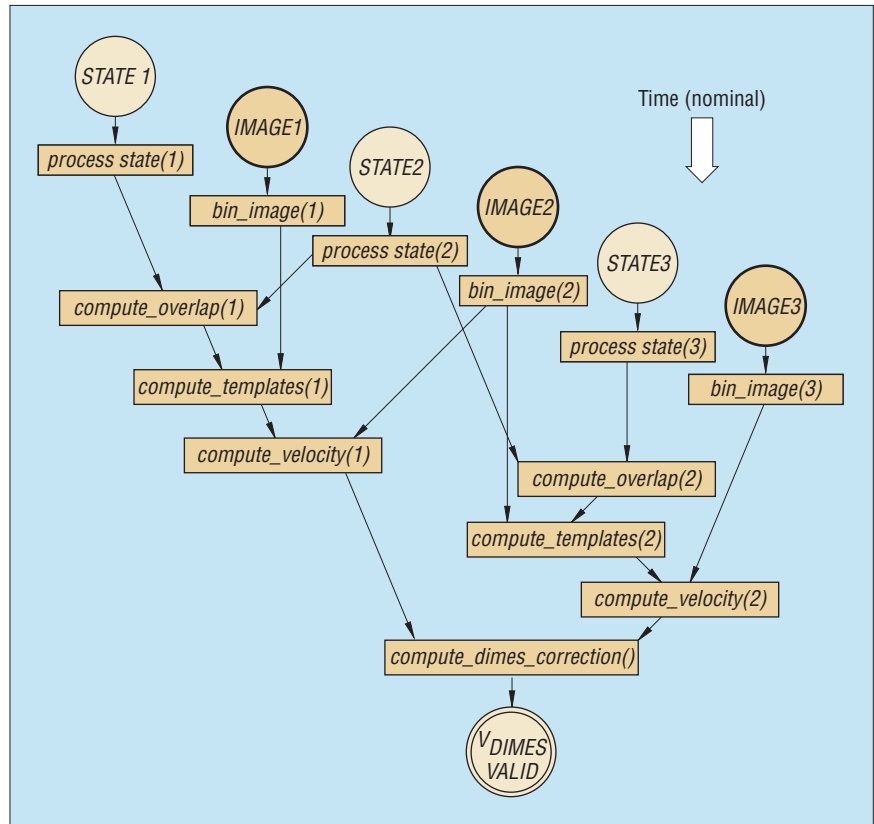

Figure 7. Dataflow machine.

by projecting the camera's field of view onto an orbital image of the landing site.

We used Mars Observer Camera Narrow Angle images for three reasons:

- The pixel resolution was comparable to DIMES image resolution.
- They were wide enough to contain a DIMES image.
- Many images were covering each MER landing site.

Figure 8 shows an example of three imaging states and associated fields of view on a single MOC-NA image. Finally, we used sensor models to generate noisy measurements of lander attitude, altitude, and biased horizontal velocity from the true EDL trajectories.

Using Monte Carlo simulation, the DIMES performance at Gusev Crater (Spirit's landing site) was 99 percent valid velocity answers and a measurement accuracy of 3.7 m/s (3 sigma). At Meridiani Planum (Opportunity's landing site), the DIMES performance was 71 percent valid velocities with 3.7 m/s (3 sigma) velocity error. Both errors are for an average altitude of 1,550 m.

Field testing

Monte Carlo simulation is important for assessing algorithm performance, but it cannot replace taking pictures with a real camera at altitude over Mars-like terrain. This is the purpose of 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.

In the DIMES field tests, we collected the required data to run the algorithm and assess its performance and then, in postprocessing, ran it through the DIMES flight software to estimate velocity. In other words, we performed no real-time velocity estimation during the tests. An engineering model DIMES camera and IMU were rigidly attached and placed on a pan-tilt unit. We placed this assembly in an inertially stabilized platform attached to the front of a crewed helicopter. A GPS unit served to record position for verifying DIMES velocity measurements. We used the difference of GPS position and a
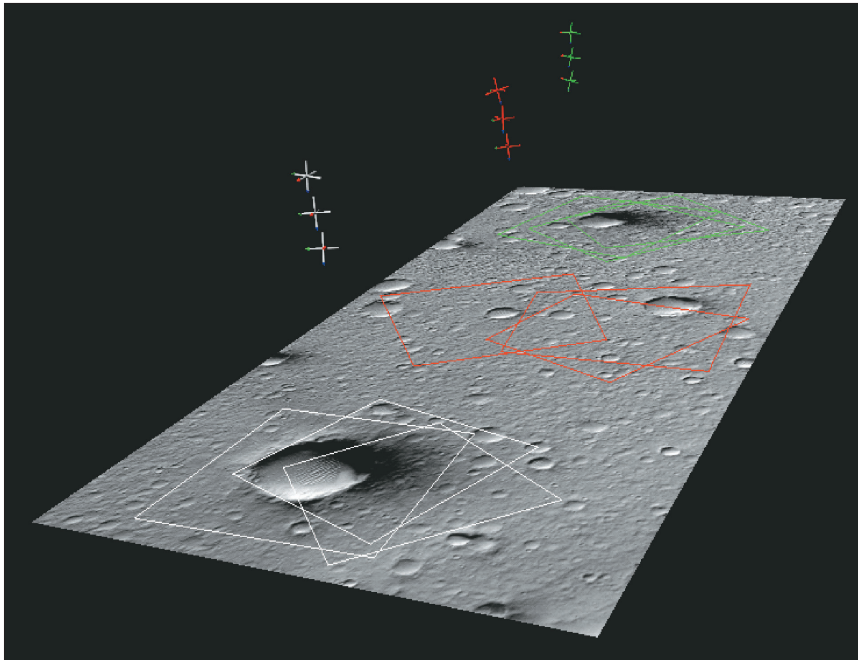
Figure 8. Monte Carlo simulation.

GPS registered digital elevation map to simulate altitude measurements. Finally, we obtained an initial estimate of camera attitude, which is needed to initialize the gyro attitude propagation, by taking images of targets on the ground in known GPS position. By aligning the 3D ground targets with the image of the ground targets, we could

Helicopter platform

Pisgah Lava Flow triples

GPS

LN200 IMU

MER EM descent camera

Gyro-stabilized gimbal

Typical DIMES image (Pisgah Lava Flow)

Digital elevation maps

Pan/tilt unit

Figure 9. Field test equipment, trajectory, and typical image.

determine the camera's attitude in GPS coordinates using standard photogrammetry techniques. Figure 9 shows the field test equipment and helicopter.

During a typical test, the helicopter flew up to 2 kilometers altitude and then slowly flew down in a zigzag trajectory until close to the ground. For some runs, we activated the pan-tilt unit to induce angular rates up to 60°/s. Consecutive images were at similar altitudes, which is not the case for EDL. However, the zigzag motion generated image triples, albeit significantly separated in time, with altitudes and overlap typical of EDL.

We performed field testing at three test sites in the Mojave Desert: the Pisgah Lava Flow, Kelso Sand Dunes, and Ivanpah Dry Lake Bed. We selected these sites because they are free of vegetation, cover a wide range of appearances, and have significant elevation variations. The final field test in early October 2002 produced the most complete set of data for assessing the DIMES algorithm's performance. During this test, we conducted three helicopter flights at different times of day at each test site. When we ran this data through the DIMES flight software, all image triples that emulated flight-like conditions resulted in velocity estimates with errors less than the 5 m/s requirement. These runs proved that DIMES worked with flight camera images independent of the terrain imaged or the time of day.

Flight system testing

We conducted flight system testing to ensure that the DIMES flight software and camera operated as expected on the flight system. Specifically, flight system testing verified the following:

- Sensor alignments and calibration data
- The temporal alignment of data from the IMU and camera
- The accuracy of velocities computed by the flight software
- The performance of the flight software in fault scenarios
- The amount of processing required to compute the DIMES velocity estimate

Verification of the camera-to-lander alignment—that is, the phasing—for both spacecraft occurred during assembly test and launch operations at Kennedy Space Center in February through April 2003. All other testing occurred on a functionally equivalent avionics and flight software test
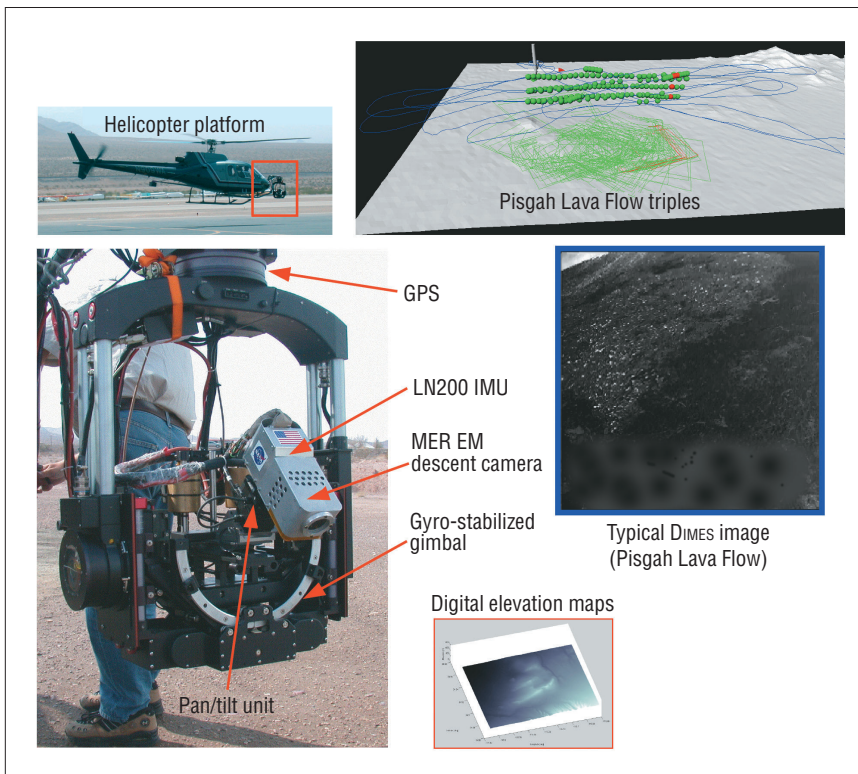
bed called the Cruise and EDL Test Bed in June through December 2003.

Developed by the Mars Exploration Rover Project, DIMES performed successfully on the MER Entry Descent and Landing system during both rover landings this past winter. Future Mars lander missions such as the Mars Science Laboratory, Phoenix, and Mars Sample Return can use DIMES as is. With minor modifications, DIMES can also serve for velocity estimation during landing on any planetary or small-body surface, making it useful for proposed space science missions such as the planned Comet Surface Sample Return and Lunar Surface Sample Return. DIMES will also work for autonomous navigation of unmanned aerial vehicles used in surveillance and other defense applications. ▱

Acknowledgments

References

1. Y. Cheng et al., "Passive Image-Based Hazard Avoidance for Spacecraft Safe Landing," *Proc. 6th Int'l Symp. Artificial Intelligence, Robotics, and Automation in Space* (iSAIRAS 01), European Space Agency Publications Div., 2001.

2. C. Harris and M. Stevens, "A Combined Corner and Edge Detector," *Proc. 4th Alvey Vision Conf.*, 1988, pp. 147–151.

3. A.E. 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), European Space Agency Publications Div., 1999, pp. 627–634.

4. S.I. Roumeliotis, A.E. Johnson, and J.F. Montgomery, "Augmenting Inertial Navigation with Image-Based Motion Estimation," *Proc. IEEE Int'l Conf. Robotics and Automation* (ICRA 2002), IEEE CS Press, 2002, pp. 4326-4333.
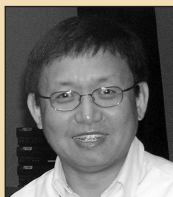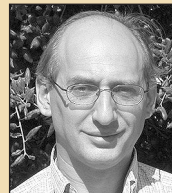
Yang Cheng, DIMES Algorithm Developer at the Jet Propulsion Laboratory. Contact him at ycheng@robotics.jpl.nasa.gov.

Jay Goguen, DIMES Science Lead at the Jet Propulsion Laboratory. Contact him at jay.goguen@jpl.nasa.gov.

Andrew Johnson, DIMES Software Lead at the Jet Propulsion Laboratory. Contact him at aej@jpl.nasa.gov.

Chris Leger, DIMES Flight Software Developer at the Jet Propulsion Laboratory. Contact him at cleger@robotics.jpl.nasa.gov.

Larry Mattthies, Supervisor Machine Vision Group at the Jet Propulsion Laboratory. Contact him at lhm@robotics.jpl.nasa.gov.

Miguel San Martin, MER Guidance and Control Project Element Manager at the Jet Propulsion Laboratory. Contact him at miguel.sanmartin@jpl.nasa.gov.
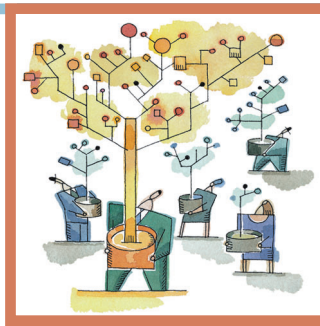
Reg Willson, DIMES Hardware Lead at the Jet Propulsion Laboratory. Contact him at rwillson@jpl.nasa.gov