AGU PUBLICATIONS

Earth and Space Science

RESEARCH ARTICLE
10.1002/2016EA000189

Key Points:
- Stereophotoclinometry and horizon matching can determine position
- This technology provides independent navigation to rovers and landers
- This technology can reduce risk and cost to missions to the surface of other planets and moons

Supporting Information:
- Supporting Information S1

Correspondence to:
E. E. Palmer,
epalmer@psi.edu

Citation:
Palmer, E. E., J. N. Head, R. W. Gaskell, M. V. Sykes, and B. McComas (2016), Mercator—Independent rover localization using stereophotoclinometry and panoramic images, *Earth and Space Science*, *3*, 488–509, doi:10.1002/2016EA000189.

Received 6 JUN 2016
Accepted 11 OCT 2016
Accepted article online 3 NOV 2016
Published online 2 DEC 2016

©2016. The Authors.
This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

Mercator—Independent rover localization using stereophotoclinometry and panoramic images

Eric E. Palmer[1], James N. Head[2], Robert W. Gaskell[1], Mark V. Sykes[1], and Brian McComas[2]

[1]Planetary Science Institute, Tucson, Arizona, USA, [2]Raytheon Missile Systems, Tucson, Arizona, USA

Abstract We present a navigation technique for a lunar vehicle based upon matching horizon imagery with a database of synthetic panoramas generated from ultrahigh-resolution topographic models created by stereophotoclinometry. We conducted two tests, one with lunar topography and fully synthetic panoramas, and another using Earth-based data and actual panoramic images. We have shown that the horizon-matching technique can search a wide area for a vehicle's location and determine that location to within 6 m using readily available imagery.

1. Introduction

Navigating on the Earth has become routine due to the support of Global Navigation Satellite Systems such as the Global Positioning System (GPS); however, navigating on the surfaces of other planets or their satellites remains difficult. Older technologies, such as triangulation, inertia navigation, dead reckoning, or odometry, must be relied on even though they are time intensive. External (currently human-based) positional information is needed for landers to identify the context of their location. Rovers frequently need this external data to update their position because inertial measurement units (IMU) drift over time. Additionally, astronauts operating on the surface are highly dependent on navigational aids to maintain positional awareness, especially when they are out of the line of sight of the lander during surface operations.

Autonomous non-GPS lunar surface navigation can enhance and expand science and exploration science activities. Every sample and image acquired, whether for lunar petrology, geomorphological analysis, or resource prospecting, requires excellent documentation of its location. As a potentially more precise and much quicker adjunct to radio navigation, our optically based method will reduce required time and effort. In addition, program risk (the cost of operations) can be reduced if the method obviates the need for a lunar GPS constellation or an extensive surface beacon network.

Rovers have operated on Mars continuously for over a decade, traversing over 50 km and collecting vast amounts of data. One continual problem is the difficulty of navigating remotely, which is done with a hybrid of automated routines and human involvement [*Arvidson et al*., 2004; *Li et al*., 2002, 2006, 2011]. The process is time consuming, reducing the rover's ability to travel long distances and placing operational constraints on science operations.

We are suggesting a technique that allows for highly precise location determination without the need of extensive human involvement. Additionally, it can be used for both the initial position determination as well as continued movement tracking. One of the many benefits is that it is quick and can be done with standard hardware, specifically a camera that can be used to stitch panoramic images together and onboard computation. This would allow a rover to periodically stop to update its position, evaluate to ensure progress is consistent with expectations, and then continue on an extended traverse. It would allow for longer traverses without the delay of mission control's involvement and consequently shorter transitions between science stations.

A study of a semiautonomous system was done previously called Visual Position for Rovers [*Cozman and Krotkov*, 1996; *Cozman et al*., 2000]. Their study performed similar horizon matching around Pittsburgh, PA, the high deserts of Chile, the U.S. Rocky Mountains, and a simulated Apollo 17 landing site. They created a complete system consisting of a rover, cameras, and processors, and used existing digital elevation models (DEM) that were openly available, typically USGS topographic maps with a resolution of 30 m. Their system required human involvement to aid in filtering images during the generation of an acceptable panorama. Their system was effective having errors on the order of 100–400 m or about 2.5 to 6.5 times larger than

Table 1. Imaging Systems Employed on Mars/Moon With Resolution Data

Optical System	Mission	Body	Resolution (m/pixel)	Pixels	System Type
MOC	MGS	Mars	1.4	2048 pixel	Pushbroom
HiRISE	MRO	Mars	0.2	20, 048 wide	Pushbroom
LROC	LRO	Moon	0.5	5064 wide	Pushbroom
HIRES	Clementine	Moon	7–20	288 × 384	Framing

the DEM resolution. Our work complemented theirs and showed that with higher-resolution images positions can be determined close to the resolution of the terrain.

2. The Mercator Project Design

Mercator is a project to advance autonomous surface navigation capability for lunar applications. The Mercator approach is to use orbital and ground imagery to obtain an estimate of the vehicle's location. Stereophotoclinometry (SPC) would use orbital imagery to construct a DEM of the relevant terrain. The DEM would be used to make a set of synthetic panoramas depicting the local terrain from a height comparable to the vehicle imaging system. During operation the vehicle imaging system would obtain a real panorama. A comparison algorithm would determine which synthetic panorama most closely matched the ground panorama. The location of the closest match would be designated the location of the ground vehicle.

The significance of this work is that it will provide a mechanism for precision navigation on the lunar or Mars surface in the absence of external navigational aids or communication with the Earth. Early demonstration of this capability will influence lunar or Mars science architectural decisions, potentially saving time and $\$10^9$ in life cycle costs of local GPS or a surface beacon network. This cost and time savings result from developing a method that reuses data acquired for other purposes and exploiting capabilities likely to be standard on lunar exploration vehicles.

For fundamental lunar science, the method will enhance understanding of lunar geodetics by producing tie points wherever camera and transponder-equipped assets travel on the Moon. Enhanced geodetics allows improved registration of lunar topography and gravity models, reducing errors in geophysical models derived from such data, and hence, constraining models of surface formation and evolution. For exploration applied science, the method will improve our preparedness for conducting geodetic measurements and optimize science return from robotic and human missions through automated sample/image site documentation. Mercator development will thus help define surface exploration field studies.

Specific science applications of the results of this study include the following: (1) Site documentation for scientific sample collection, (2) navigation for autonomous vehicles undertaking science investigations, (3) location determination of science vehicles after landing guidance failure, (4) safe crew return on circuitous geologic field study traverses, and (5) precise navigation to cached supplies supporting field geologists.

2.1. Navigation Requirements

In order to implement the Mercator algorithm, the following inputs are required: imagery of the horizon (360°), tilt of the imaging system, the height of the imaging system above the local level, and a standard reference direction, e.g., true north. Navigation system analysis indicates that vehicle attitude errors will flow into the position error. The impact is strongly dependent upon the surrounding terrain and the look direction, precluding a single-attitude error requirement. The project largely focused its resources on testing the algorithm and in section 5 we discuss the errors and identify types of terrain that prove problematic for the Mercator approach.

The Mercator system makes extensive use of a DEM of the region upon which a rover or astronaut is traveling. The resolution of the DEM has significant impact as to how accurate the localization can be. Table 1 describes recent data sets that have been obtained for both the Moon and Mars. A DEM can be constructed using SPC at a resolution comparable to the resolution of the imagery, which is used by the Mercator algorithm and with imagery of the horizon to localize a position.

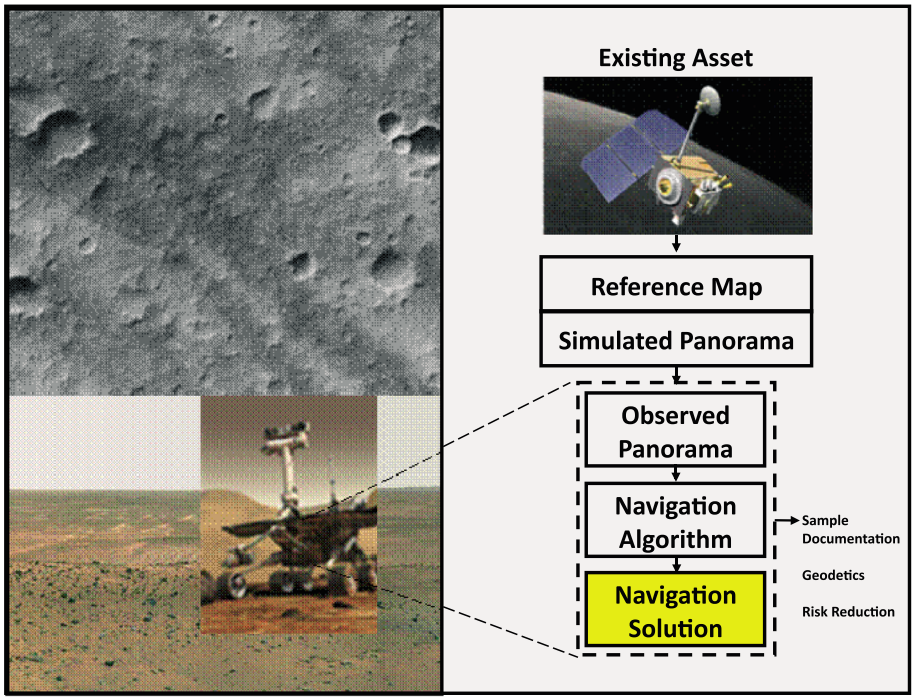

Figure 1. Top level concept of operations for Mercator.

2.2. Concept of Operations

The Concept of Operations (CONOPS) developed for Mercator is shown in Figure 1. It is assumed that any landed vehicle on the Moon or other terrestrial body will set down in an area with extensive preexisting orbital imagery in accordance with NASA practice since the 1960s (Table 1). Using the relatively new technique of SPC the orbital imager is converted to a Digital Elevation Model (DEM). Using Mercator algorithm, panorma.f, the DEM is used making a series of synthetic panoramas. These are compared to a real panorama generated by the onboard landed imaging system acquired when the vehicle is stationary. The comparison will yield a solution map, from which the location of the landed system can be derived. This position can be augmented with additional data such as odometry and radio tracking.

2.3. Notional Vehicle Design

The notional ground vehicle would carry an instrument suite to gather the information required to implement Mercator. The instrument suite would include an imaging system capable of acquiring 360° horizon image data, an Inertial Measurement Unit (IMU) capable of determining vehicle acceleration, translation and tilt relative to the local gravity vector, and a star tracker capable of determining the direction of true north and tilt. For the purposes of this study we determined to use an imaging resolution of 0.1° per pixel. We further determined that vehicle tilt knowledge would be required to about 2°. The direction of true north would need to be determined to within 1°.

Our minimum rover design included an IMU capable of determining tilt to 2°, well within the capabilities of the widely used IMU systems, such as the Northrop Grumman LN-200. Optical sensors begin with a set of cameras to provide images that would be stitched together to form a panorama. Cameras such as these could either be a 360° camera, a set of cameras covering 360°, or a single camera that rotates over 360° (Table 2). It is true that radio tracking can be used to constrain the location of a lander on another planet given sufficient time, but it is instructive to consider that the Viking 2 landing site was not localized until the arrival of Mars Global Surveyor two decades later.

Additionally, and more importantly, our design would have a star tracker. A star tracker can provide the rover the ability to determine its attitude in three dimensions. For horizon matching, it is critical to have a precise attitude because we use these data to align the panoramas, i.e., ensure north is in the same place for both

Table 2. Existing Optical Sensors Readily Exceed the Performance Standards Used in This Study

Camera System	Sampling (mrad/pxl)	Field of View (deg)	Images to Scan Horizon	Star Tracker	Angle Accuracy (arcsec)
This study	1.75	-	-	This Study	3600
IMP	0.98	14	29	Andrews Pyxis	<10
MER PanCam	0.28	17	24	Ball CT-633	5–25
MER NavCam	0.82	45	9	Sodern SED 26	<8
MER HazCam	2.1	124	4	Terma HE-5AS	<1

synthetic and real panoramas. We also need attitude to remove any tilt in the real panoramas. We evaluated several star trackers and noted that in each case the performance significantly exceeded our design requirements (Table 2).

The information flow for this study and for a landed flight vehicle is shown in Figure 2. It is assumed that any landed craft would have an initial position known to within 5 km, well within both radio tracking and landing precision on the Moon. Once the initial position is found the panoramic matching updates can be made during long-range traverses to correct for IMU drift. If required, Mercator could function in a "lost-on-the-moon" mode where there are no assumptions about starting location. In that extreme case

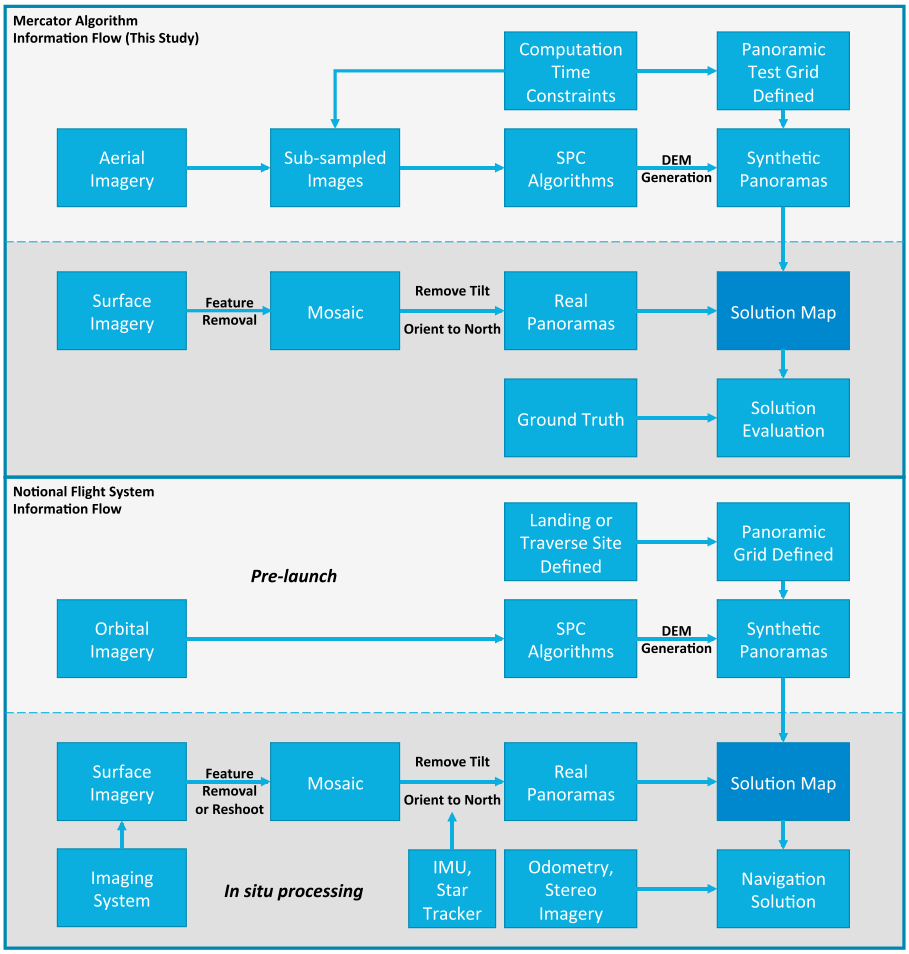

Figure 2. Mercator information flow as implemented in this study and as it could be implemented in a flight system.

the star tracker and IMU would provide a latitude estimate, bounding the search area for the Mercator algorithm. Further implementation of this mode delves into system design issues beyond the project scope.

The Mercator algorithm is fundamentally based upon the "dropoff" scenario where the rover can be placed anywhere within a 10 km by 10 km region. This type of problem is the one most likely to be used during surface landing missions, and then it can be used throughout a rover's traverse. By putting realistic limits on the possible locations upon which the rover could start, it reduced the area of which we had to compute ultrahigh-resolution topography. While the methodology is insensitive to the physical boundaries, it can become computational limiting to do an excessively large area, which provides no real benefit; i.e., we could calculate the entire lunar surface, but it is unnecessary.

2.4. Generation of Topography—How SPC Works

The first step of Mercator is the development of a topographic model. For our testing, we are using the topography generated using the SPC technique [*Gaskell et al.*, 2008]. SPC is a suite of software routines that uses both stereo imaging and photoclinometry to derive topography. The routines are able to generate topography with a ground sample distance (resolution) near the pixel resolution of the source images. If the images have significant overlap, then they can produce a stereo-derived topographic model. If the images have significantly different illumination conditions, then it can produce a photoclinometry-based model. However, as most planetary missions are, there are many more images that do not qualify as stereo pairs, so a combination of the two methods provide topographic solutions based on a much wider base of images.

At its root, SPC is geometric stereo, meaning that it starts with spacecraft geometry to calculate the height of individual control points (aka landmarks) on the surface using two or more images taken from different directions. The advantage of SPC is that it has integrated both the albedo and shape of the surface in an iterative process to increase precision and resolution. If there is no change in illumination conditions, SPC is reduced to multiimage photogrammetry. If there is an insignificant stereo angle, SPC works as 2-D photoclinomerty with multiple images.

2.4.1. Requirements

Traditional stereo requires the following:

1. Images that are taken with similar lighting conditions;
2. Stereo angle between 10° and 40°;
3. Images have nearly identical image resolution;

Stereophotoclinomery works best with the following:

1. A minimum of three images (typically $>$30, this study averaged four);
2. Two stereo images;
   a. Emission angles 45° (desired 35°–48°, limit 5°–60°);
   b. Stereo angle 90° (desired 70°–110°, limit 10°–120°);
   c. Incidence angle 0° (desired 0°–20°, limit 0°–60°);
3. Three photoclinometry images;
   a. Emission 0° (desired 0°–20°, limit 0°–60°);
   b. Incidence angles 30° (desired 20°–50°, limit 0°–60°);
   c. Variation in illumination geometry of 90° (desired 40°–90°, limit 30°–180°). This refers to having different positions of the Sun with respect to the observed target.

In general, geometric stereo is more accurate for long-scale topography, while photoclinometry is better for short scale [*Kirk*, 1987]. SPC uses both of these components to mitigate the errors that are caused by each of them. Stereo creates a position in 3-D space using 3–5 pixels from an image. SPC typically sets control points (landmarks) with spacing of 50 pixels.

SPC will use photoclinomerty to fill in the heights of the points between the landmarks. The horizontal ground sample distance (resolution) is typically on the order of the resolution of the images, and we have successfully generated topography at double the image resolution.

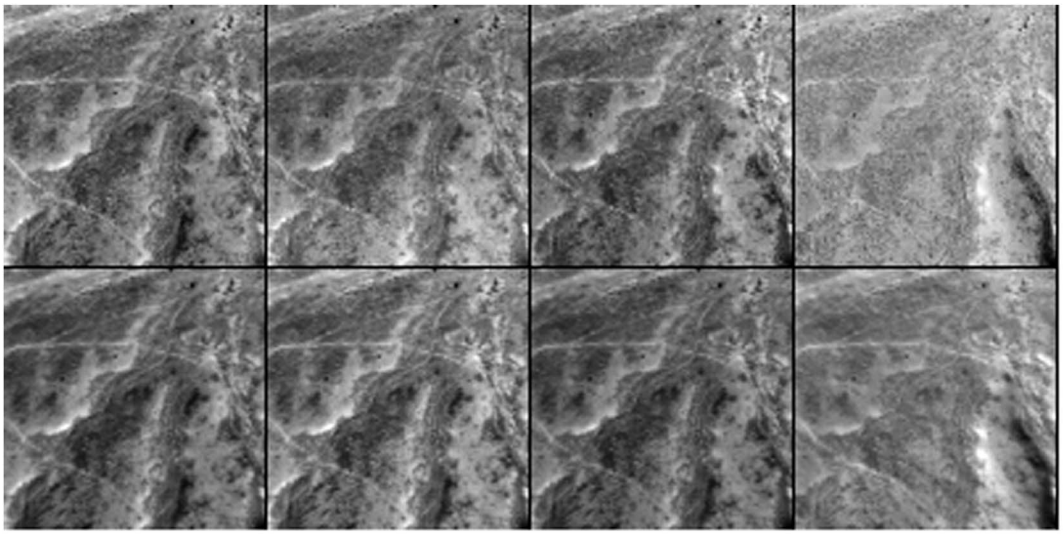

Figure 3. (top row) Source images and (bottom row) rendered images from DEM of Mercator test terrain. In these four images, there is variation in illumination conditions, resulting in a significantly different looking terrain. The bottom row are images based on the shape model illuminated for its corresponding image. By projecting all the images to the same orientation and scale, and illuminating the model accordingly, we can more precisely register images over a significantly wider range of emission angles, thus reducing topographic error. These "maplets" are 30 m across.

The calculations of the points between landmarks are solved in 99×99 grids (aka maplets). Maplets are positioned such that there is 30% overlap between each one, so most heights are solved by multiple maplets. The heights are solved for iteratively until the heights of overlapping maplets converge. This constrains the center and all edges of the maplet to an absolute height solved for by stereo, reducing systematic photoclinometric error. As such, we derived a closer ground sample distance (resolution) than stereo alone can provide, but without less error than a purely photoclinometric solution can provide.

SPC uses a three-step iterative process to derive a DEM: register images, warp the model, and update camera position/pointing. We start with an initial shape model that has very low resolution that provides a starting point for SPC. This shape can come from limb measurements, a radar shape model, spherical harmonics, or if need be, can be derived via a mathematical representation of a triaxial shape. In this case, it was set to a flat surface.

2.4.2. Register

The first step is to register the images to a reference frame (aka MAPLET), which are defined by a control point (aka landmark) and are 99×99 pixels in side. Each image that falls within a MAPLET is orthorectified and projected onto the shape model (Figure 3, top). Associated with each of these maplet views is the DEM with albedo, which is illuminated at the same solar geometry (Figure 3, bottom).

We use both manual and automated tools to coregister the images within each maplet such that the center point of each maplet is in the same location as all the rest. Because the images are orthorectified, it allows us to have different observing conditions for each image, increasing our ability to match identical features, even if the original viewing geometry makes them hard to identify. There are minor improvements in the coregistration of images as the shape model is improved because this reduces projection error and, correspondingly, any misregistration.

2.4.3. Morph

Once the landmarks are registered, we use that information to update the topography or shape model. While the height of the center of the landmark is derived from geometric stereo, we use photoclinometry to solve for the heights of all other pixels of the landmark. We use 2-D photoclinometry using multiple images to remove systematic and minimize nonsystematic error (see section 5.4).

Simple photoclinometry is 1-D and requires the albedo to be a constant so that the only variation in pixel intensity is the topography. For SPC, it creates a full 3-D model that includes albedo. As described previously, for each image within a maplet, we illuminate the shape model to correspond to its observing conditions. Because we

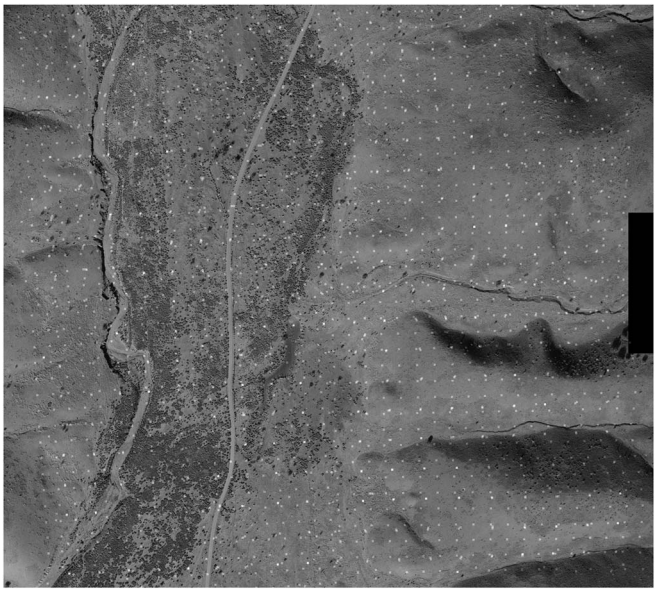

Figure 4. SPC can generate images that depict all of the landmarks that have been registered onto it. The program GEOMETERY will use the detailed sample/line position of each landmark to calculate the position and pointing of the camera.

are fully controlling for topography and albedo, any variations in each image's pixel value (digital number or DN) from its representative shape model's value is an error in the model. The deviations are turned into corrections for both albedo and topography. A configuration file within SPC provides weighting that are applied to the corrections for albedo and topography. The weighting only effects how fast the model converges to a solution not the solution itself.

2.4.4. Camera Position/Pointing

The next major step of SPC is to take the updated heights/positions of the surface's control points (aka landmark) and use it to improve the position and pointing of the spacecraft. Figure 4 shows a bright square at the center of each control point. Within an image, each of these control points (landmark) provides the exact sample/line position of the landmark. These data, along with the angular field of view of each pixel, allows SPC to determine the position and pointing of the spacecraft that minimizes error. Note that a narrow field of view makes it difficult to break the degeneracy where a displacement can be either position or pointing. This is handled by weighting each of these terms based upon the expected errors in positions and pointing. Typically, these estimates are provided by the navigation team. These correction of camera position/pointing and the assumptions made by those generating the topography are where the major differences are between different stereophotogrammetry groups.

For example, during the terrestrial portion of this work, our starting position for the images was the aircraft position at the time the picture was taken, with an estimated 3-D error of ~30 m. During our processes, SPC refined the location of the images by making translations so that the control points have the least amount of error between their position in the DEM and the position predicted by their position on each image. On average, SPC shifted the camera's position (aircraft position) by 9.2 m. The camera pointing was much less constrained with significant errors in aircraft heading, roll, and pitch. We started with an estimate for pointing error of 100 mrad. During our processing, we updated the pointing of the camera and calculated a solution of 41 mrad (or 2.4°) change from our original estimates.

2.4.5. Processing

As the topography was processed, the steps of register/morph/update are iterated until the residuals of the solution approaches a minimum. We used a variety of tools to evaluate misregistration of images, the presence of image artifacts, and control points with low image coverage. Our main analysis uses the program RESIDUALS, which calculates the position on the shape model of each landmark/image combination. For each landmark, RESIDUALS evaluates each image used in determining the height of the control point, which includes using limb observations. RESIDUALS will use the best solution for the position of the control point, then calculates the position each image thinks the control point should be, and the difference between the two is the residual error calculated as root-mean-square (RMS). This is the basis for SPC's "formal uncertainty" reflecting how consistent the model is with the data provided. The RMS error shows how close all of the images are aligned to a consistent model. It is a fully geometric solution based upon the registration of the landmarks (or control points) and the position of the spacecraft.

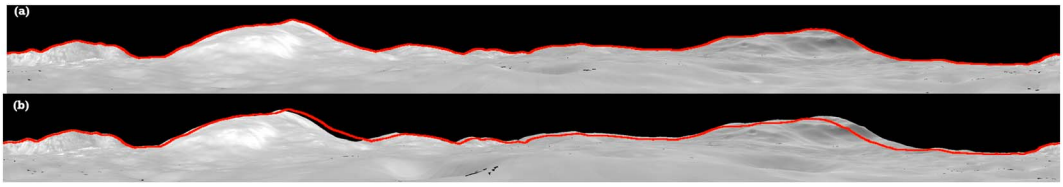

Figure 5. Horizon imagery and traces. The red light denotes the automated horizon-tracing algorithm's results. (a) The Mercator algorithm searches for model locations such that it minimizes the least squares mismatch between the actual horizon trace and all of the horizon traces created from the DEM. (b) Illustration of a mismatch.

2.4.6. Output Products

The main output product for SPC is a shape model, a 3-D set of position vectors that define the heights, which can be converted into a vector/plate model. For this study, we focused on a single piece of terrain with a trivial curvature. For DEM generation of high-resolution topography, we use a 2-D representation that is relative to a single plane, called a BIGMAP. The BIGMAP contains a vector to the center pixel of the DEM, vectors to describe the normal plane, a ground sample distance, and the width of the BIGMAP. The BIGMAP contains a 2-D matrix that contains the height of every pixel above/below the normal plane. Additionally, it also contains a scaled albedo (or scaled average surface reflectance) for each point in the 2-D matrix. SPC generates a BIGMAP by averaging all the maplets that fall within its lateral limits.

2.5. Navigation Technique

The fundamental technique for the Mercator project is horizon matching. It compares synthetic panoramas generated from an ultrahigh-resolution DEM with the panorama that a lander/rover derives from its own imaging camera. The best fitting match would be the best solution for the location of the lander/rover. To implement this technique, we need two data sets: database of synthetic panoramas and a real panorama taken from the unknown location (i.e., images from a rover).

The first step is to use SPC to generate an ultrahigh-resolution DEM of the target region where the lander/rover will operate. This can be done with a variety of orbiting assets and much of that data are already in hand. Once SPC generates a terrain model, we use a Fortran program that does ray tracing of the DEM and generates what the terrain would look like from an observer on the ground (Figure 6). It uses both the albedo and heights for a detailed image of the surface. Both the generation of the terrain model and the generation of the synthetic horizon profiles would be done on the ground before lander/rover operations.

The second data set is the panorama taken from the rover. It is expected that the rover's panorama will come from a series of images that is stitched together to form a 360° panorama. Such an image will need to be resampled to match the resolution of the synthetic panorama.

Then we extract the horizon profile from the real panoramic image (Figure 6). The horizon-tracing program will scan, column by column, from the top until it finds the horizon line using a criterion of five pixels with a DN that are above the background. This is done to reject cosmic rays, hot pixels, or other noise. It also evaluates the horizon profile for a continuous curve ensuring that there are no drastic height changes. A height skip of five vertical pixels is unlikely to be normal topography, and would be an indication of a cosmic ray that was struck by the camera nearly vertically. Regardless, if a discontinuity is detected, the panorama will be rejected and a second one will be taken (Figure 5).

Once we have both data sets, we can search for the rover's location. By doing a systematic search over the synthetic panoramas, we can identify the specific synthetic panoramas that have the best match with the real panorama. The root-mean-square error is calculated from the sum of squares of deviation between the real and synthetic, equation (1). The value of the height for each column of the real horizon (h_{real}), is subtracted from the corresponding height of the synthetic horizon, (h_{syn}), squared and summed. For this project, we created panoramas with 3600 horizontal pixels, n, corresponding to one tenth of a degree pixel resolution, or 1.7 mrad. Epsilon, the error, is the average amount of deviation between the two panoramas in pixels. This gives an accurate measurement of deviation.

$$\varepsilon = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(h_{\text{syn}} - h_{\text{real}}\right)^2} \tag{1}$$

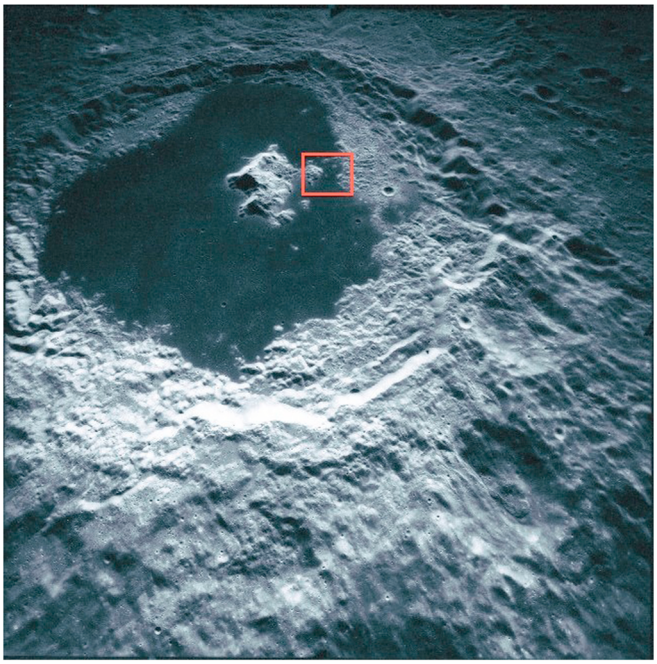

Figure 6. Lunar crater Tsiolkovsky. Red rectangle denotes lunar test area.

3. Tsiolkovsky Test, Lunar Imaging Data

We performed an initial validation of Mercator's panorama matching technique. We used a preexisting ultrahigh-resolution data set of the Moon as a test site. One author (Gaskell) had previously generated a topographic model of the north side of Tsiolkovsky crater, 20.4 N 192 E, on the far side of the Moon (Figure 6). This crater has a dark basaltic floor and bright anorthite outcrops in the central peak. The testing region was 10 km × 25 km.

3.1. Tsiolkovsky Test Setup

The first piece of the test was generating "synthetic" panoramic images that represented our virtual world. For the initial testing, we used the same DEM as the rover images, but we decreased the resolution from 5 to 10 m/pixel. We used this downsampled map to generate the synthetic panoramas.

We calculated the horizon profile for each synthetic panorama, building a database of the horizon profiles for the entire region with a sampling of every 100 m, generating almost 25,000 panoramas that were traced and stored. This allowed for very fast testing because the time-intensive computing was already done for the test region.

For the second data set, we selected 20 "ground truth" locations for the rover, using locations on hilltops, valleys, flat plains, and craters (Figure 7). Because we did not have actual ground truth panoramas, but one created from the same DEM as the synthetic panoramas, we introduced error, which would avoid perfect matching. We used the original ground sample distance, 5 m, to generate the panoramas rather than the downsampled 10 m DEM used for the synthetic panoramas. This resulted in minor errors being introduced from the smoothing effects. Additionally, we offset the ground truth panoramas so that none of them fell on the 100 grid that was used for the synthetic panoramas.

As part of this test we evaluated the effects of equipment and the feasibility of different kinds of equipment. We simulated the rover having a basic 1024 × 1024 CCD camera with a 100 mm lens. We used MATLAB to add noise into images created from the 10 m resolution topography. The noise included a point spread function, hot pixels from the CCD, error from dark current forming hot pixels, and the effect of cosmic rays. We generated 24 images that we mosaic together to form a panorama (Figure 8). These panoramas were evaluated for their usability by the existing code. While most noise that was added did not cause a problem, cosmic rays could cause difficulties for the horizon-tracing algorithm such that we added a 5 pixel threshold. Modifications were made for the horizon tracing to flag an error if the horizon profile failed being continuous, indicating a problem requiring human evaluation.

We calculated the horizon line for the simulated ground truth image (the rover's panoramic image) and ran a nonoptimized test to determine the best fit by testing it against every synthetic panorama. A full test consisted of testing all 25,000 synthetic profiles. Due to the large number of possible positions a rover could be in, we did not generate a panorama for every point, but sampled the testing space ever 10 pixels, or 100 m.

3.2. Tsiolkovky Crater Results

We conducted the test and compared the locations determined by the Mercator algorithm with the locations used to generate the rover's panoramas (Table 3). We show that for most of the locations, Mercator detected

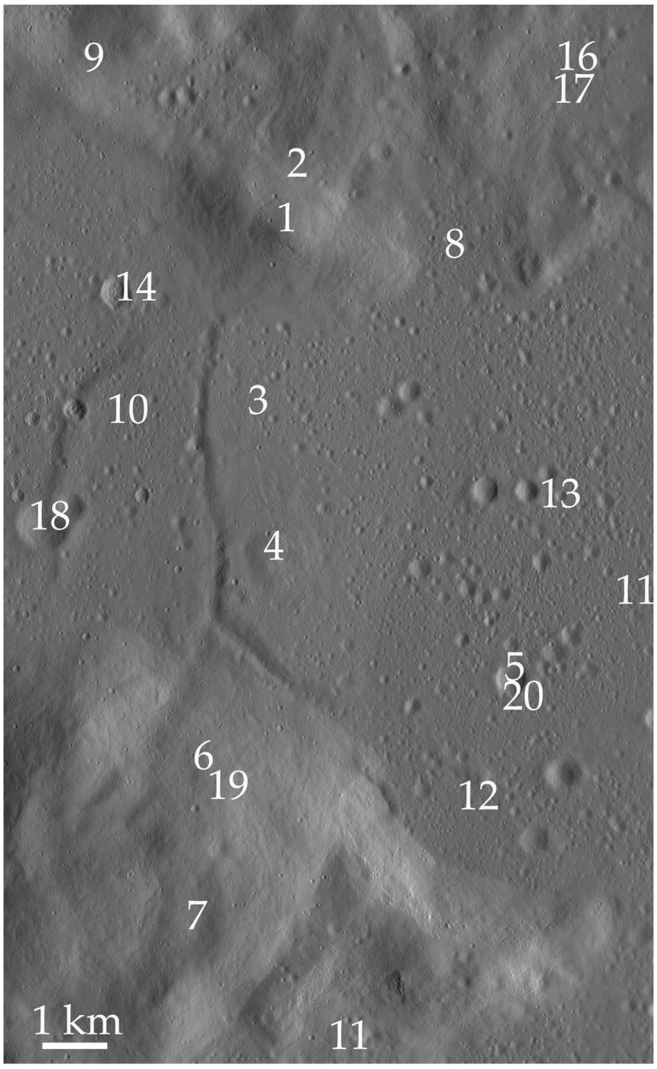

Figure 7. Lunar test study area. Numbers denote test locations (see Figure 9).

the closest possible solution. We did have failures for craters, which was not expected. The small-scale features of craters and variations in their rims are too small to be accurately resolved with this scale DEM (10 m). Thus, the horizon profile would lack the detail needed to detect differences among craters. When we tested a crater, Mercator was unable to distinguish which crater it was in. We expect, although we have not tested, that increased detail, features of the crater walls, specifically the crater's overall size and the rover's distance from each wall, will fix this problem.

We also had three of our test sites at the top of peaks (tests 6, 7, and 19) which provided an extensive view of the surrounding terrain. However, these peaks were close to the edge of the DEM so the ray tracing fell outside of the DEM topography. Thus, the system did not have a usable horizon and failed. This shows that operation near the edge of the DEM may result in errors due to the boundaries of the DEM that could be fixed by increasing the size of the DEM.

4. Ground Truth Test, Terrestrial Imagery Data

The second test we conducted was to provide as real world of a test as possible of the Mercator system with actual panoramic images of the test site. This would replicate the type of errors that a flight mission would be subjected to. We used aerial photography and a team that collected panoramas from the ground.

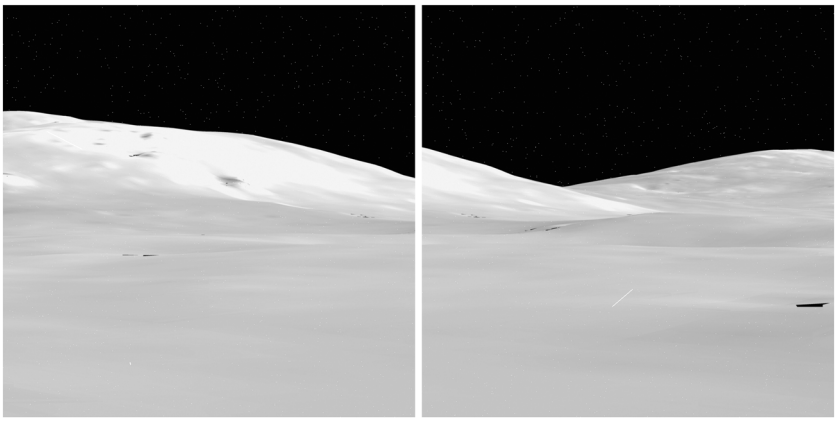

Figure 8. Synthetic lunar images used to build a panorama of the test area. These images have overlap and would be connected to construct panorama.

4.1. Ground Truth Test Location

A key component of this real world test was to simulate the environments that would exist for operations and to identify which types of terrain were problematic. Because of the operational risk and difficulties of landing within a rugged terrain, the locations selected for landers have minimal topography. This translates into test locations that are relatively featureless with low and distant hills. Previous testing shows that horizon matching is effective where they are clearly identifiable and unique features [*Cozman et al.*, 2000], so the Mercator algorithm was specifically designed to operate on terrains with small variations and no significant features.

An additional constraint was the desire to have no easily recognizable terrain features because it would make triangulation possible. Triangulation is a technique in which one can identify one's location by plotting lines of bearing on a map from two landmarks that are accurately identified. Basically, using a compass, a person measures the direction to two mountains that are uniquely identified. Then, on the map, one draws a line of a reciprocal bearing from the two mountains, and his location is where the two lines cross.

We wanted to avoid a test region where this technique could be used. Triangulation is one of the easiest ways to navigate relying on identifying one or two clearly discernable terrain features, such as mountain peaks. This almost always relies on some level of human involvement to aid in identifying which peaks are which. Further, we did not want the computer algorithm to use such key features to the exclusion of other data, such as smaller hills.

For out test site, we used a hilly region near Sonoita, Arizona, about 60 km SE of Tucson, 31.73°N–110.58°W, elevation 1400 m. This region has extensive drainage channels that are wide and well eroded. Vegetation is limited to grasses, mesquite bushes, and yucca plants. The low amount of vegetation reduced the error due

Table 3. Locations and Solutions for Lunar Test in Tsiolkovky Crater

Test Location	Error (m)	Location	Test Location	Error (m)	Location
1	46	Hilltop	11	91	
2	111	Valley north of ridge	12	64	Foothills south of peak
3	42	Plains east of rill	13	36	Near big hill
4	32	Mound east of rill	14	Fail	Middle of a plain
5	Fail	Misshaped crater	15	57	Inside of crater
6	Fail	Horizon beyond map	16	51	Inside large crater
7	Fail	Horizon beyond map	17	30	Region with high rim
8	22	Plains at foothills	18	20	Within crater flood plain
9	56	Valley #2 north of ridge	19	Fail	Horizon beyond map
10	42	Plains west of rill	20	Fail	Inside of crater

Table 4. Aerial Flight Data

Name	Altitude[a] (feet)	Aircraft Path	Time Collected (local)	Number of Images	Original Resolution (cm/pixel)	Downsampled Resolution (cm/pixel)
Flight A	17,700	north-south	10:55	18	25.4	126.7
Flight B	17,800	east-west	11:32	18	25.4	124.7
Flight C	11,300	east-west	2:40	50	13.7	63.4

[a]Referenced to mean sea level.

to the lack of seasonal changes and the effect of wind. Additionally, the aerial photographs and panoramic images were taken nearly at the time, just over 1 month during winter.

The test region had some significant mountains in the far distance, most significantly Mount Wrightson, elevation 2881 m and 25 km to the west. Most of the other mountains were obscured by the local hills and valleys, which was why this test site was selected. However, for the mountains that were visible in the distance, they were removed from the panoramas by hand. This ensured that they did not contribute to the horizon profile at the test location. It is possible to support these mountains, but it would require generating a much larger DEM. More significantly, having a single prominent feature would violate our constraint on prominent features.

4.2. Terrestrial Topography

The first part of the test was to generate an ultrahigh-resolution shape model of this region. This required high-resolution images taken from different angles and times of day. The technique of SPC will solve for both the slope and albedo (surface reflectance) of the terrain. We wanted to ensure we tested both the photoclinometry components as well as the stereo components. Stereo imaging is used extensively in planetary science, e.g., on Mars Pathfinder, MER, Curiosity, but it requires a restricted set of stereo pairs to be effective. The power of SPC comes from its ability to use images in a variety of orientations, resolution, and illumination conditions.

We contracted with Cooper Aviation of Tucson, AZ, to collect the images at two different times of the day (Table 4) at and two different altitudes on 29 November 2011. Flights A and B provided 1:26,000 scale and flight C provided images at 1:13,000 scale. The images were taken with a Wild RC30 camera with a Wild Universal Aviogon/4-S 153 mm lens camera, then scanned by Cooper Aerial. They were delivered as 22,837 × 23,217 pixel TIF images. We downsampled the images by a factor of 4 to aide processing. The working images have a resolution of 125 cm/pixel and 63 cm/pixel, respectively (Figure 9).

The imagery covered a 10 km × 10 km region centered at 31.73°N 110.58°W. There was a 60% overlap down track and a 30% overlap cross track. We sequestered every other image in order to limit the amount of overlap that was provided to SPC. This was done to evaluate the effectiveness of SPC in dealing with images from different altitudes, locations, and times of day (Figure 10). The full resolution images and sequestered images could be used in the future to evaluate how much these data would improve the topography (Table 5).

4.2.1. Terrain Calculations

The next step was to create synthetic panoramas for every location in our 4 km × 4 km test region. Panoramas were generated using the "panorama" routine in SPC (Figure 5). For this project, we created panoramas with 3600 horizontal pixels, corresponding to one tenth of a degree pixel resolution, or 1.7 mrad.

For the generation of synthetic panoramas, we created panoramas every meter of the 4 km × 4 km BIGMAP (Table 1). Ultimately, we set the observer height to be 1.25 m, while higher than small rovers is reasonable for this test. It is important to note that higher-resolution topography would allow for a lower observer height. We used the cluster computer at the Planetary Science Institute to conduct the batch computation of the 16 million synthetic panoramas. Next, we used the automated traceHorizon routine to identify the horizon profile of each panorama and stored the results (Table 6).

4.3. Terrestrial Real Panoramas

The next major stage of the project required ground truth is the actual images from the test site. On 2 January 2012, our team went out to the test site to collect images for the panoramas. We used a Canon Digital Rebel

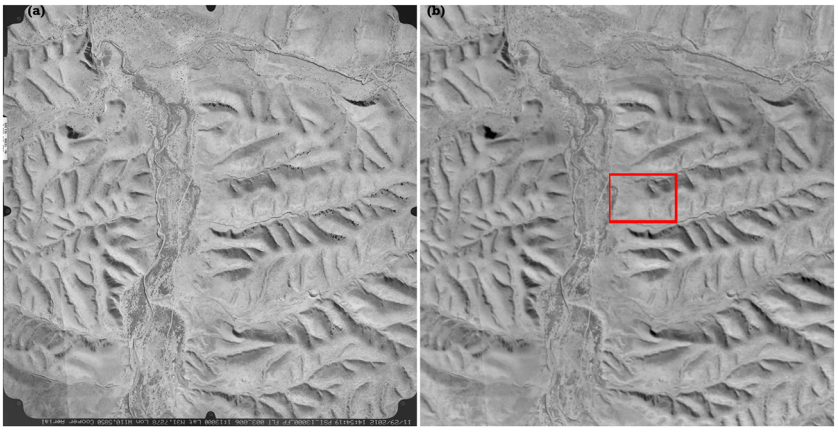

Figure 9. (left) Aircraft imagery and (right) synthetic terrain with albedo of the study area southeast of Tucson, AZ. The rectangle demotes the test site (Figure 11).

Ti1 with a 24–55 mm lens with a 1.25 m tripod to collect the images. The camera was mounted so the focal plane was near the tripod's axis of rotation. We took nine images with the lens at 24 mm, providing for extensive overlap for stitching. We also collected a GPS set of coordinates for each location. We processed panoramas for 13 sites (Figures 10 and 11).

The images were stitched using the software PTGui, which is a commercial graphical interface for the open source Helmut Dersch's Panorama Tools. We hand-selected the control points for every image based on the horizon profile of each image to reduce distortion. From this, we generated panoramas with 18,700 horizontal pixels. We shifted the panorama such that true north was at the center of the image to align with the synthetic panoramas. Then the panoramas were downsampled to match the synthetic panoramas with a pixel width of 3600.

Once the panoramas were aligned, we used Photoshop to remove the sky. Unlike images of the Moon or most solar system bodies that this technology would be used upon, Earth has clouds that make an automated removal routine difficult as discussed by *Cozman et al*. [2000]. Further, glare from the Sun with a low level of contrast in the horizon added to the difficulty. Finally, there were long grasses, yucca stems, and leafless trees that were well resolved in the panoramas that were not resolved in the imagery used for SPC. Thus, the shape model did not maintain a signature for these objects while a horizon profile would. While not generally significant, in some places we removed vegetation by hand.

Once the sky had been removed, we dealt with the next problem, tilt. We did not take the instrumentation necessary to ensure the camera had no tilt and hence needed to remove it at this stage. We created a routine that minimized the error by varying the tilt and axes for the image based upon the known horizon shape for each test location. Then we mathematically shifted the horizon profiles to remove the effect of camera tilt.

4.3.1. Batch Comparison

Now that the images had both bearing and tilt corrected, they could be compared to the 16 million synthetic panoramas. Because such a detailed search is time consuming, we chose to do a wide field search of the panoramas at 10 m spacing. These gave preliminary results that were good, Table 2. An example of the quality of fit is Figure 12.

While the 10 m results are good, they can be improved. Specifically, we want to ensure that we do not select a location that is only a local minimum just because we sampled near the best part of the local minimum. To deal with this, we tested subregions that had good matches. We created a routine that identified about five non-contiguous regions that had the best matches. We tested these subregions at 1 m sampling, usually only a few thousand test cases. The higher-resolution test cases almost universally identified a better solution for our data, and in one case, it identified the global maximum rather than a local maximum that the lower resolution search had.

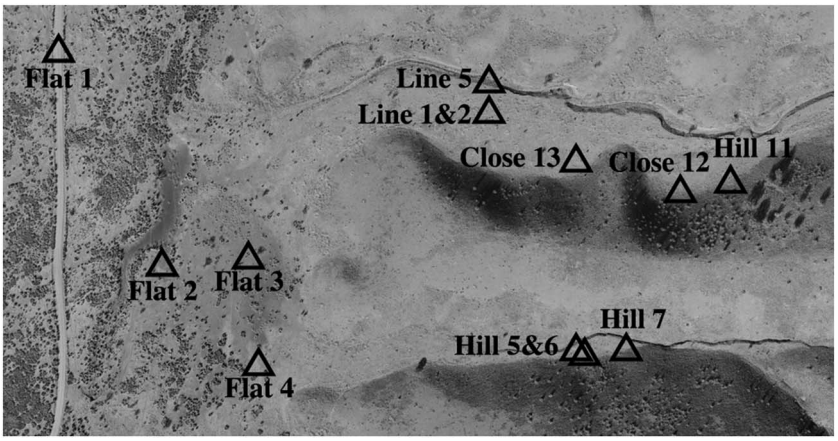

Figure 10. Test site within the study area. Triangles denote where panoramic images were taken by the ground team. GPS coordinates were obtained.

Once we ran the batch comparison, we evaluated how close the "best" solution was to the correct location (Figure 12). These results show that Mercator is capable of identifying its location to an error of 3 to 10 times the resolution of the topography.

4.4. Terrestrial Results

The performance of the Mercator algorithm has been reviewed and is robust. For our data set, we have been able to locate the correct position within an average error of 6.2 m with no data locations more than twice that distance away. That means that within our test, position can be determined to within an average of 6 m. Table 7 shows a qualitative assessment of the different types of terrain upon which Mercator was tested.

4.4.1. Flat, Distant

A suite of testing was performed in a flat region with hills (topographic highs) in the distance, locations 1–4. For these tests, we were in the middle of the wide and shallow valley that had a meandering streambed and a road. The hills that were controlling the horizon were about 0.5 km away. Tests done from these regions produced a large number of possible matches; however, the best matches were at the correct solution, with an average deviation of 8.2 m. As seen in Figures 13a and 13d, there is a wide and broad field of good matches with an increasing solution toward the actual point.

4.4.2. Hill

The specific testing site was chosen to establish the limits of what Mercator could accomplish. One of the major scenarios was to find an area where there were numerous hills that had limited distinguishing features. Most of the test regions, especially away from the riverbed, were hilly without significant prominent features —basically, things looked very similar. We needed to avoid having a single large and easily identifiable feature, such as a mountain or butte, because areas with those types of features would result in a test that reduced to triangulation. While this would clearly be useful, it did not test the more difficult aspect, repetitive, and nondescript relief (rolling hills).

This region enabled us to test that small variations in the position and height of hills relative to other hills could indeed constrain the location (Figures 13b and 13e). The tests that were done for this region, 5–7, were successful. The average error was only 5.0 m, which is on the order of our ability to measure and register the location to our reference maps.

Table 5. DEM and Panorama Data

	DEM Resolution	DEM Area	Panorama Spacing	Simulated Camera Height	Mean Error
Step 1	30 m/pixel	10 × 10 km	N/A	N/A	N/A
Step 2	10 m/pixel	4 × 4 km	every 2 m	2 m	75.8 m
Step 3	1.5 m/pixel	4 × 4 km	every 2 m	2 m	17 m
Step 4	0.5 m/pixel	4 × 4 km	every 1 m	1.25 m	6.2 m

Table 6. Summary Errors—Residual Errors Are Comparable to Errors in Registering the Aerial Imagery to the GPS Grid

	Type	Pass2	Pass3	Pass4 Every 2 m	Pass 4 Every 1 m
1	Flat	12.4	3.6	31.1	6.7
2	Flat	18.0	2.2	4.2	7.6
3	Flat	19.4	13.6	2.2	3.0
4	Flat	150.6	11.7	6.4	5.0
5	Rill	11.7	20.0	3.6	6.3
6	Hill	9.1	45.0	13.3	3.2
7	Hill	14.6	13.3	35.2	2.8
8	Line	34.7	17.0	14.0	10
9	Line	13.5	17.8	12.0	9.5
10	Line	16.4	54.1	9.8	11.7
11	Hill	12.6	10.8	5.4	5.0
12	Hill, Close	21.6	12.0	5.4	6.7
13	Hill, Close	36.4	8.9	3.0	3.0
Average		28.5	17.7	11.2	6.2

4.4.3. Hill, Close

The next test we tested the effect of being close to a hill. In these tests, we located two similar looking crescent-shaped hills that had been eroded (Figure 13f). These tests were to stress Mercator in the following criteria. First, a large part of one half of the distant horizon was blocked. Second, the close horizon was very large, such that any deviation in our shape model would create a drastic effect in how high the hill went up, which could result in significant error and possible location mismatch. Finally, and most importantly, because there were two nearly identically crescent-shaped hills, we thought it was possible that Mercator might find the wrong area.

Mercator was able to identify the correct location for all three of the testing locations. For tests 12 and 13, there was a secondary solution of lower quality at nearby crescent-shaped hill; however, the correct hill solutions were significantly better. Unless intimately familiar with the local terrain, it is unlikely that a human would be able to distinguish which of the two hills they occupied; the Mercator algorithms had no problems. The average error was 7.7 m.

4.4.4. Rill (Simulating Small Crater)

In our previous testing, we saw that Mercator had insufficient horizon cues at the 5 m resolution to have successful location determination within craters. To see if Mercator could deal with the lower height, especially relative to the surrounding territory, we performed a test within a sharp gulch (or what might be called a tiny rill on the Moon). The gulch was about 1.5 m wide about 0.5 m deep (Figures 13c and 13g). The gulch had grass growth that decreased the clarity of the surface. As such, SPC did not compute the depth or sides of the gulch as sharply as they should have been.

Our original tests, which were based on low-resolution topography (pass 2,) did not perform well. The vertical relief was too large for the horizontal ground sample distance, which resulted in a discontinuous horizon profile. We identified that not only was the location of the camera too high but the resolution of DEM we were using was too low, which was 1.5 m. Subsequently, we increased the DEM resolution to 0.5 m, which resulted in a much better match. We ascribe this improvement to two factors. First, the camera height (1.25 m) was

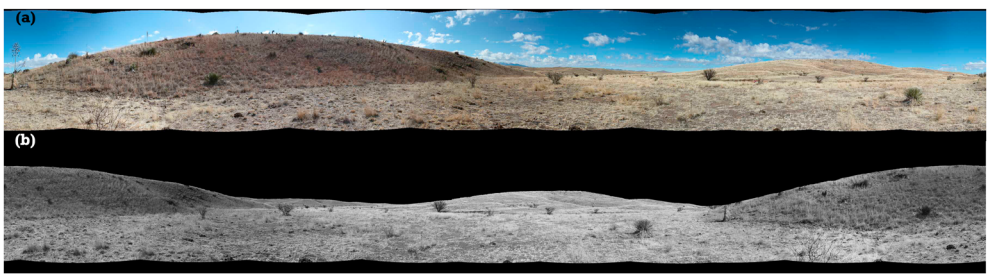

Figure 11. Example (a) raw and (b) processed mosaic. We removed clouds and foliage that extended into the sky and shifted the image so north was at the middle of the panorama (pixel column 1800).

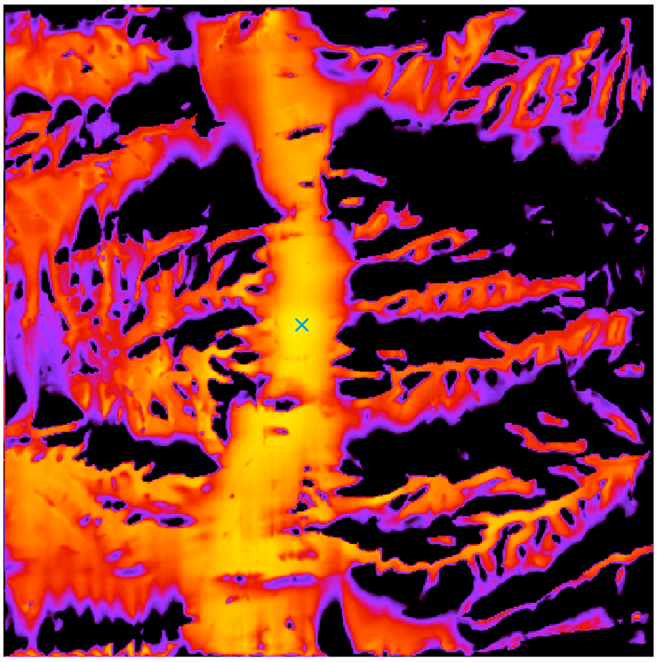

Figure 12. Solution Map—For each test site, we take the real panorama's horizon profile and compared it to a synthetic horizon profile from every location on the DEM (over 16 million possible locations on the DEM). This image is a spatial plot of the error between the real and synthetic panoramas' profiles. The solution map color codes the error. Black values indicate very poor agreement (larger residual errors). Bright values indicate small errors, so a good agreement. Because the synthetic panoramas are discretized, a perfect solution is essentially impossible. The width of the image is 4 km.

lower than the DEM resolution. Secondly, before increasing the resolution, the nearest hill had the wrong height; i.e., it was lower than it should have been. Once the DEM resolution was increased, the hill raised in height about 2 m (Figure 14).

4.4.5. Line Test

One concern with the Mercator algorithm is getting false positives (high correlation) that are not the correct solution. To test this, we took a sequential series of locations to allow for combining solutions and improving the accuracy of Mercator. We assumed that Mercator could result in many locations where there would be local minimums of okay matches. In fact, it is possible that the best local match could end up not being the correct solution if the terrain is similar.

Because finding the correct solution is key to navigation, we conducted a test to see if using two ground locations a known distance apart greatly improved Mercator's ability to determine its location. The idea is that it may be possible to have two locations where the horizons are similar because the shape and slopes of different hills are similar; however, if you move 5 m in a known direction, it becomes extremely unlikely that the

Table 7. Terrain Types, Pros, Cons, Performance With Current Algorithm

Terrain Types	Pros	Cons	Performance
Hill	Frequently provides well constrained solution	Can get local minima that may not be the global minima	Good
Flat	Unlikely to have local minima	Creates large number of reasonable matches	Excellent
Hill, Close		Extreme topo may result in horizon going outside camera FOV Small errors in the DEM results in large errors	Good
Rill		Needs high resolution Small errors in the DEM results in large errors	Marginal

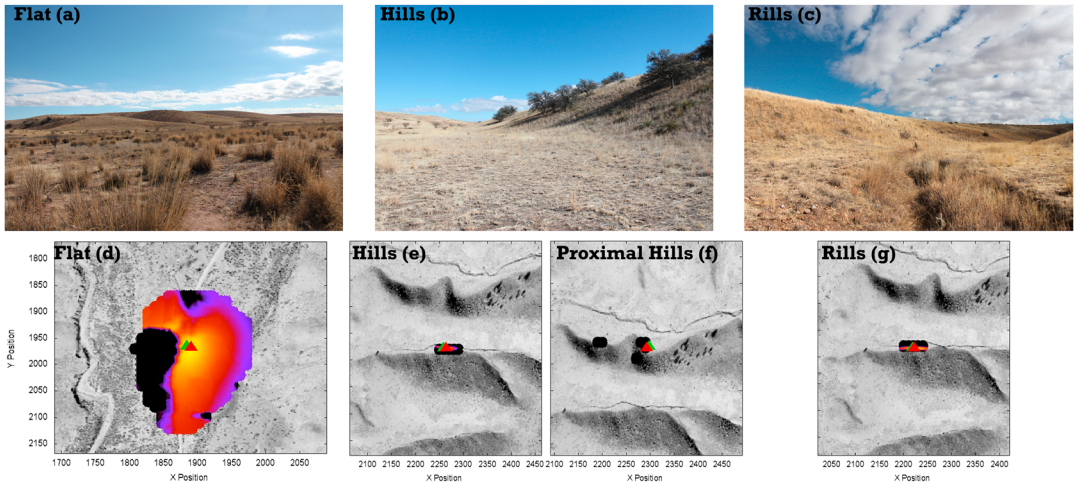

Figure 13. Test site showing (a–c) representative images from the ground and (d–g) corresponding aerial images with superposed solution maps. Note the broad peak of the solution space in flat terrain in contrast to other terrains studied. The green triangle is the correct solution, and the red triangle is the best solution. X and Y pixel scale is in meters.

similarity would remain. That would require the hills to be the same size and distance in all directions, a scenario that is implausible.

For this test, we took a panorama at a well-defined location in a hilly region. Then we took additional panoramas at 2 and 4 m north of our starting location. These panoramas simulated a rover that gets an initial location then moves to see which matches remain consistent—only the correct match would remain a good fit.

This test showed that this technique works well; however, it was not a full test. We did not have a situation where the local minima were not the correct location, so we did not fully test the scenario. Regardless, when we tested other "good" matches, and then tested the location 2 and 4 m farther north, the other matches' solutions became significantly poorer.

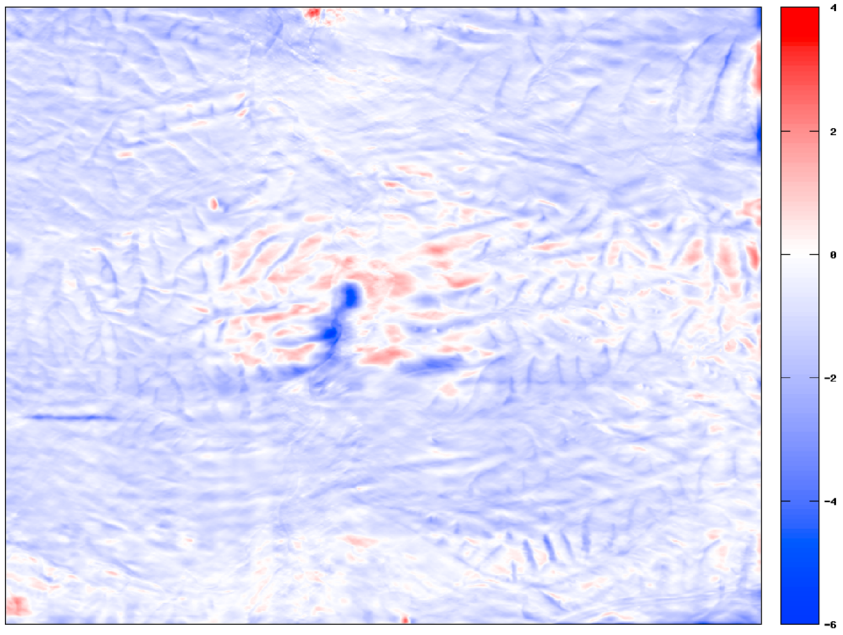

Figure 14. Impact of improving DEM from 1.5 to 0.5 m. The adjustment is less than 1 m overall, but in some places the smaller ground sample distance resulted in an improvement of the shape of up to 4 m (2–3 times the image resolution). This image shows the subtraction between the two DEMs. Increase in terrain is denoted in red, with full red at 4 m. Decrease in terrain is denoted in blue, with full blue at −4 m.

5. Discussion

5.1. Ground Truth Error

There are a variety of error sources for the positional information. Most of the ground points were taken with a basic GPS, which had an error of approximately 2 m for each data point. Additionally, we had to convert the latitude and longitude into the *x* and *y* coordinates system used by SPC. This required plotting the data in Google Earth as latitude/longitude, then identifying the corresponding location on an image projection of the BIGMAP search region. The translation from latitude/longitude pairs to SPC *x*/*y* position was the largest source of error and was strongly influenced by the lack of nearby features. We estimated the error to be 4.1 m for registering the GPS latitude/longitude to image features to the DEM coordinate system. This is a significant portion of error from the Mercator panorama matches. Results suggest that Mercator's location error was on the order of the errors inherent in the resolution used and registration method.

5.2. Tilt

One significant source of error that we had to remove was the tilt of the images taken with the camera. During panorama image collection, there was no technical equipment to ensure the images were perfectly horizontal. We removed the tilt by measuring the amount of tilt required to match real horizon with the correct synthetic horizon. Then we applied the correction, usually between 2 and 5°. From there, we conducted the search over the entire test region.

Ensuring that there is no tilt is a critical element of this technique. The test shows a tilt of a few degrees decreases the accuracy of the solution by tens of meters, or even creates a false positive if there is a similar terrain elsewhere. For this reason, the optimum design for equipment using Mercator would have a star tracker. By using a star tracker, we can affix the spacecraft's orientation to identify what tilt there is relative to the stars, and thus, the surface of the object in either center of figure, center of mass, or local normal. Additionally, it will accurately determine "true north" (or a reference direction) to ensure that the panoramas are aligned to true north and tilt removed. Use of a star tracker would require an initial latitude/longitude position to solve for the ambiguity between tilt and orientation on the body.

5.3. Improved Resolution

As mentioned, during our testing, we improved the resolution of the DEM from 1.5 m to 0.5 m. This was based on the same imagery, only continuing the processing to create a higher-resolution topography. The process of increasing the resolution was based on two steps, creating 1.0 m resolution maplets followed by 0.55 m maplets. We created 362 additional maplets at the 0.55 m resolution that covered the test region with 50% overlap with each other. We used automated tools to register the images to the existing topography then fixed registration problems individually by hand until all the maplets were registered producing a map that had agreement among the maplets with an RMS error of 3.8 m.

The working resolution of the images from 17,000′ MSL are half the pixel size (1.25 m) of the DEM we created, The working resolution of the images from 12,000′ MSL are only slightly lower than the DEM (63 cm). We note that working with a DEM with subpixel resolution and image registration continued to improved the topography of the test region. Figure 14 shows the difference in height between the 1.5 m DEM and the 0.5 m DEM. In our specific testing region, one hill increased in height by 2 m, which improved the shape of the synthetic panoramas, which decreased the location error from 11 m to 6 m.

5.4. Vegetation

Another source of error is how grass and other plants affected the surface topography. The grass, yucca, bushes, and trees in this region and time of year have sparse leaves and branches. As such, they provide a partially solid object that casts a diffused shadow. If a tree or bush is significantly full, it would cast a full shadow allowing SPC to neglect the shadowed region in its calculations and solve it as a boulder. However, for our testing the shadows that were created were diffuse, only decreasing the amount of sunlight rather than obscuring it. SPC is not designed to account for this effect, and as such, there are small topographic bumps around the base of trees. They are not significantly high, but they do cause the ground to be higher than if they were not present.

For our purposes, the additional height was minor because the objects were mostly diffuse. Further, the trees and bushes were seldom on the horizon itself. Thus, the horizon profiles were not significantly impacted by

Table 8. Types of Error Within Photoclinometry

Problem	Systematic	Type Error	Calculations	Mitigation
Image noise	Nonsystematic	Gaussian	error/sqrt (*N*)	Multiple Images
Offset or pedestal	Systematic	1-D only	N/A	Calibration
Atmospheric clouds	Nonsystematic	Removal by hand	N/A	Hand removal blemishes
Atmospheric dust	Systematic	Gaussian	error/sqrt (*N*)	Multiple Images
Albedo	Systematic	1-D only	N/A	N/A
Orthogonal	Systematic	1-D only	N/A	N/A
Photometric function	Systematic	1-D—Slope 2-D—Gaussian	error/sqrt (*N*)	
Misregistration	Nonsystematic	Stereo	1/tan (theta)	
S/C Position	Nonsystematic	Stereo	1/tan (theta)	Bundle adjustment or position/pointing update

trees and shrubs. After reviewing both the synthetic and real panoramas, it does not appear that trees and shrubs affected the results. We did note that the grass in the gulch (or rill) did affect its depth and width, and especially reduced the clarity of the gulch's walls.

5.5. Clouds and Dust

On Mars, we might expect dust storms and clouds to obscure or confuse the horizon. During dust storm cases we do not expect the algorithms as currently developed to function or to function with the precision we achieved using atmosphere removed data. Under such conditions navigation updates would be delayed, leading to less accurate geodesy tie point determination. Other navigational uses would be more tolerant of larger errors, though at some point the vehicle would need to stand down until conditions improved. For example, if a vehicle was navigating in dangerous terrain where a few tens of meters error would be catastrophic, a clear view of the horizon would be a prerequisite for operations.

5.6. Errors Within the Photoclinometry Portion of SPC

5.6.1. Image Errors

Jankowski and Squyres [1991] defined seven types of error that effects photoclinometry and broke them into three categories (Table 8). The first category is error from the images themselves. These errors include variations in the image due to noise, such as detector read noise, analog-to-digital converter noise, and photon-statistic ("shot") noise. This type of error is typically a Gaussian noise error between pixels. While most data used from missions are calibrated images, SPC sometimes uses the uncalibrated images, which means that they contain additional noise, such as hot pixels and cosmic rays. We handle this error in two ways. First, systematic errors, such as hot pixels, can be removed for all images by using a global "blemishes" file that masks those pixels out. Second, random errors on individual images, such as cosmic rays, can be removed on an individual basis with a similar blemishes function.

5.6.2. Nonsystemic Errors With Photoclinometry

The second category of noise is nonsystematic noise. Single-image photoclinometry is highly sensitive to this type of error. The variations in pixel DN results in topographic fluctuations on a pixel-by-pixel scale (i.e., there will be bumps in the DEM that are not really there). High-quality calibrations are essential for single-image photoclinometry.

Stereophotoclinometry is able to minimize these errors by the use of multiple images. SPC uses a large number of images per landmark, frequently over 50. The nonsystematic nature of the noise results in a Gaussian distribution of error, in which the mean value of this error converges on zero and effectively removed. Thus, because we have the camera in numerous orientations with different pictures, the same pixel is never in the same place so artifacts such as cosmic rays, shot noise, hot or weak pixels are softened.

These tests do not have the same number of unique images to minimize noise error. Most landmarks have four to six images. We evaluated the pixel-to-pixel noise that is present in our images. These images were scanned and stored at TIFF images. We used the film alignment marks, which are uniformly black, to calculate

the standard deviation of the imaging process. The standard deviation of these blank regions is 0.52 DN of a 255 DN dynamic range (8 bit). This means that the pixel-to-pixel error is 0.203% DN. Using this noise error, then the use of four images, would reduce the noise to 0.1%. Using six images would reduce the effect of the noise to an error of 0.08%. This amount of error (six images) would result in a point-to-point noise-based variation of a slope of 0.05° for flat areas, and based on the resolution of the imagery, the height error would be 5.2 mm.

An additional source of error within images is contained in a background offset (pedestal) used to ensure that images values vary around a background value rather than being truncated at 0 (i.e., it provides an offset to avoid having acceptable random variations truncated at 0). This offset is easily removed during calibration. For SPC, any offset is automatically removed because the albedo values of the images are scaled from 0 to 2 with the mean at 1.

Finally, recorded photons are recorded into discrete DN values. Older imaging system used 8 bits that provide a limited DN range (0 to 255). Most imaging systems today use 14 or 16 bit images, allowing 16,384 and 65,536 values, a greatly expanded range of values. The error in slope that would occur due to quantization for a 16 bit image is less than 0.02° in most circumstances. However, the images used in Mercator are 8 bit images, meaning that the quantization error would have a mean of 0.5 DN, which is similar to the noise error. The error in slope caused by this would be 0.12° for typical DN values. Again, using multiple images, the error will be reduced to a Gaussian distribution, and for typical landmarks it would be a 0.05° slope error or a 5.2 mm height error.

5.6.3. Systematic Errors With Photoclinometry

5.6.3.1. Atmospheric

Spacecraft missions that have used SPC are done on objects without an atmosphere. However, the ground test for this project was done on Earth where atmospheric effects can play a role. If the atmospheric effects are consistent over the entire surface, then they can be efficiently removed as a scaling term of the data (e.g., the reduction in flux would be the same as having a shorter exposure). However, atmospheric effects that vary with position can create significant problems. This would require preprocessing those images to remove the effects by hand, or more likely, removing the degraded portion of the images using the blemishes tool. Fortunately, for this study the flights were conducted on a clear day without appreciable dust, haze, fog, or significant thermal heating to produce optical variations. For this study, we are assuming the effect to be minimal because no systematic effects have been detected.

5.6.3.2. Albedo

Single-image photoclinometry makes the assumption that the entire surface has a near-uniform albedo. As such, the variations in DN are solely based on topographic variations. This assumption has greatly reduced the effectiveness of photoclinometry because so few surfaces have such constrained surface properties.

Stereophotoclinometry has as a core element that it solves for the "albedo" of every pixel of each maplet. While stereo can get a solution with only two images, SPC requires three when we solve for the interstitial topography. This is because we solve for *x* slope, *y* slope, and albedo. Note, this albedo is not a bolometric Bond albedo or geometric albedo. It is the average surface reflectance of a specific point on the surface, basically the average *I*/*F* based on all images that cover that point. Because SPC solves for the albedo of the surface, the error normally produced by single-image photoclinometry is completely removed.

5.6.3.3. Orthogonal

Photoclinometry has previously focused on solving for slopes along a single line—1-D photoclinometry. The researcher will have to identify the direction of the slope, then conduct the study along the slope, which should be orthogonal to the strike of the surface. However, any deviation from the actual maximum slope, or a change in the strike orientation, would result in an error.

Stereophotoclinometry has focused on solving the problem by conducting the photoclinometry using a 2-D grid. Because we solve for both slopes in two orthogonal directions (*x* slope and *y* slope), we can fully constrain the topographic variations that come from *x* slope and *y* slope.

5.6.3.4. Photometric Function

Photometric functions are an important correction to the image data because light does not reflect at the same ratio in all directions. For single-image photoclinometry, photometric functions can be a source of significant error. While other types of error are mostly random noise, photometric error can result in errors that compound over the entire solution. These invalid corrections would result in DN values that are consistently

too high or too low and produce invalid slopes that appear to be consistent with the data. In other words, photometric corrections would result in terrain that is consistent with the images, but not representative of the actual topography because it overestimates or underestimates the slopes.

Multiimage photoclinometry solutions, such as SPC, are less affected by error of photometric functions. SPC makes photometric adjustments when illuminating the model to which the imagery is compared. The software has several different models built into it and is changed if a spacecraft mission determines a better model. However, the main model used by SPC, and the one used in this study, is a model that combines the effects of Lambertian and the Lommel-Seeliger photometric functions [*McEwen*, 1996] While this model is derived for the Moon and Mars, it has been used in numerous projects because it reproduces the observed surface effectively [*Gaskell et al*., 2008, 2010, 2011; *Gaskell*, 2012, 2013a, 2013b, 2013c, 2013d; *Gaskell et al*., 2014]. Further, the choice of a photometric function is noncritical when using multiimage photoclinometry because errors in the adjusted images do not propagate systematically. Multiimage photoclinometry uses observations from numerous illumination and observation geometries, and as such, the phase, incidence, and phase angles vary in a potentially random manner within a limited angular range. These variations result in the different images DN values for each pixel having a random photometric error applied to them. As long as a large number of images do not have the same illumination/observation geometry, then the errors can be treated as noise.

An additional constraint that SPC uses are overlaps. In general 1/9th of all points are in two or more maplets. As such, all of these points have an additional set of equations to condition their height. We process SPC until the maplets have the overlapped region in agreement. As such, any error is limited to a small random walk of no more than 8 pixels. This results in very high pixel-to-pixel precision (or local accuracy).

6. Conclusions and Recommendations

This work shows that a high-quality resolution is critical for navigation by panoramas. Mercator was able to fix the location of actual locations with in 6 m, which is close to the error of the latitude/longitude derivation system that was used. It was robust, able to distinguish minor variations in the horizon lines that were too subtle for human operators to notice.

This system would enable static landers (provided they have a camera) to determine their position readily. Operations for rovers that utilize Mercator would be greatly simplified. Mercator would provide an independent means of establishing a rover's position, one that is significantly faster than that done by humans. With this autonomous navigation approach, support for movement operations would be significantly reduced and the vehicle would most likely realize an increase in range and operational efficiency. Mercator would be especially useful for missions that will operate on the lunar far side, out of direct line of sight with the Earth. It could also provide a safety measure during a rover's transit because it could easily test that it is making the expected progress during the transit without human involvement. Finally, astronauts could use Mercator to easily determine their position without the use of an extensive navigational support system, such as a GPS-like constellation.

One possible improvement is the horizontal resolution of the panoramic images. Currently, the panoramas that we tested had 3600 height points, or 10 points per degree. For actual navigation, it may be useful to increase the horizontal resolution of the panoramas, which would allow for a higher level of accuracy. However, it is likely that unless the DEM and other steps are improved, the current horizontal resolution of the panoramas is sufficient.

The work, specifically the work on the Moon, shows that craters significantly reduce the ability of horizon matching for location identification. It is likely that a better DEM would allow both the identification of which crater the rover was in, as well as how close it was to each wall. However, this has not been tested and should be done in future works.

Acknowledgments
This work is supported by NASA contract NNH10CD16C. We thank Leonard Vance and Ross Rosenwald for their support as project managers during phases of this project. Wayne Jouse is also thanked for the technical support providing the constraints to our design.

References

Arvidson, R. E., et al. (2004), Localization and physical properties experiments conducted by Spirit at Gusev crater, *Science*, *305*(5685), 821–824, doi:10.1126/science.1099922.

Cozman, F., and E. Krotkov (1996), Position estimation from outdoor visual landmarks for teleoperation of lunar rovers Appl. Computer Vision, Proc. 3rd IEEE Workshop, 156–161.

Cozman, F., E. Krotkov, and C. Guestrin (2000), Outdoor visual position estimation for planetary rovers, *Auton. Robots*, *9*, 135–150.

Gaskell, R. W. (2012), SPC shape and topography of Vesta from DAWN imaging data AAS/Division for Planetary Sciences Meeting Abstracts 44, 209.03.
Gaskell, R. W. (2013a), Gaskell Mimas Shape Model V2.0 NASA Planetary Data System 206.
Gaskell, R. W. (2013b), Gaskell Phoebe Shape Model V2.0 NASA Planetary Data System 207.
Gaskell, R. W. (2013c), Gaskell Tethys Shape Model V1.0 NASA Planetary Data System 208.
Gaskell, R. W. (2013d), Gaskell Dione Shape Model V1.0 NASA Planetary Data System 209.
Gaskell, R. W., et al. (2008), Characterizing and navigating small bodies with imaging data, *Meteorit. Planet. Sci.*, *43*, 1049–1062.
Gaskell, R. W., O. S. Barnouin, and D. J. Scheeres (2010), The NEAR Shoemaker landing on Eros Lunar Planet. Sci. Conf. 41, 2093.
Gaskell, R. W., E. E. Palmer, N. Mastrodemos, O. S. Barnouin, L. Jorda, and A. H. Taylor (2011), Mercury and Vesta—Preliminary shape and topography Abstracts A1576 presented at 2011 Fall Meeting, AGU, San Francisco, Calif., Dec.
Gaskell, R. W., L. Jorda, E. Palmer, C. Jackman, C. Capanna, S. Hviid, and P. Gutiérrez (2014), Comet 67P/CG: Preliminary shape and topography from SPC AAS/Division for Planetary Sciences Meeting Abstracts 46, 209.04.
Jankowski, D. G., and S. W. Squyres (1991), Sources of error in planetary photoclinometry, *J. Geophys. Res.*, *96*(E4), 20,907–20,922, doi:10.1029/91JE02209.
Kirk, R. L. (1987), A fast finite-element algorithm for two-dimensional photoclinometrry PhD thesis, California Institute of Technology.
Li, R., F. Ma, F. Xu, L. H. Matthies, C. F. Olson, and E. A. Raymond (2002), Localization of Mars rovers using descent and surface-based image data, *J. Geophys. Res.*, *107*(E11), 8004, doi:10.1029/2000JE001443.
Li, R., et al. (2006), Spirit rover localization and topographic mapping at the landing site of Gusev crater, Mars, *J. Geophys. Res.*, *111*, E02S06, doi:10.1029/2005JE002483.
Li, R., et al. (2011), MER Spirit rover localization: Comparison of ground image- and orbital image-based methods and science applications, *J. Geophys. Res.*, *116*, E00F16, doi:10.1029/2010JE003773.
McEwen, A. S. (1996), A precise lunar photometric function Lunar Planet. Sci. Conf. 27, 841.