Pose Estimation Using Starfield Occlusion

John Kua

CMU-RI-TR-08-21

May 2008

Submitted in partial fulfillment of the requirements for the degree of Master of Science in Robotics

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Abstract

This thesis introduces and demonstrates a novel localization method for the unique environment that is found in permanently dark lunar craters. The method makes possible near GPS-quality localization without any artificial navigation infrastructure, performing an order of magnitude better than existing methods. This is inspired by the observation that as one traverses across the bottom of a crater, the view of the starfield changes. The view is obscured by the walls of the crater and as one moves, one wall occludes more of the starfield and the opposing wall reveals more of the starfield, demonstrating the direct relationship between the different views and the locations in the crater. By matching a view taken at a position to map-computed views of the starfield, the method can estimate location.

Similar vision-based methods use the information provided by the skyline, the boundary between the viewed sky and the terrain. However, the skyline is not directly visible in dark crater environments and must be inferred from the occlusion of the starfield, degrading the accuracy of these methods. Alternatively, instead of using the skyline, celestial navigation techniques can be used which employ the information in the starfield. This work will demonstrate that position estimates computed from the inferred skyline alone or from the starfield alone are generally far less accurate than the proposed method, which fuses the two methods to take advantage of both information sources.

While many localization methods tend to be grid based, this method utilizes a significantly faster pseudo-gradient descent method to search the position space. This is made possible by the simplification of the map, a simplification that is supported by general lunar crater morphology.

This work presents an analysis of the method and highlights three significant factors that affect the performance of the localization: the ability to detect the starfield, the geometry of the crater, and the relative position in the crater. It is shown that for most lunar craters, this method should perform well. This method was successfully tested in a simulation environment modelling the mission target, Shackleton Crater, with localization accuracy better than 35 meters.

Acknowledgements

This work would not be possible without the significant amount of support that I've been blessed with and I'd like to express my deepest thanks and gratitude:

- First to my advisor, Red Whittaker, for your words of wisdom and for allowing me the freedom to pursue my interests.
- Also to my other committee members, David Wettergreen and Greg Barlow, who provided me valuable insight into the presentation of my work.
- To Ross Diankov, who survived MRD with me and helped develop the initial concept behind this paper.
- To the Scarab lunar rover design team: Paul, Deb, Jim, and Phil, who taught me so much about designing robots in a lunar environment.
- Then of course to Andrew, my officemate in year two, and Pras, my cubemate in year one, for allowing me to bounce my ideas off them. For putting up with my shenanigans and indeed participating in and instigating them. For foosball, Bomberman, Hitman, Sumotori Dreams, and the Create pull. For SIGBOVIK papers, late-night exercise ball fights, Segway runs, and car repairs. And for the snowmen, who will rule us all.
- To the RI incoming class of 2006 – it's been fun.
- Finally, to my mother, father, and sister, who have always supported me in my journey. Thanks.

For the Mars Exploration Rovers Spirit and Opportunity,
who have shown us how to perform beyond our design limitations

and, indeed, to all the intrepid explorers of space,
human, animal, and mechanical, who have inspired in us
le reve̋ d'étoiles – *the dream of stars*

Table of Contents

List of Figures

Chapter 1

Introduction

This paper introduces a novel method for estimating the pose of a rover within dark lunar craters, including the orientation and position relative to a map, without the use of artificial references such as beacons or some other positioning system. This is valuable for planetary rovers, which generally do not have access to extensive localization infrastructure such as the satellite-based and radio-based positioning systems which are available on Earth. Even with its close proximity, Earth-based GPS systems are not accessible on the Moon.

The thesis put forth is that the starfield and the occlusion of the starfield can be used to determine an estimate of the pose of a rover within a crater, given a map of the crater. The starfield is imaged using a wide-angle camera or star tracker with a 180 degree field of view. Because the rover is inside a crater, the wall of the crater will occlude the camera's view of the starfield such that unique sets of stars are seen in different locations of the crater. By matching these sets to computed sets, a location estimate can be determined. Additionally, by identifying the stars in the image with the use of a star catalog, the attitude of the rover can be determined very accurately. The accuracy of this method is controlled by a number of factors, including the number of imaged stars, the geometry of the crater, and the relative position of vehicle in the crater.

The motivation for this concept is the proposed exploration of the craters at the poles of the moon in order to look for water ice [1]. Due to the small axial tilt of the moon, these craters are permanently dark, creating cold traps which could retain ice deposited by past cometary impacts. This ice, if it exists, would be valuable for future exploration and settlement of the moon. Previous orbital surveys have returned data that indicate the presence of hydrogen at the pole, which may be in the form of water ice. Clementine used bistatic radar in 1994 to map the craters at the poles. The data from that experiment showed reflection characteristics in the permanently dark portions of the craters that suggested the presence of ice deposits. NASA's Lunar Prospector, launched in 1998, mapped the presence of hydrogen, a strong indicator of the existence of water, at both poles with the use of a neutron spectrometer. However, based on data from a later Arecibo radio telescope study, critics have suggested that the reflectance characteristics captured by Clementine are simply due to surface roughness, not ice [2]. In addition, during the controlled crash of Lunar Prospector into a polar crater at the end of its mission, water was not detected in

Figure 1.1: Image of Shackleton Crater, captured by the ESA SMART-1 mission

the plume caused by the impact. However, the estimate of the likelihood that water would be detected by this experiment, even if ice was present on the moon, was below 10% [3].

In order to directly confirm the presence of water ice on the moon, missions have been proposed and studied to send a rover like the Scarab prototype [4] into one of these polar craters and sample the crater floor in a grid with spacing on the order of a few kilometers. A prime candidate for such a mission is Shackleton Crater, shown in Figure 1.1, located at the south pole. Of course, while a simple positive or negative indication of the presence of water ice is valuable, knowledge of the location of these sample sites is critical in order to evaluate the general spread of the ice deposits and locate the ice in the future.

However, within these permanently dark craters there are a number of challenges that make localization and navigation difficult. There is no existing navigation infrastructure on the moon, such as the global positioning system (GPS) satellite constellation and the radio navigation beacon systems available on Earth. Measuring the local magnetic field for use as a heading reference is not possible due to the extremely weak field radiated by the moon. As the craters are completely dark, conventional vision-based navigation methods which depend on visible terrain features cannot be used. In addition, the much lower gravity ($1/6^{\text{th}}$ Earth normal) on the moon means that inertial navigation systems, both dead-reckoning and stellar-aided, are much less accurate than they are on Earth as accelerometers are less accurate in estimating tilt relative to the gravitational vector in low gravity environments.

The need, then, is to develop a method of localization that does not depend on these sources of navigation information.

Chapter 2

Background

The localization problem discussed in this thesis is called the "drop-off" problem [5, 6], where the rover must determine its location within a map without direct knowledge of its initial location. This scenario is used as it is very common for planetary rovers and conforms with crater mission profiles in which a rover lands within a crater with only extremely coarse, if any, location information based on its landing trajectory. As discussed in the introduction, a localization method is needed which does not require artificial, magnetic, or visual references from the terrain as these are not available in polar crater environments. Active sensing modalities such as long range radar or lidar combined with a particle filter certainly are functional possibilities, however these methods require significant amounts of power, which is severely limited on rover missions of this kind, which cannot operate off solar power and must carry a power source such as a radioisotope thermal generator (RTG) or fuel cells. Pure inertial navigation, whereby measurement of linear and angular acceleration is used to determine changes in position, does not require external references of any kind to propagate an estimate forward, however, it requires an initial position fix which cannot be provided. In addition, without future position fixes, the estimate is likely to drift significantly as measurement errors accumulate rapidly over time. Clearly, there is a need to find some external reference source to solve this problem.

2.1 Stellar-aided Navigation

One accessible external reference is the starfield. The starfield provides an excellent heading reference, and with knowledge of the movements of the stars and a good clock, the attitude of the vehicle can be determined with accuracy on the order of arcseconds. By adding a measurement of the local gravity vector (which provides a reference of the local vertical), position in two dimensions (latitude/longitude) can be determined as well. The "goodness" of the clock depends on the rotation speed of the planetary body the vehicle is situated upon – on Earth this is on the order of one second or better accuracy, on the moon this requirement is significantly less stringent. These measurements are the basis of celestial navigation, which has long been used aboard ships on Earth. Traditionally, sailors measured the elevation of stars above the horizon with the use of a sextant, combining the estimation of the local vertical into the same measurement. For land-based navigation, where the horizon is not visible, the local vertical is estimated by means of a plumb bob, or electronically by an accelerometer setup. Modern celestial navigation techniques eschew hand measurement of star elevations, instead imaging them with a camera for more accurate results.

The lunar environment makes imaging the starfield relatively easy compared to Earth-based observation. The absence of an an atmosphere means that there is no distortion or attenuation of the starfield. In addition, the slow rotation rate of the moon about its axis (sidereal rotation of 27.3 days or 0.549 degrees per hour compared to Earth's 15.041 degrees per hour) allows for longer sensor exposures, resulting in the detection of fainter stars. However, the lower gravity of the moon, at $1/6^{\text{th}}$ Earth normal, reduces the ability of accelerometers to measure the local vertical accurately. Sigel [7] has shown that celestial navigation on the moon with contemporary hardware will have a theoretical worst case error of 1.2 km, with the most significant contributing factor being the accuracy of the measurement of the local vertical.

Extending from the basic celestial navigation system, Malay, et al. [8] have analyzed a stellar-aided inertial navigation system for lunar exploration. In this system, an extended Kalman filter is used to combine celestial navigation estimates with measurements from a high quality inertial measurement unit (IMU) to produce a better position estimate. In simulation, they computed an accuracy of 400 m for a stationary vehicle located at the equator of the moon.

2.2 Skyline-based Navigation

Polar crater missions can take advantage of the unique location within a crater. The walls of the crater occlude the sky in a measurable way and this information can be used to compute a position estimate. A number of scholars [9, 10, 11, 12] have developed visual localization approaches for Earth-based scenarios that utilize the information contained in the skyline. The skyline is the boundary where, from the observer's perspective, the sky meets the terrain. The localization is done by matching the viewed skyline to skylines generated from a topographic map. This is possible as long as the terrain has elevation changes that are captured in the skyline, as shown in the Apollo 17 panorama in Figure 2.1, taken in the mountainous Taurus-Littrow region. If the terrain is flat and featureless, as in the Apollo 11 panorama in Figure 2.2 which was captured on Mare Tranquilitatis, there is very little information in the skyline and matching to a map yields no practical information.

If there is sufficient information in the skyline, the next issue is the size of the search space. If the map is continuous or very high resolution, then it is not possible to compute the skyline at every location, let alone compare the resulting skylines to viewed skyline. Many algorithms approach this problem by matching the skyline data to a discrete grid of locations in the map in order to reduce the size of search space. Even with a grid the space is still fairly large – for a 20 km x 20 km map, 10 m spacing results in four million grid points. Generating skylines for each of these positions typically requires a significant amount of processing.

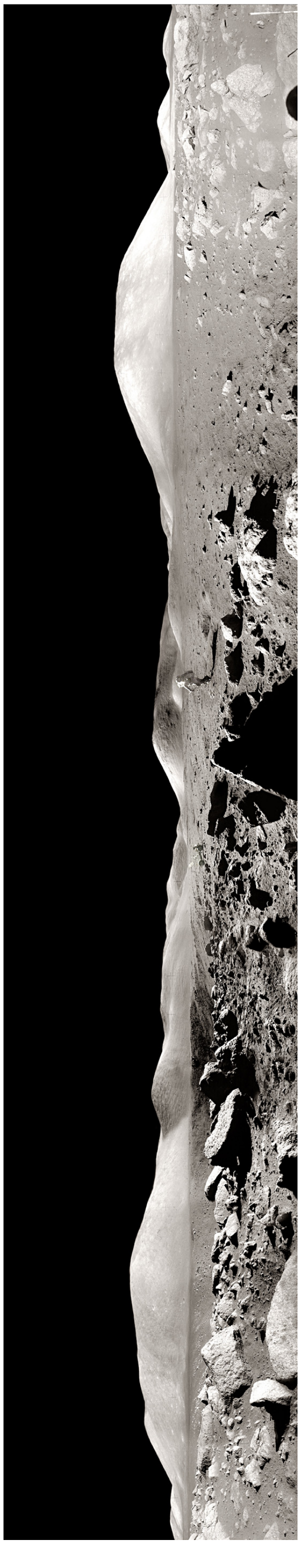

Figure 2.1: Lunar panorama captured by the Apollo 17 mission – note the multiple large curves in the skyline

Figure 2.2: Lunar panorama captured by the Apollo 11 mission – note the very flat, featureless skyline

Talluri and Aggarwal [9] developed a two stage approach whereby four points on the skyline, sampled at the four cardinal directions of the compass (North, South, East, and West), were matched to four points estimated from a grid of positions within the topographic map in order to reduce the search space. Full skylines are compared for this much smaller search space, with a curve matching algorithm to determine the best match, which is the final estimate. In simulation, they achieved a mean accuracy of 26 meters with uncorrupted images of the skyline and 55 meters when using images with zero-mean Gaussian noise added. These are excellent results, however, some have argued [11] that the four point method is too simplistic and will cause problems in more general terrain.

Stein and Medioni [10] approach this problem by subsampling both the skyline and the map. The skyline is simplified by fitting line segments to it and the map is subsampled to a 75m grid. In addition, they speed up computation during the localization process by precomputing the skyline estimates at every point in the grid, trading online processing time for memory storage space. With their method, they achieved accuracies between 300 and 600 m on real imagery.

Cozman [11] utilized a similar method to Stein and Medioni, except that he uses the full skyline rather than the fitted line segment representation. In addition, he represents the problem in a Quasi-Bayesian formulation, computing both a position estimate and a measure of estimation quality. His experiments in a variety of environments, including urban, desert, and lunar skylines achieved accuracies on the order of hundred of meters, which scaled with the resolution of the map.

Naval et al. [12] uses a completely different approach than the above. Their method is feature-based and matches features between the image and the map to generate a pose estimate. In addition, their method does not require an orientation estimate – it is determined in the process. While reasonable results were achieved (mean accuracy of 373 m), the correspondence search required with this method consumed an average of approximately 90 minutes of processing time for a single location.

The author previously [13] used a skyline based approach in a scale model crater environment. This method attempted to combine vehicle odometry with skyline information in order to both estimate the map and localize the vehicle. In this work, it was assumed that the environment was dark, as in a lunar polar crater, with only the starfield visible. Under this scenario, the skyline is not directly visible and must be inferred by the occlusion of the starfield. The skyline was estimated to be located along the outer most stars in the starfield. Experimentally, this method worked well in the simplified case of travelling in a straight line and estimating only the crater depth and rover position, both of which are estimated *a posteriori*. The results, if scaled to a crater with the same diameter as Shackleton (20 km), would have a mean position error of 800 m and a depth estimation error of 110 m. In the more general 3D scenario, inaccuracies in the heading estimate and the fact that the stars do not lie precisely on the skyline prevented good results from being computed from the experimental data.

This thesis will show that by using a combination of techniques gleaned from both stellar-aided navigation and skyline-based navigation, a system can produce quite good localization results which cannot be matched by either technique alone for lunar polar crater missions.

Chapter 3

Crater Localization

In a lunar polar crater, there are a number of environmental factors that change the nature of the localization problem and prevent the direct use of the techniques discussed in the previous chapter. The constant availability of the starfield in this environment, due to the lack of atmosphere and the polar location, is very convenient and would suggest a celestial solution. However, as discussed in Section 2.1, that while the stellar-aided inertial navigation method provides a reasonable solution, 400 m accuracy is not sufficient for 1 km scale traverses. That aside, the ability of the starfield to provide highly accurate heading is very useful.

The crater topography provides walls that give good location information in the skyline. With the launch of the Lunar Reconnaissance Orbiter in October 2008, maps of the lunar surface with surface resolution of 35 m and depth resolution of 0.1 m will soon be available thanks to the Lunar Orbiter Laser Altimeter (LOLA) instrument onboard. This will allow for the computation of estimated skylines, as in the skyline-based techniques discussed in the previous chapter. However, due to the permanently dark nature of the polar craters, the skyline is not directly observable. It can only be inferred by means of starfield occlusion. It will shown that this inferred skyline is insufficient to directly provide a good position estimate.

Crater morphology studies [17] tell us that the crater floors and rims are relatively smooth and flat on a large scale. In addition, the shape of a crater's plan view is generally circular in nature. The method will take advantage of this by simplifying the terrain model as a squat, open-top cylinder, as shown in Figure 3.1. The floor is, of course, not flat near the edges of the crater, which is generally more bowl-like. However, this approximation is reasonable when not attempting to scale or descend the steep crater walls. For comparison, a general crater model from the Lunar Sourcebook is shown in Figure 3.2. It should be noted that the height of the central peak is not very tall, generally 1-3% of the depth of the crater.

As a vehicle moves within the crater, it will move away from one side of the crater and towards the opposite side. As this occurs, the view of the starfield will change as the starfield occlusion by the walls changes. Examples of starfield views at different locations in a crater are shown in Figure 3.3. There is a direct, 1:1 mapping between sets of visible

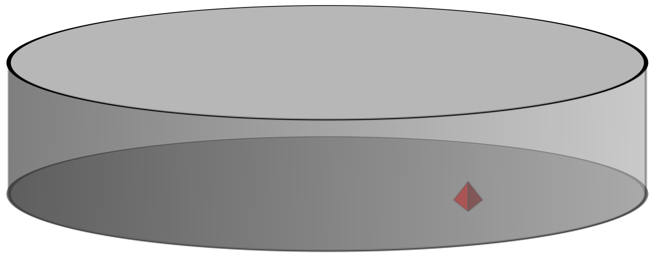

Figure 3.1: The simplified crater model

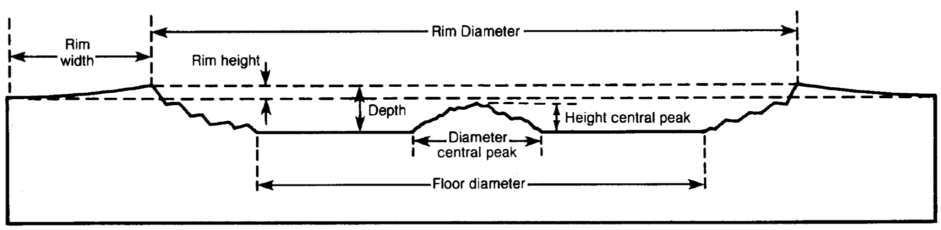

Figure 3.2: The general lunar crater model

stars and regions of the crater, and hence the position in the crater can be determined by matching the visible stars against the stars that are computed to be viewable at a location in the crater.

Because the stars have been studied in great detail by astronomers, star catalogs are available which characterize their emission spectra and movements very accurately. With this information, stars can be matched and identified in an image perfectly using standard star tracker algorithms which use pattern matching for this process [18]. Since the stars can be perfectly identified in every starfield, image-based methods such as iterated closest point (ICP) need not be used to compare starfields – instead, lists of stars can be compared, which greatly accelerates the process. By matching to a star catalog, the measured values of azimuth and elevation can be corrected to the known values, eliminating quantization and random noise in the measurements.

3.1 Method

This method requires a camera capable of imaging the entire starfield, such as a fisheye camera with a 180 degree field of view. This camera is pointed upwards at the sky. Using this image, a two step process is used to estimate the position of the vehicle. Initially, a rough position estimate is computed by using an estimated skyline generated as the convex hull of the visible stars. Because the map is circular and of known diameter and depth, the position can be directly computed. This initial estimate is used to seed the main localization algorithm, which matches the starfield generated at the estimated position to the viewed starfield. If they do not match, the starfields are projected on a 2-D plane and

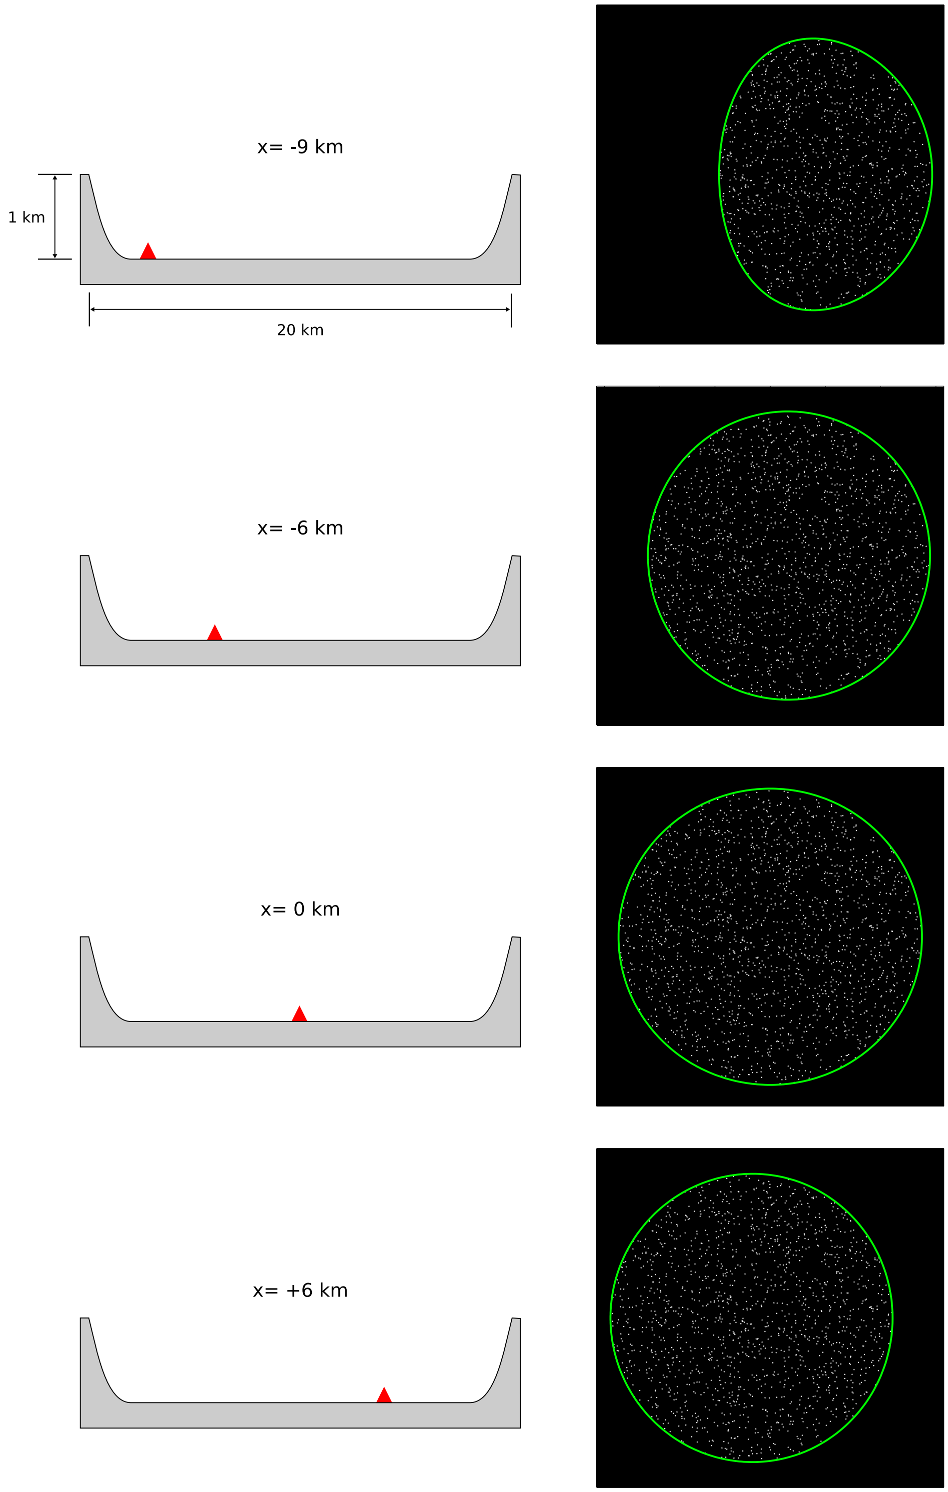

Figure 3.3: Viewed starfields at different positions along the centerline of the crater. The green line indicates the position of the edge.

their centers of mass computed. The resulting vector from one center to the other is used as a pseudo-gradient to determine the the direction and distance the position estimate should be changed, as in a gradient descent method. This is repeated until the starfields match or the algorithm is otherwise terminated.

3.1.1 Initial Estimate

In order to determine a rough position estimate, a computation is made to determine which stars form the convex hull of the viewed starfield. From this, the position within the crater can be computed, using the assumption that the stars are located on the crater edge and form the skyline. The distance along the floor to the crater wall is computed from the trigonometric relationship between the star elevation angle and the known crater depth:

$$d = \frac{\text{crater depth}}{\tan\theta}, \tag{3.1}$$

where θ is the elevation to the star. This generates a set of polar coordinates of points that form a circle. Converting to Cartesian coordinates, the standard form of a circle is employed:

$$(x-h)^2 + (y-k)^2 = r^2, \tag{3.2}$$

to compute a least squares fit which is setup as follows:

$$Aq = B \tag{3.3}$$

$$\begin{bmatrix} 1 & -2x_1 & 1 & -2y_1 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & -2x_n & 1 & -2y_n \end{bmatrix} \begin{bmatrix} h^2 \\ h \\ k^2 \\ k \end{bmatrix} = \begin{bmatrix} r^2 - x_1^2 - y_1^2 \\ \vdots \\ r^2 - x_n^2 - y_n^2 \end{bmatrix}. \tag{3.4}$$

r is the known radius of the crater, (x_n, y_n) are the points on the circle in Cartesian form, and (h, k) is the coordinate of the center of the circle relative to the current position. By taking the pseudoinverse of A, q can be solved as:

$$A^+B = q. \tag{3.5}$$

By negating the coordinate of the circle center:

$$\text{position} = (-h, -k), \tag{3.6}$$

an estimate is obtained of the position relative to the center of the crater. As the stars are generally not precisely on the crater rim and are some small elevation above the skyline, this is only an estimate. In a test of this initial estimate at 100,000 uniformly random sampled points throughout a crater under typical conditions (10:1 aspect ratio crater, magnitude 6 star detection limit), the mean accuracy was 637 m with a standard deviation of 275 m and a maximum value of 1.735 km. A plot of the position error at each sampled location is shown in Figure 3.4. From the plot it is observed that the method appears to do better near the center and worse near the edges of the crater. This is because as the rover travels

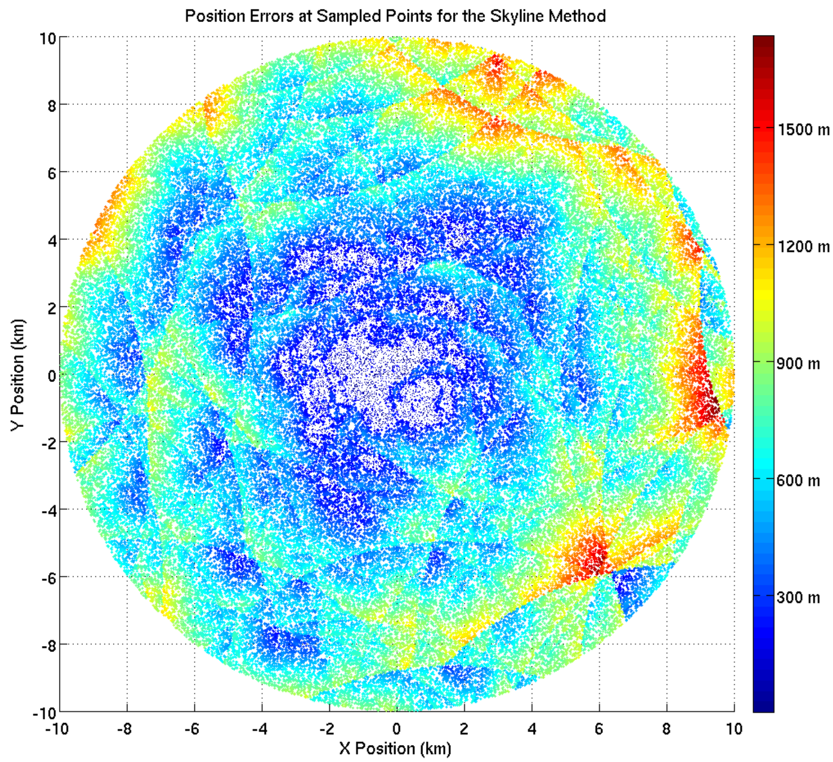

Figure 3.4: Position error at sampled points for the skyline estimate – Point size and color scaled to position error

away from a point on the crater edge, the error in distance to that point induced from the angular error in the skyline estimate increases exponentially, as shown in our previous work [13]. The combination of these errors for each point on the inferred skyline results in poorer performance near the crater walls as the errors induced by the points on the distant wall far outweigh the increase in accuracy from the points along the near wall.

This algorithm depends on knowing the absolute heading of the vehicle, which is assumed to be obtained from the starfield.

3.1.2 Estimate Refinement

To refine the initial estimate, a pseudo-gradient descent method is used. Pseudocode for this method is shown in Algorithm 1. The pseudo-gradient is computed by first generating a view of the starfield at the estimated position. Projecting the stars onto a plane, the center of mass is computed for both the computed view and the actual view. The vector from the center of mass of the estimated view to the actual view is used for the pseudo-gradient and is used to update the current estimate to determine the new estimate. An example of this is shown in Figure 3.5. The measure of goodness of fit is computed by comparing the generated view against the actual view. The number of stars which do not match are counted and used as the goodness of fit measure. The algorithm terminates when the starfields match. The advantage of this method is that by using the pseudo-gradient to guide the search, significantly fewer evaluations are made compared to grid-based methods.

Algorithm 1 Localization algorithm

```
function LOCALIZE(estPosition, obStars, starfield, map)
    obStarsCartesian = azElTo2DCartesian(obStars)          ▷ Project stars onto a plane
    obStarsCOM = mean(obStarsCartesian)                     ▷ Precompute center of mass

    i = 0
    repeat
        estCraterEdge = genObservedCraterEdge(estPosition, starfield, map)
        estStars = genObservedStars(starfield, estCraterEdge)      ▷ Estimate starfield
        starError = count(obStars ≠ estStars)                      ▷ Compare starfields

        estStarsCartesian = azElTo2DCartesian(estStars)
        estStarsCOM = mean(estStarsCartesian)
        vector = estStarsCOM − obStarsCOM                   ▷ Difference in centers of mass

        estPosition = estPosition + vector                  ▷ Position estimate update

        i = i + 1
    until starError == 0 or i ≥ MAX_ITERATIONS

    return position
end function
```

Due to the discontinuous nature of the problem, the algorithm does not always terminate. The reason for this is that the pseudo-gradient is determined by the starfield, which is discrete in nature. When the difference between the estimated and actual starfields is large, the pseudo-gradient is a good estimate of the gradient. However, when the difference is small, on the order of a few stars, the pseudo-gradient is a poor estimate. This is because a single star difference on the edge of the observation does not provide much information – the correct direction can be as much as ± 90 degrees away from the azimuth of the star. This can cause the algorithm to wander and not terminate.

In the implemented version of this algorithm, a limit is placed on the number of evaluations to ensure termination. In practice, the algorithm was terminated by the evaluation limit in approximately 0.6% of 100,000 randomly sampled locations under standard conditions (10:1 aspect ratio crater, magnitude 6 star detection limit). A check of those locations that were terminated by the evaluation limit showed that their mean error was not significantly different from the overall mean.

Another way of solving this problem would be to evaluate a grid of points around the estimate when the difference in the starfields becomes small and selecting the best match.

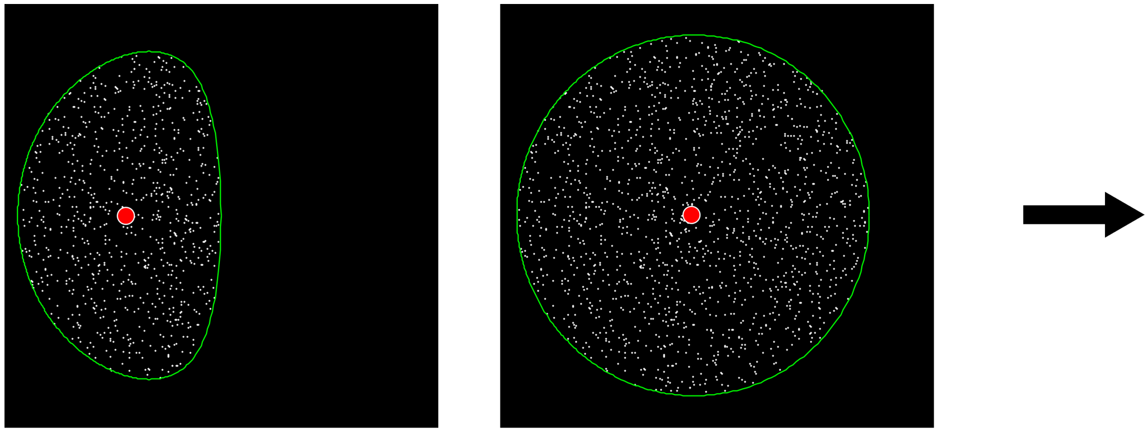

Figure 3.5: Example showing pseudo-gradient computation. The red dots indicate the centers of mass of the two views. The vector created from those points form the pseudo-gradient.

3.2 Analysis

In this section, the simulator used to test this method is introduced. In addition, analysis is presented which examines the various parameters that affect the performance of this localization method. These parameters include the ability of the camera to detect the starfield, the geometry of the crater, and the relative position within the crater. Finally, some of the limitations of this method are discussed.

3.2.1 Crater Environment Simulator

The crater environment simulator, which was developed in MATLAB, computes the visible stars in a given location. This simulator first generates the entire starfield across a 180 degree field of view, which is stored as a pair of elevation and azimuth values for each star. The algorithm for the starfield generation is discussed below in Section 3.2.2. The stars are considered to be at an infinite distance away and the crater is considered to be small enough relative to the size of the moon to be locally flat. Using these assumptions the stars always appear at the same azimuth and elevation at any location within the crater. These are reasonable assumptions – as the stars are extremely far away the parallax angles are very small. Even the closest star, Proxima Centauri, at 4.2 lightyears away, only has a parallax angle of 0.772 arcseconds across the two extremes of the moon's orbit around the sun. For the locally flat assumption, across a 20 km lunar crater an observer would experience a rotation of 0.66 degrees in the starfield. This is small enough that an initial position estimate would not be significantly affected and would yield a result which would allow the estimate refinement to compensate for the small amount of rotation. In addition, the rotation rate of the moon itself is neglected in this simulation. In a fielded system, the rotation rate of the moon is slow enough that it could easily be compensated for with

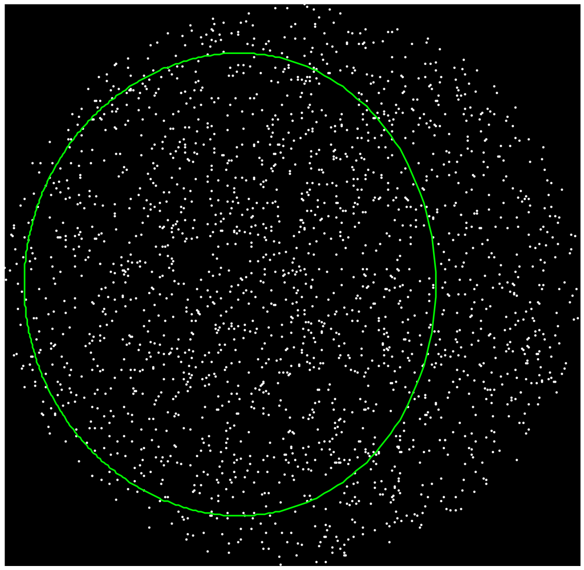

Figure 3.6: Example showing viewed starfield computation. The green line is the crater edge projected onto the entire starfield. Stars within the crater edge are the viewed stars.

knowledge of the moon's rotation rate and an accurate source of time (with accuracy on the order of seconds). The crater itself is modelled as a simple cylinder, as described in the beginning of the chapter, and stored as a pair of diameter and depth values.

In order to generate the viewed starfield at a given position, the observed crater edge is first computed as it would appear if it was visible. This is done by evenly sampling 360 points along the crater rim and computing the elevation to those points from the given position. These crater points and the starfield are then projected onto a 2-D plane. The stars that lie within the polygon formed by the crater points are considered to be the stars viewable from the given position. An example of this projection is shown in Figure 3.6. While this was tested with a perfectly flat rim, this method should extend to situations where there are vertical changes in the rim, as long as the circularity of the crater is not affected.

3.2.2 Starfield

This method depends on the ability of the camera to detect changes in the viewable starfield, which is a function of the sensitivity of the camera to detect the starfield. In star trackers, this is given as a star detection threshold, typically expressed in terms of star magnitude, a measure of the brightness of a star on a logarithmic scale. The scale is referenced to the brightness of the star Vega, at $M_v = 0$, with the formula for computing the difference in magnitude between two stars as:

$$m_1 - m_2 = -2.5 \log(f_1/f_2), \tag{3.7}$$

where m is the apparent magnitude of the star and f is the observed flux of the star [14]. The greater a star magnitude value, the dimmer the star. Polaris, the pole star, has a visual magnitude of approximately 2, while a star of magnitude 6 is near the limit of viewability by the naked eye.

Instead of using actual starfields, this analysis employs randomly generated starfields in order to prevent any particular features of the starfields from biasing the results. Liebe [15] presents a formula for estimating the average number of stars visible given the magnitude limit of the sensor, computed by fitting to a star catalog:

$$N_{FOV} = 6.57e^{1.08M} \frac{1 - \cos\left(\frac{A}{2}\right)}{2}, \tag{3.8}$$

where N_{FOV} is the number of stars in the field of view, M is the magnitude limit, and A is the field of view. For our method we assume a fisheye lens with a 180 degree field of view. With a magnitude limit of 6, this equation gives us 2,142 visible stars in the sky. This is, of course, only an estimate – the stars are generally not evenly distributed in the sky. There are more stars viewable in the galactic plane than out of it. However, for the sake of our analysis, the number of stars given by this equation are uniformly distributed over the unit hemisphere. The distribution is performed by sampling two random numbers, r_1 and r_2, uniformly over $[0, 1]$ and computing the azimuth, ϕ, and elevation, θ, as follows [16]:

$$\phi = 2\pi r_1 \tag{3.9}$$

$$\theta = \sin^{-1}(r_2). \tag{3.10}$$

Sample plots of the generated starfield are shown in Appendix A in Figure A.1.

What this analysis determines is how the camera detection performance affects the ability of this method to localize well. Using our simulator, the distance a rover must travel before a change is seen in the starfield can be determined. This is done by computing the visible starfield at a number of locations along a centerline of a crater. This is a good metric as the key information used in this localization algorithm is the starfield, and the distance between starfield changes is, in a sense, the resolution limit for this information source, as there is no new information until the viewed starfield changes. This metric should directly correlate to the efficacy of the localization method. For this analysis, both average and maximum distances travelled between starfield changes were determined for different star detection limits. A Shackleton-type crater with a 10:1 aspect ratio was assumed and the star detection limit was varied from magnitude 2 (Polaris) to magnitude 7 (below the viewing threshold of the human eye). The results of this are shown in Figure 3.7. From that plot, it is observed that there is a definite correlation between the number of stars viewed and the information in the starfield. As one would expect, the more stars that are detectable, the more information there is in the starfield occlusion.

A by-product of this analysis shows that the distance between changes in the viewed starfield increases as the vehicle approaches the center of the crater, as shown in a plot of the distances between starfield changes as the vehicle travels across the crater in Figure 3.8. This is because at the center of the crater, the change in occlusion from the starfield walls is at a minimum.

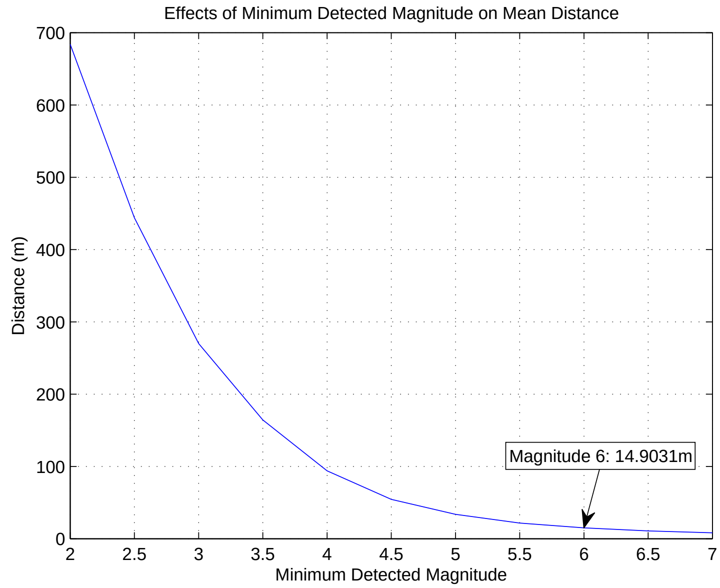

(a) Mean distance between changes in the viewed starfield

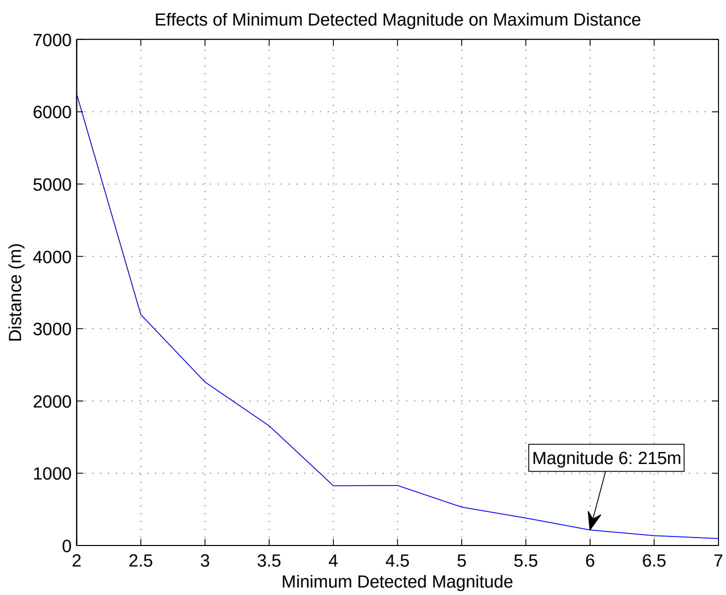

(b) Maximum distance between changes in the viewed starfield

Figure 3.7: Effects of changing the star detection limit

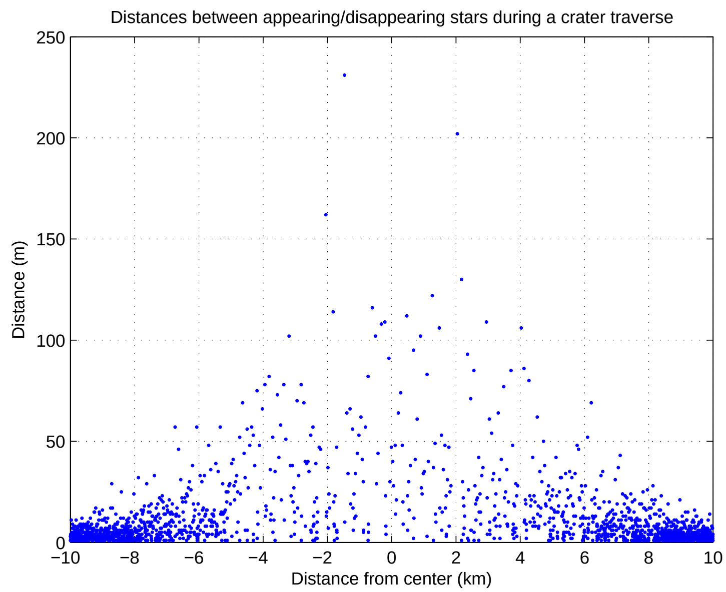

Figure 3.8: Plot of the distance between starfield changes along a crater traverse

3.2.3 Crater Aspect Ratio

Another important environmental variable for this method is the crater aspect ratio, the ratio between the diameter and the depth of the crater. Running the simulator for a range of aspect ratios, the result shown in Figure 3.9 is obtained. This simulation operated with a star detection limit of magnitude 6. As one might expect, the greater the aspect ratio, corresponding to a shallower crater, the greater the degradation in performance. At this end of the plot, the effect of the much smaller crater walls occluding the starfield is significantly reduced and approaches a flat plain. However, at the other end of the plot it is noted that decreasing the aspect ratio below 1:1 causes the performance to again suffer. The reason for this performance dropoff is that as the crater becomes deeper, the view of the starfield becomes smaller. Taken to the limit, this corresponds to standing at the bottom of an infinitely deep well, only seeing stars that are directly above. Fortunately for the purposes of lunar exploration, a significant number of craters have aspect ratios in the operable region. The motivating case of Shackleton Crater, with an aspect ratio of 10:1, falls well within the region. For lunar craters, the Lunar Sourcebook [17] uses the formula:

$$d = 0.196D^{1.010}, \tag{3.11}$$

to describe the relationship between crater depth, d, and diameter, D for simple craters with diameters less than 15 km, and the formula:

$$d = 1.044D^{0.301} \tag{3.12}$$

for complex craters with diameters between 12 and 275 km. The formula for simple craterforms gives us aspect ratios of approximately 5:1 and the formula for complex craterforms gives us a range of aspect ratios from approximately 5:1 to 50:1. Plots of diameter to aspect ratio are shown in Appendix A in Figure A.2. The two ends of the lunar crater aspect ratio range are marked in Figure 3.9, and as can be seen, a significant portion of the likely crater range is workable using this method, though the very widest craters near 50:1 begin to show maximum distances of approximately 1 km.

3.2.4 Position Dependence

Because this localization method is driven by the combination of the skyline information and the starfield, there is some dependence on position in the accuracy of the position estimate. As discussed in the Section 3.2.2 previously, the average distance travelled before seeing a change in the starfield is longer near the center of the crater. To show this dependency, 100,000 locations were uniform randomly sampled within a crater and the localization algorithm was run at those locations under standard conditions (10:1 aspect ratio crater, magnitude 6 star detection limit). Figure 3.10 shows the result, which clearly shows a degradation of performance near the center of the crater.

3.3 Shackleton Crater Results

For the case of the targeted crater, Shackleton, simulations were run for 10,000 uniform randomly sampled locations under ten different randomly generated starfields. The different starfields were used to ensure that particular features in a single starfield do not skew the results. A star detection limit of magnitude 6 was used. Shackleton crater was defined with the simplified crater model to have a diameter of 20 km and a depth of 2 km. In simulation, a mean localization error of 34.2 meters, with a standard deviation of 31.1 meters was achieved. The maximum position error was 358 meters. These results are comparable to the results achieved by Talluri and Aggarawal and are an order of magnitude better than the other stellar-aided and skyline-based methods. Table 3.1 shows the results for each of the ten runs. Plots of the position errors for each of the ten starfields are shown in Appendix B. On average, the algorithm required approximately 0.5 seconds per location estimate on a machine equipped with a 3.2 GHz Pentium D. It should be noted that this simulation uses preidentified stars. In the field, some processing will be required for starfield identification.

3.4 Limitations

As well as this method works in simulation, there are a few issues that could potentially affect performance in the field. One implementation detail is the camera mount. This method requires seeing the entire starfield and the skyline, which can be implemented with a 180 degree field of view camera in most cases as long as the vehicle does not pitch or roll. Even if the vehicle does pitch or roll, if the camera's field of view is wide enough and depending on the aspect ratio of the camera, it may still be able to capture the starfield. In a Shackleton-type crater, the aspect ratio is such that with a 180 degree FOV camera, the vehicle can pitch/roll a maximum of 5.7 degrees at the edge and a maximum of 11.3 in the

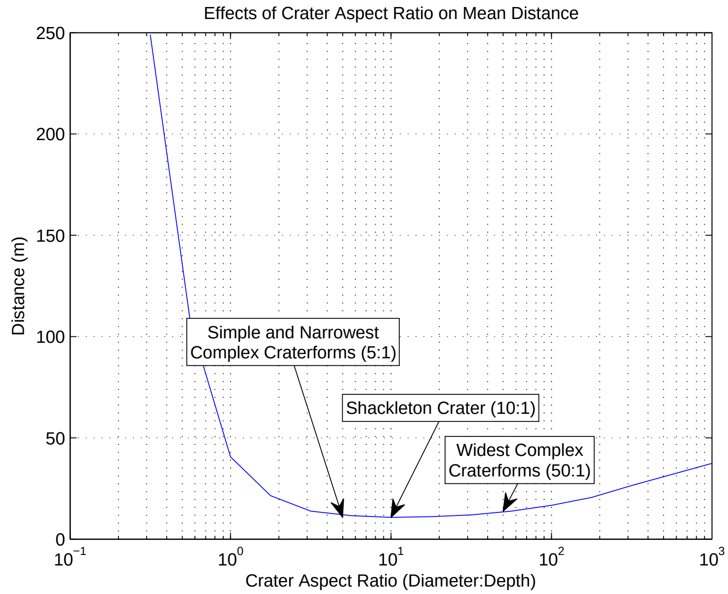

(a) Mean distance between changes in the viewed starfield

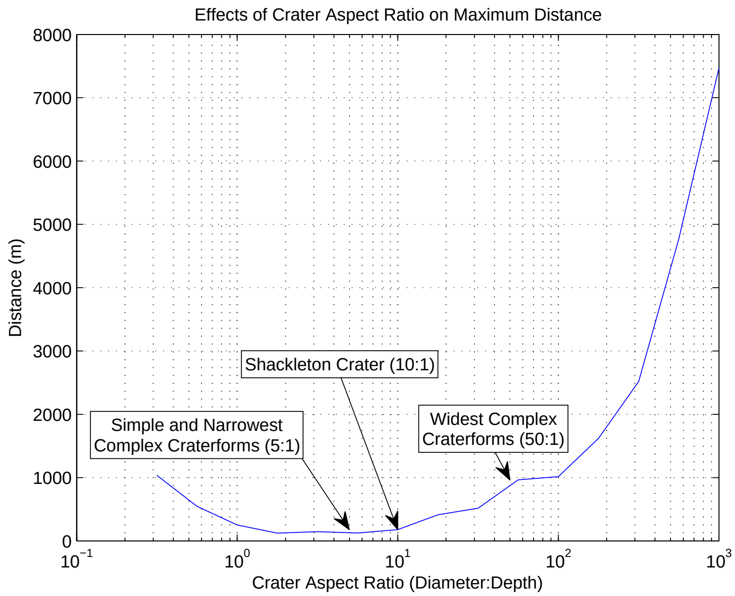

(b) Maximum distance between changes in the viewed starfield

Figure 3.9: Effects of changing the crater aspect ratio

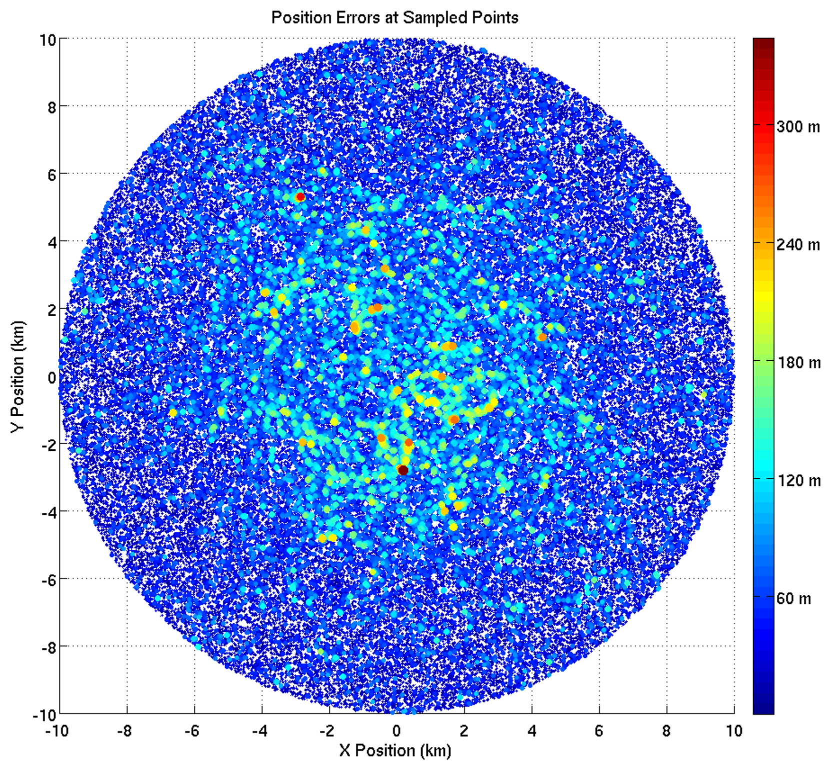

Figure 3.10: Position error at 100,000 sampled points – Point size and color scaled to position error

	Position Error (m)		
Trial	μ	σ	Max
1	33.3	29.8	317.9
2	33.7	29.9	327.3
3	34.6	31.7	344.3
4	33.6	30.2	271.2
5	34.3	30.2	242.9
6	34.1	32.3	300.7
7	36.2	32.9	325.6
8	33.9	31.2	300.2
9	34.1	31.2	358.4
10	34.4	31.5	356.6
Overall	**34.2**	**31.1**	**358.4**

Table 3.1: Results of ten simulation runs in a Shackleton-like crater

center of the crater while still seeing the entire starfield. Alternatively, the sensor can be gimbled to compensate, however this will require additional mechanical complexity which is expensive on space exploration vehicles.

One assumption used in the current method is that the unoccluded starfield is the same throughout the crater. For a crater the size of Shackleton, which has a diameter of 20 km, this is a reasonable assumption as the rotation in the starfield at opposite sides of the crater is a fairly small 0.66 degrees. However, in larger craters this could become more significant. This could be overcome by using the initial position estimate to make the proper compensation in the starfield before refining the estimate with the primary localization algorithm.

Another factor that is not considered here is the possibility of local obstacles which are large enough to occlude the starfield, but are not on the map. This can occur if the vehicle is small enough and/or the map is of significantly coarse resolution. The additional occlusion source is likely to bias the current localization algorithm away from the unexpected occlusion due to the algorithm's center of mass method of computing a pseudo-gradient. One way of correcting for this would be to focus on the stars along the skyline, attempting to match the skyline stars rather than the entire starfield. This would also help in the case of a rover tilted such that it is not capturing the entire starfield. Alternately, if the vehicle could scan the obstacle and obtain its shape, the occlusion caused by it could be directly compensated for.

Chapter 4

Conclusion

This thesis has shown a novel localization method for the unique environment that is found in permanently dark lunar craters. This method combines techniques from starfield-aided and skyline-based localization methods to achieve results which are not possible with either technique in isolation. In addition, the method takes advantage of the information available in the starfield as well as the occlusion of that starfield. It is shown that position estimates computed from the inferred skyline alone are generally far less accurate than the proposed method, which matches the visible stars to the stars which are computed to be visible from position estimates. A pseudo-gradient descent method is used to search the position space, which is much faster than slower grid-based methods. This is made possible by the simplification of the map to be cylindrical, a simplification that is supported by general lunar crater morphology. Importantly, the availability of a precision star catalog is used to match starfields and compensate for noise, as well as extract a heading estimate.

In the analysis of this method, it has been shown that there are three significant factors that affect the performance of this algorithm. The ability to detect the starfield is certainly a factor – as demonstrated, the sensitivity of the camera and the accuracy of our localization is directly related, with increasing camera sensitivity improving localization performance. In addition, the geometry of the crater, specifically the ratio between the crater diameter and the crater depth, was also shown to affect the localization performance. Instead of a direct relationship, there is a "sweet spot" aspect ratio – craters with aspect ratios away from the sweet spot had increasingly poor performance the further away it was from the sweet spot. It was shown that most lunar craters fall within the usable range on that curve. Finally, the observation is made that the localization method is sensitive to the relative position in the crater. The further away the vehicle is from the center of the crater, the better the localization accuracy with this method.

Finally, this method was implemented in a simulated crater environment, where results were obtained for the mission target, Shackleton Crater. These results showed an order of magnitude improvement over existing methods and allows for near-GPS quality positioning ($\sim$35 m) without additional infrastructure. This is a powerful tool for robotic sampling and survey missions in dark craters and should be considered for all missions of this type.

4.1 Future Work

While this localization method has been shown to work well in a simulated crater environment, there certainly is a significant amount of work that should be done to further prove and extend this algorithm. Significantly, this algorithm should be tested with a real crater floor map while still maintaining the simplified crater model to determine the applicability of the model. And while testing this algorithm on Earth is likely to be quite difficult due to the presence of an atmosphere and the significantly faster rotation rate, it may be possible to compensate computationally by analyzing the streaks caused by the movement and mechanically with a tracker mount.

One way of improving the localization performance would be to combine the algorithm with a inertial navigation system through a Kalman filter. In this way, the inertial navigation system could provide dead reckoning updates as the vehicle travels between changes in the viewed starfield, which should significantly improve the accuracy of the system.

Another way of improving the performance of the system would be to take advantage of the complimentary nature of the initial skyline estimate and the occluded starfield refinement, which is not done in the current work. The skyline estimate, while on average performs worse than the starfield estimate, performs better near the center of the crater – the opposite of the starfield estimate. The system could simply use the estimate returned by the initial skyline estimate if that estimate is determined to be near the center of the crater.

In addition, currently the method uses static images of the starfield to perform skyline inference and matching. It should be possible to extract better skyline information by taking multiple images of the starfield over time to take advantage of the apparent rotation of the starfield. This would also allow for multiple position estimates with the starfield matching algorithm, which could be aggregated to produce a better estimate.

Finally, while this work exploited the relatively simple topography of the crater environment to speed up the algorithm, it should be possible to use this method with a grid-based localization rather than the current pseudo-gradient search to estimate the position of the vehicle in more general darkened terrain and terrain that is partially lit. The starfield matching method should still be valid in such as setting given a decent map. This would allow the extension of the use of this method to general darkened terrain such as the night side of the moon, rather than just darkened lunar craters, however at the cost of performance.

Appendix A

Additional Figures

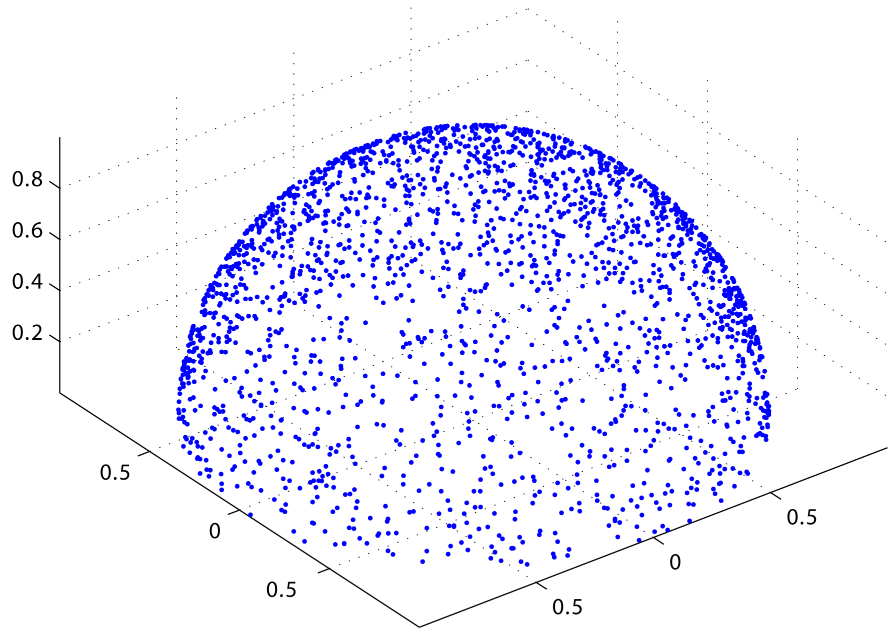

(a) Starfield projected on the unit hemisphere

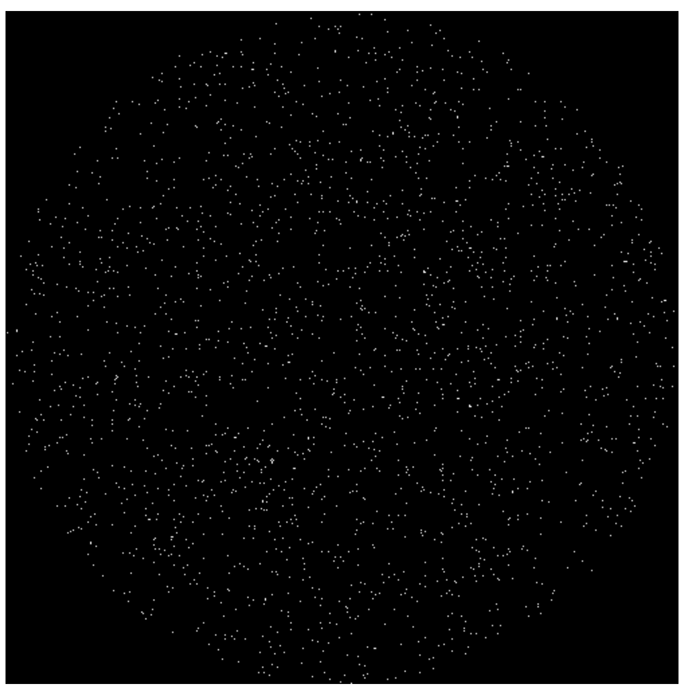

(b) Fisheye projection of the starfield

Figure A.1: Simulation generated starfield for a magnitude 6 detection limit

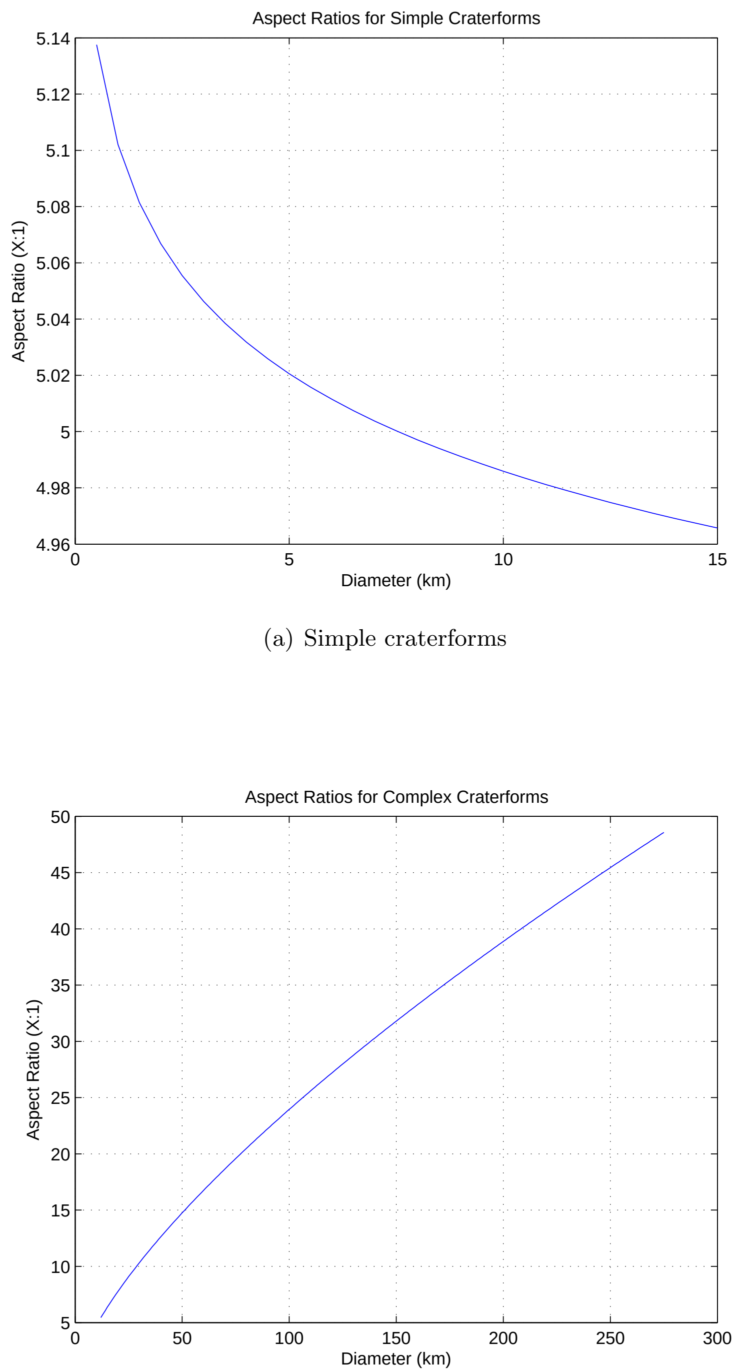

(a) Simple craterforms

(b) Complex craterforms

Figure A.2: Lunar crater aspect ratio ranges

Appendix B

Simulated Shackleton Crater Results

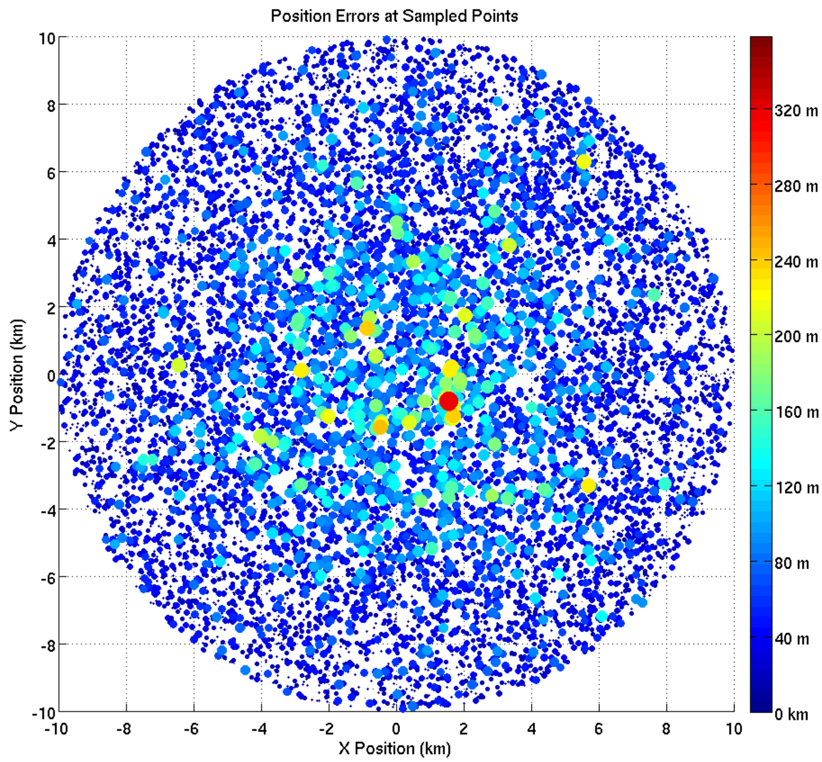

Figure B.1: Position error at 10,000 sampled points in a Shackleton-style crater, starfield 1 – Point size and color scaled to position error

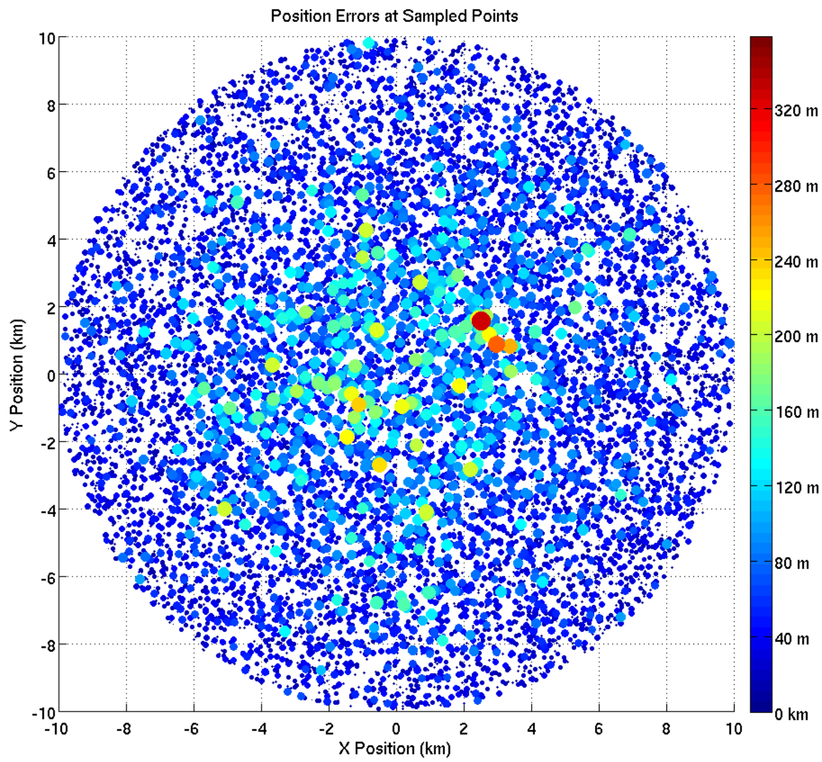

Figure B.2: Position error at 10,000 sampled points in a Shackleton-style crater, starfield 2 – Point size and color scaled to position error

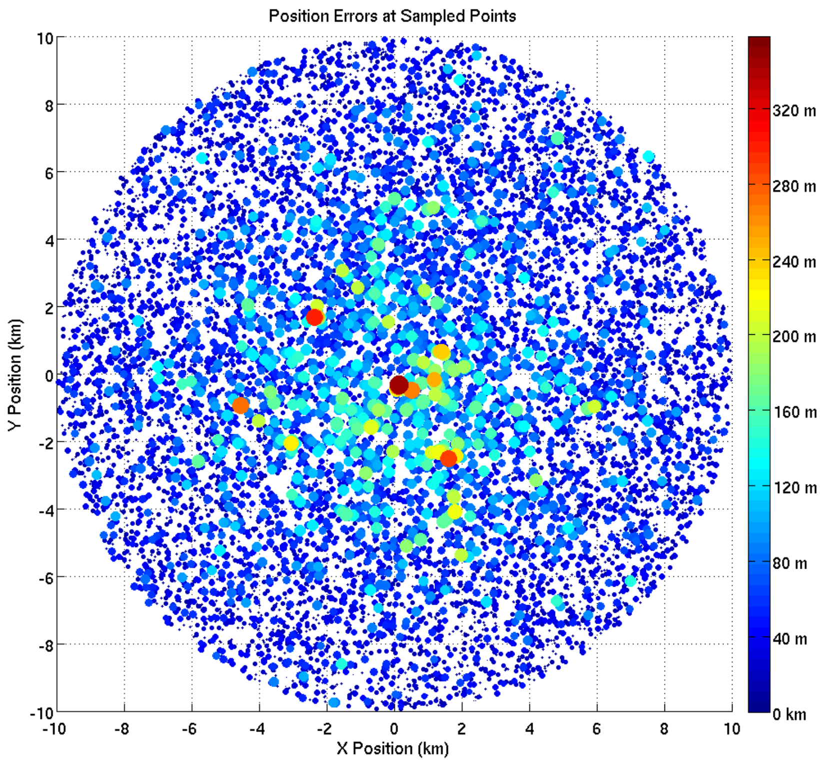

Figure B.3: Position error at 10,000 sampled points in a Shackleton-style crater, starfield 3 – Point size and color scaled to position error

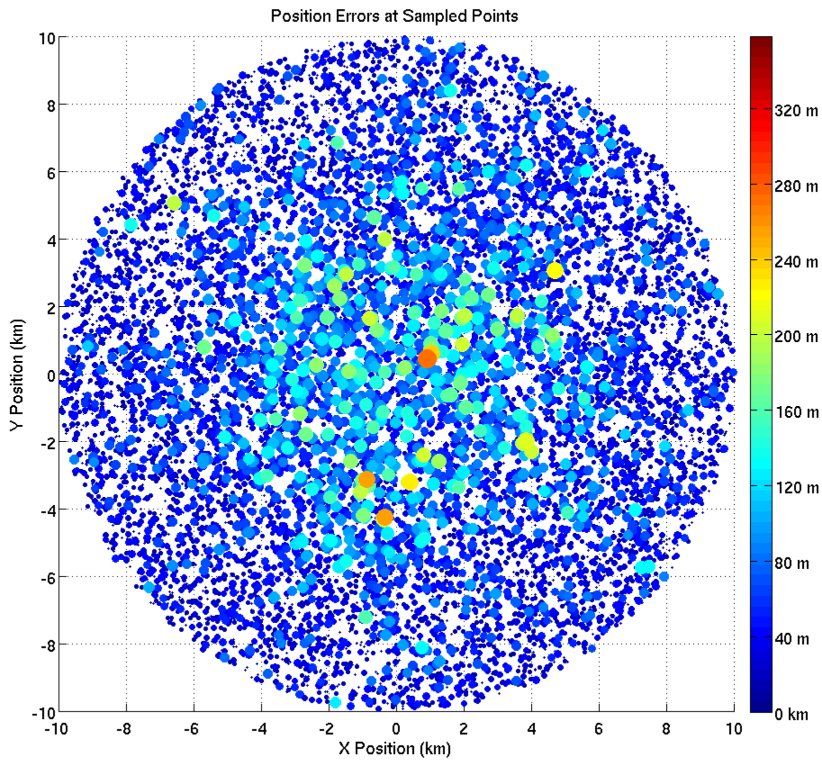

Figure B.4: Position error at 10,000 sampled points in a Shackleton-style crater, starfield 4 – Point size and color scaled to position error

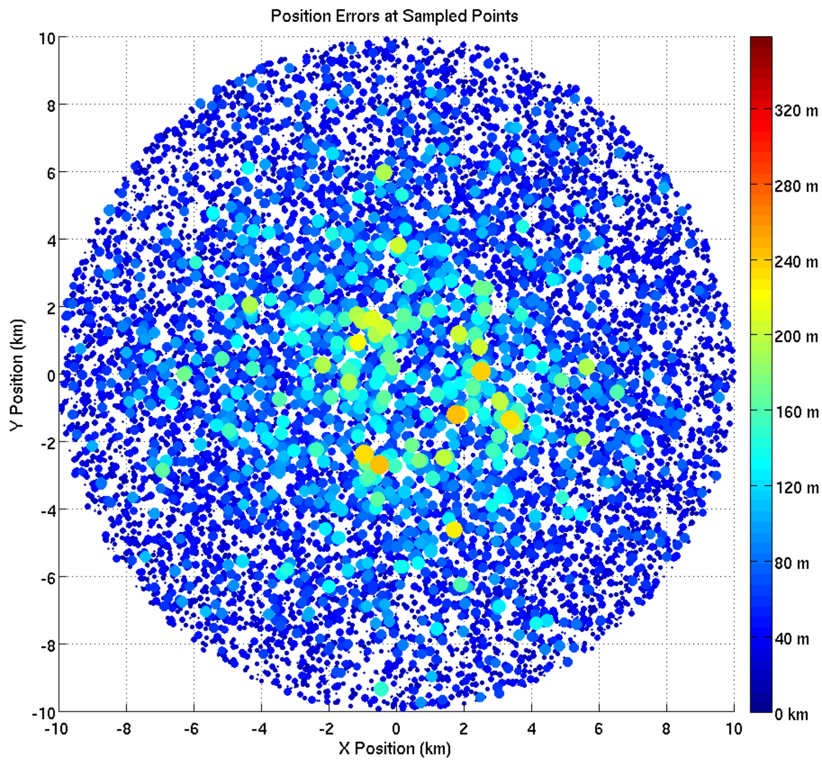

Figure B.5: Position error at 10,000 sampled points in a Shackleton-style crater, starfield 5 – Point size and color scaled to position error

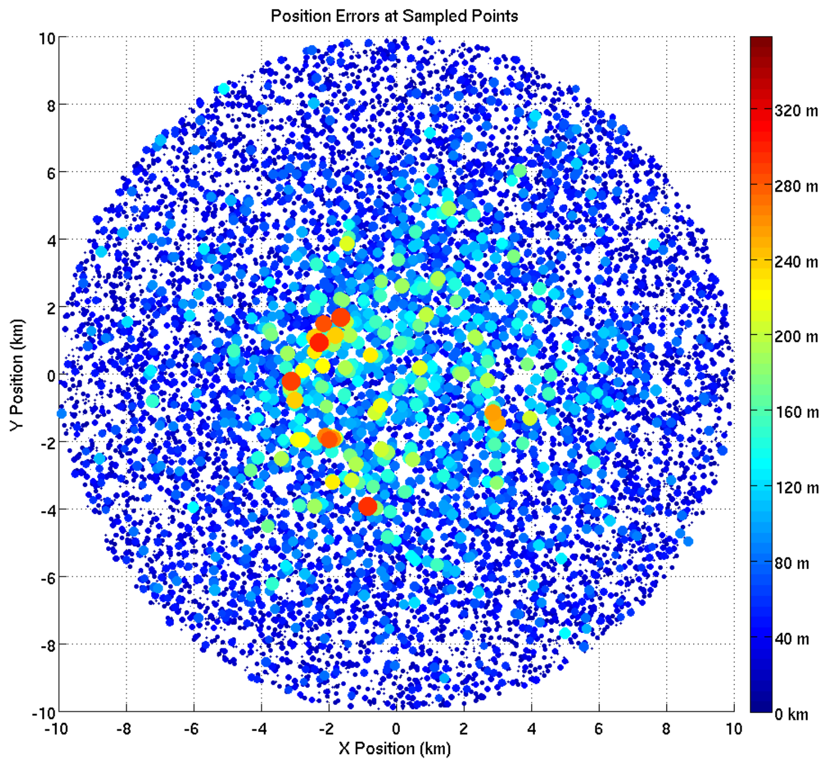

Figure B.6: Position error at 10,000 sampled points in a Shackleton-style crater, starfield 6 – Point size and color scaled to position error

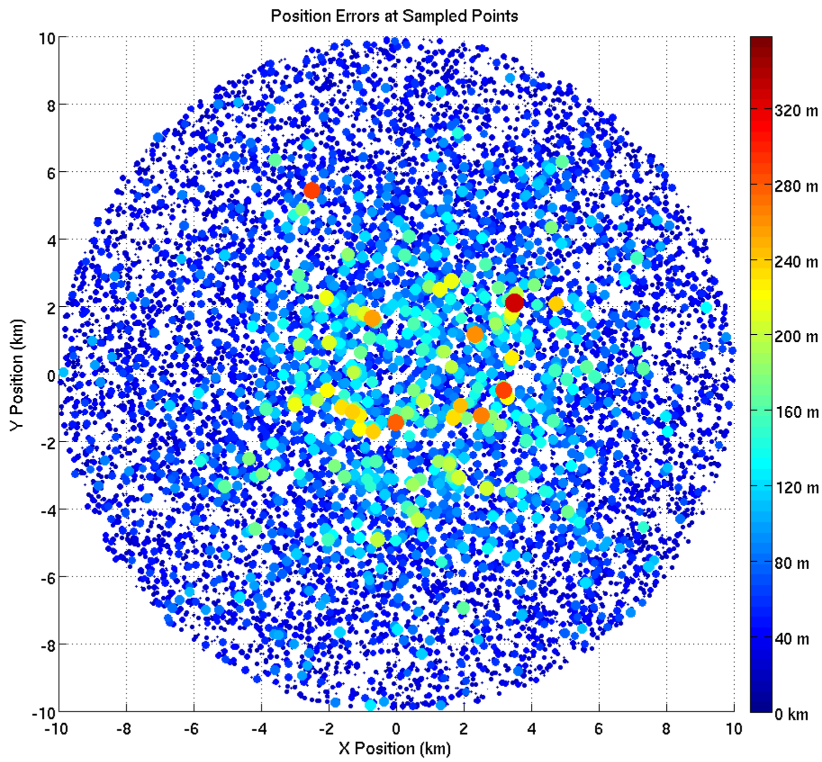

Figure B.7: Position error at 10,000 sampled points in a Shackleton-style crater, starfield 7 – Point size and color scaled to position error

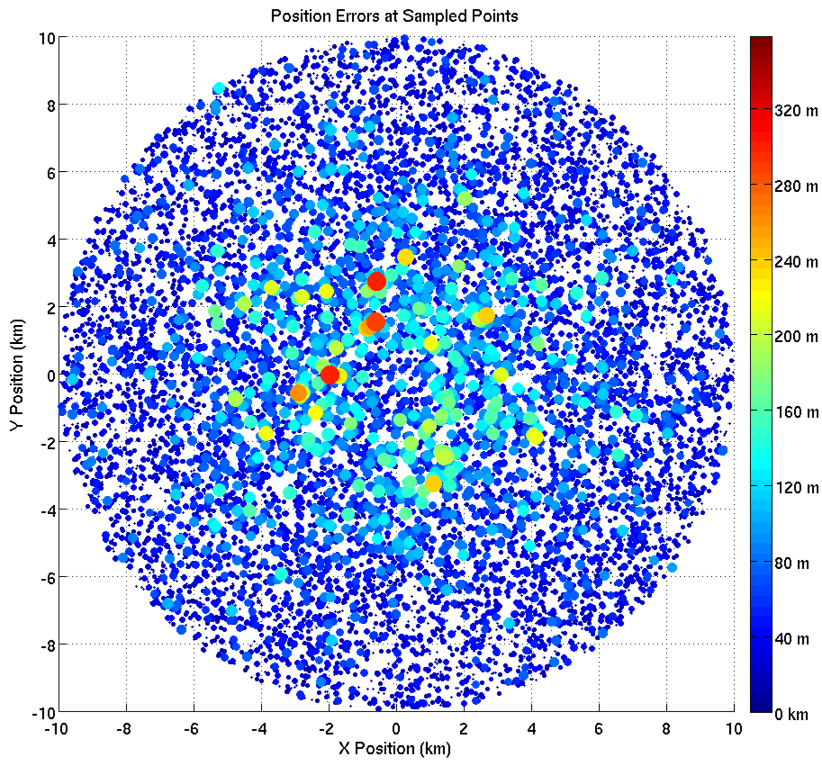

Figure B.8: Position error at 10,000 sampled points in a Shackleton-style crater, starfield 8 – Point size and color scaled to position error

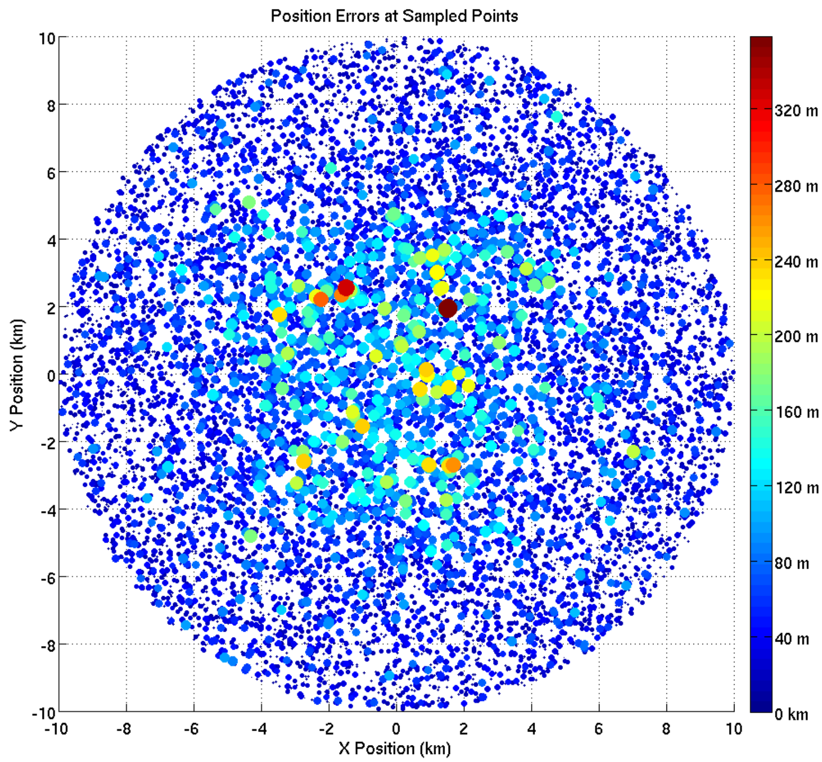

Figure B.9: Position error at 10,000 sampled points in a Shackleton-style crater, starfield 9 – Point size and color scaled to position error

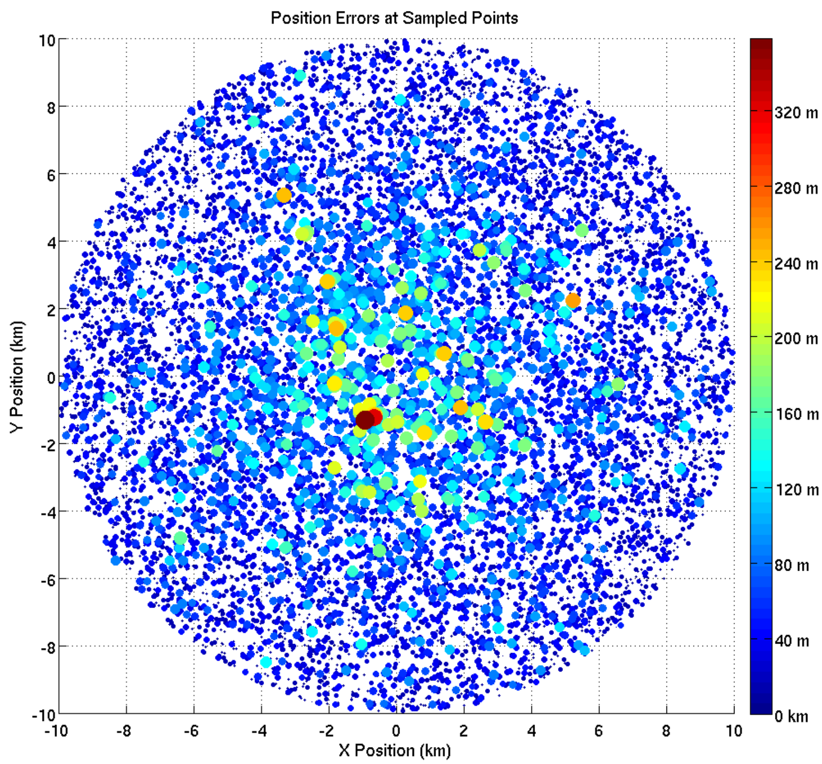

Figure B.10: Position error at 10,000 sampled points in a Shackleton-style crater, starfield 10 – Point size and color scaled to position error

Bibliography

[1] Spudis, P., "Ice on the moon," *The Space Review*, `http://www.thespacereview.com/article/740/1`, November 6, 2006.

[2] Williams, D. R., "Ice on the moon," NASA Article, `http://nssdc.gsfc.nasa.gov/planetary/ice/ice_moon.html`, NASA Goddard Space Flight Center, Greenbelt, MD, March 26, 2007.

[3] NASA Press Release 99-119, "No water ice detected from Lunar Prospector impact," `http://nssdc.gsfc.nasa.gov/planetary/text/lp_pr_19991013.txt`, October 13, 1999.

[4] Bartlett, P. W., Wettergreen, D., and Whittaker, W., "Design of the Scarab rover for mobility & drilling in the lunar cold traps," *International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*, Los Angeles, February 2008.

[5] Thompson, W., Pick, H., Bennett, B., Heinrichs, M., Savitt, S., and Smith, K., "Map-based localization: The 'drop-off' problem," *DARPA Image Understanding Workshop*, pp. 706-719, Pittsburgh, PA, September 1990.

[6] Thompson, W. B., Henderson, T. C., Colvin, T. L., Dick, L. B., and Valiquette, C. M., "Vision-based localization," *DARPA Image Understanding Workshop*, pp. 491-498, Maryland, April 1993.

[7] Sigel, D. A., and Wettergreen, D., "Star tracker celestial localization system for a lunar rover," *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Diego, CA, pp. 2851-2856, October 2007.

[8] Malay, B. P., Gaylor, D., and Davis, G., "Stellar-aided inertial navigation systems for lunar and mars exploration," *2005 Flight Mechanics Symposium*, NASA Goddard Space Flight Center, Greenbelt, MD, October 2005.

[9] Talluri, R., and Aggarwal, J., "Position estimation for an autonomous mobile robot in an outdoor environment," *IEEE Transactions in Robotics and Automation*, Vol. 8, No. 5, pp. 573-584, October, 1992.

[10] Stein, F., and Medioni, G., “Map-based localization using the panoramic horizon,” *IEEE Transactions on Robotics and Automation*, Vol. 11, No. 6, pp. 892-896, December 1995.

[11] Cozman, F. G., “Decision making based on convex sets of probability distributions: Quasi-Bayesian networks and outdoor visual position estimation,” Ph. D. dissertation, The Robotics Institute, Carnegie Mellon University, CMU-RI-TR-97-49, December 1997.

[12] Naval, P. C., Jr., Mukunoki, M., Minoh, M., and Ikeda, K., “Estimating camera position and orientation from geographical map and mountain image,” *Society of Instrument and Control Engineers 38th Research Meeting of the Pattern Sensing Group*, Tokyo, Japan, pp. 9-16, January 1997.

[13] Aguirre, A., Goldfarb, M., Kua, J., Soto, V., Viswanathan, K., Weissenberger, F., Wilson, B., Whittaker, W., “Estimating robot pose in polar lunar craters by finding the edge,” Tech report CMU-RI-TR-06-55, Carnegie Mellon University, December 2006.

[14] Schulman, E. and Cox, C. V., “Misconceptions about astronomical magnitudes,” *American Journal of Physics*, Vol. 65, No. 10, pp. 1003-1007, October 1997.

[15] Liebe, C. C., “Accuracy performance of star trackers – a tutorial,” *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 38, No. 2, pp. 587-599, April 2002.

[16] Dutré, P., *Global Illumination Compendium*, `http://www.cs.kuleuven.ac.be/~phil/GI/TotalCompendium.pdf`, September 29, 2003, p. 19.

[17] Heiken, G. H., Vaniman, D. T., and French, B. M., eds., *Lunar Sourcebook*, New York: Cambridge University Press, 1991.

[18] Liebe, C. C., “Pattern recognition of star constallations for spacecraft applications,” *IEEE Aerospace and Electronic Systems Magazine*, Vol. 8, No. 1, pp. 31-39, January 1993.