Landing Site Dispersion Analysis and Statistical Assessment for the Mars Phoenix Lander

Eugene P. Bonfiglio[1], Douglas Adams[2], Lynn Craig[3], David A. Spencer[4] and William Strauss[5]
Jet Propulsion Laboratory, Pasadena, CA, 91109

Frank P. Seelos[6] and Kimberly D. Seelos[7]
JHU Applied Physics Laboratory, Laurel, MD, 20723

and

Ray Arvidson[8] and Tabatha Heet[9]
Washington University, St. Louis, MO 63130

The Mars Phoenix Lander launched on August 4, 2007 and successfully landed on Mars 10 months later on May 25, 2008. Landing ellipse predicts and hazard maps were key in selecting safe surface targets for Phoenix. Hazard maps were based on terrain slopes, geomorphology maps and automated rock counts of MRO's High Resolution Imaging Science Experiment (HiRISE) images. The expected landing dispersion which led to the selection of Phoenix's surface target is discussed as well as the actual landing dispersion predicts determined during operations in the weeks, days, and hours before landing. A statistical assessment of these dispersions is performed, comparing the actual landing-safety probabilities to criteria levied by the project. Also discussed are applications for this statistical analysis which were used by the Phoenix project. These include using the statistical analysis used to verify the effectiveness of a pre-planned maneuver menu and calculating the probability of future maneuvers.

I. Introduction

HIGH resolution HiRISE images from Mars Reconnaissance Orbiter (MRO) put the Phoenix project in the unique position of being the first Mars lander to ever have knowledge of the exact location of lander-sized rocks near the landing site prior to entry. The lander, targeted for the northern region of Mars was essentially a build-to-print of the failed Mars Polar Lander. The science mission that began when Phoenix touched down in May will be conducted over a 90 sol period (prime mission) and may continue for another 150 sols as part of an extended mission. The instrument suite on Phoenix is capable of performing chemistry analysis of the atmosphere, surface, and subsurface and, in particular, will be used to search for water ice in the northern hemisphere of Mars.

Landing safety analysis is the culmination of work performed in many different disciplines. Surface hazards must be characterized and hazard maps created. The entry state uncertainty must be determined by the navigation team. And finally, the entry, descent and landing must be modeled to get the associated landing statistics. What made Phoenix unique from previous missions that attempted to assess the landing safety statistics (such as MER and Pathfinder) is the surface hazard characterization that was performed. The predicted landing ellipse in the Phoenix

[1] Senior Engineer, Guidance, Navigation & Control Section, M/S 264-623, AIAA Member.
[2] Senior Engineer, Section 352, M/S 303-422, Nonmember.
[3] Senior Engineer, Guidance, Navigation & Control Section, M/S 264-623, AIAA Nonmember.
[4] Phoenix Deputy Project Manager, AIAA Nonmember.
[5] Senior Engineer, Guidance, Navigation & Control Section, M/S 301-220, AIAA Nonmember.
[6] Postdoctoral Research Fellow, Space Department – Planetary Exploration group, M/S MP3-E104, Nonmember
[7] Postdoctoral Research Fellow, Space Department – Planetary Exploration group, Nonmember
[8] James S. McDonnell Distinguished University Professor, Dept. of Earth and Planetary Sciences, Nonmember
[9] Doctoral Student, Dept. of Earth and Planetary Sciences, Nonmember.

landing region, known as 'Green Valley', was literally covered with HiRISE images. Each image spanned about 6 km X 20 km - some of which were taken only 2 months prior to landing. The 30 cm/pixel images were run through automated rock counters to determine the location of every distinguishable rock. This rock counting was performed by Doug Adams, Yang Cheng and Andres Huertas at the Jet Propulsion Laboratory. This knowledge, combined with MOLA (Mars Orbiter Laser Altimeter) slope information and the geomorphology of the landing region was used by the landing site selection team to determine Phoenix's Certified Safe Zone (CSZ). All of this information was essential in creating hazard maps and calculating probabilities of landing in hazardous regions which were used in landing site safety assessment during operations. The resulting assessment was used in determining whether the spacecraft would perform planned trajectory correction maneuvers (TCMs).

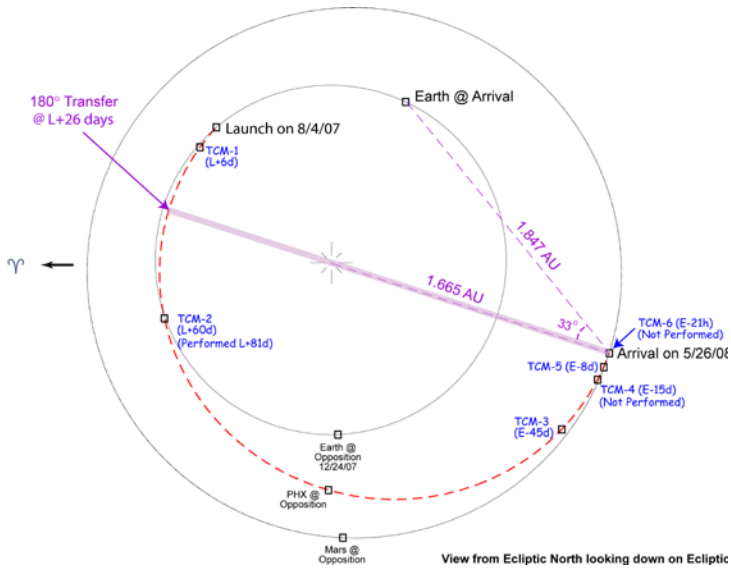

Figure 1 Phoenix Interplanetary Cruise Trajectory

The planned Phoenix trajectory (Fig. 1) contained 6 opportunities to fine-tune the targeted entry point in order to meet the entry flight path angle requirement of ±0.2° (3σ)[1]. The TCM dates were chosen such that they would provide navigators with enough knowledge of the trajectory to perform the maneuver and, therefore, minimize the entry error. Table 1 gives a list of the TCMs and their associated dates. Also listed is the actual maneuver size as calculated by the navigation team during operations[2]. The 'X' maneuvers were planned maneuvers to be used in contingency situations.

As each maneuver date approached, the navigation team estimated the error associated with their calculated entry state[3]. These navigation errors, along with aerodynamic uncertainties were used in a Monte Carlo analysis of the EDL trajectory to calculate the landing dispersion on the ground. The dispersions, combined with the hazard maps, were used to calculate various probabilities that were measured against landing-safety criteria set by the project. The resulting assessment was used in determining whether or not to perform a planned trajectory correction maneuver. Landing safety analysis also played a key role in setting up the contingency maneuver TCM-6XM. This maneuver, planned for 8 hours prior to entry, was to be used only if TCM-6 did not perform sufficiently or if the navigation team later determined that a maneuver was necessary after TCM-6 had been canceled. The 'M' in '6XM' stands for 'menu.' TCM-6XM was a special implementation of TCM-6X requiring 25 pre-calculated/pre-tested 'off-the-shelf' maneuvers that could be quickly assessed (from a landing safety point of view) in emergency situations.

Table 1. Planned TCMs and the Corresponding Maneuver Size

TCM Number	Event Date	Maneuver Size
TCM-1	L+6 Days	18.5 m/s
TCM-2	L+81 Days	3.6 m/s
TCM-3	E-45 days	1.4
TCM-4	E-15 days	Cancelled
TCM-5	E-8 days	.05 m/s
TCM-5X	E-7 days	Cancelled
TCM-6	E-24 hours	Cancelled
TCM-6X(M)	E-8 hours	Cancelled

The following sections discuss the creation of the hazard maps and their role in landing site selection. They also detail the evaluation of the TCM-6XM options and how the real-time landing safety assessment played out in operations leading up to atmospheric entry.

Throughout this paper, the terms 'operations phase' and 'development phase' are used. Operations phase implies the use, in analysis, of actual data that was sent by the Phoenix spacecraft during flight. Work done in the development phase generally means that assumption or predictions were made of future data/performance.

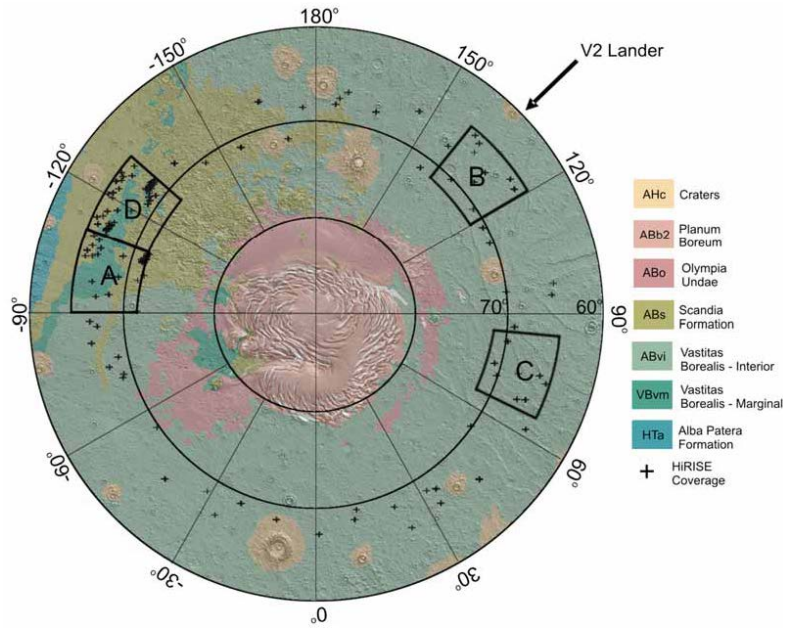

Figure 2 Phoenix Landing Regions A, B, C and D

II. Landing Site Selection Overview

The landing site selection was a multi-year process involving data from many different instruments and Mars orbiters. While science viability was a major consideration in selecting a landing site, this paper concentrates on the landing site safety assessment that was performed in the effort to determine a surface target for the navigation team. Safety assessment was performed considering two main hazard types. The first hazard type was slopes, for which there was a requirement to keep lander tilt to less than 16 degrees. The second hazard was associated with landing on rocks during the touchdown event. While there was no specific requirement with respect to rocks, the goal was to find the least rocky landing site that would meet the science requirements.

A. Initial Regions of Consideration

Initially, 4 large regions (~300 km X ~300 km) were canvassed for candidate landing sites. The regions, labeled A, B, C, and D are shown in Fig. 2. Since it was determined that the regions were all equally valid science targets[4], they were graded based on estimated rock abundance and suitability for Entry, Descent and Landing (EDL). Region B was initially selected because indirect measurements implied low rock abundances and the low altitude of the region would provide the EDL trajectory with margin in their development analysis.

In the fall of 2006, when MRO returned the first HiRISE images, large boulder fields were observed in Region B for the first time. While there were no specific requirements with respect to rock abundance in the proposed landing site, rocks were, by far, the most dangerous ground-based hazard to the spacecraft. Given this new information about Region B and the susceptibility of the lander to rocks, a new search began which leveraged knowledge gained from HiRISE images to find new candidate areas with low rock abundance.

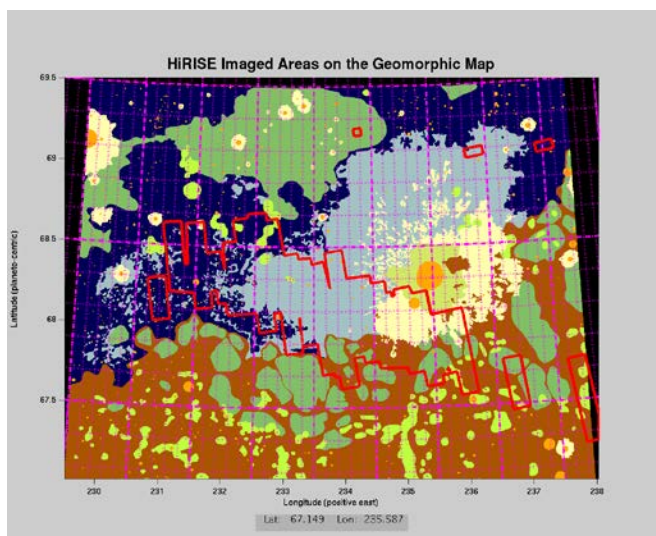

Figure 3 HiRISE Imaged Areas (Red Outline)

The new search relied heavily on HiRISE and imaged a number of different targets in the northern region of Mars. The black plus signs in Fig. 2 are the locations imaged by HiRISE (as of the fall of 2007) during the process of finding a new landing site. The team finally settled on an area that became known as 'Green Valley' in Region D. Green Valley obtained its name from the map classification process the team went through to find candidate landing sites. With pixels on the maps being classified as red, orange, yellow, or green based on their level of rock abundance, Green Valley came back as having relatively low rock abundance and turned up green on the map.

B. HiRISE Imaging Campaign, Rock Counting and Hazard Classification

Once the search had been narrowed down to the smaller area containing Green Valley, a more targeted imaging campaign was planned with the goal of obtaining high resolution images in the portion of the Phoenix map where the lander was most likely to touchdown. Figure 3 is a map of the Phoenix landing region with the terrain has been classified according to its geomorphology[5]. The red contours outline the portion that was covered by high resolution HiRISE images. Each image was processed by automated rock counting software. Since the resolution of the HiRISE images is 30 cm/pixel, the rock counting software was not able to detect every single rock that could be hazardous to the lander (i.e. rocks greater than about 50 cm in diameter). The software was only able to reliably resolve rocks of about 1.5 meters or larger. Size-frequency relationships for rock distribution (i.e. rock abundance as a function of the number of rocks of a particular size) were used to estimate the number or rocks present that were smaller than 1.5 meters[6]. The accuracy of the rock counting software and the size-frequency relationships were verified using ground based images taken by MER and Viking landers. The height-to-diameter ratio of the rocks was assumed to be 0.5.

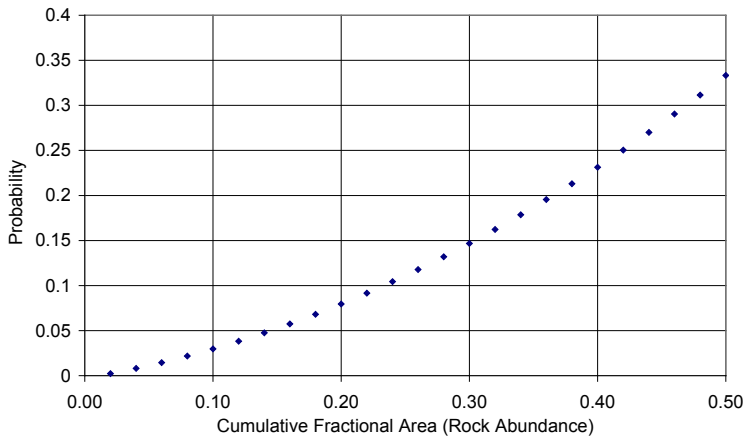

Figure 4 Probability of a Rock Causing a Degraded/Failed Mission

Table 2. Hazard Color Assignments and Associated Rock Abundance

Hazard Color	1.5 Meter Rock Count	Cumulative Fractional Area	P(Fail)
Green	1-3	7%-10%	1.94%-2.99%
Yellow	4-8	10%-14%	3.41%-4.88%
Orange	9-19	14%-21%	5.23%-8.56%
Red	>19	>21%	>8.56%

After the software was verified, every high resolution image that was taken within the Phoenix map, displayed in Fig. 3, was processed and the associated cumulative fractional area (fraction of the ground covered by rocks) was calculated for each 100 meter X 100 meter square (1 hectare). The rock abundance was then related to the probability of a rock causing a degraded or failed mission using the curve in Fig. 4. In Phoenix's case, a degraded mission was one in which a rock has impacted the underside of the lander, caused instability, or prevented either solar array from opening. The data in Fig. 4 was calculated by Lockheed Martin from a Monte Carlo simulation of 2000 touchdown states that were generated using an EDL trajectory simulation software program.

Each hectare square region was given a corresponding hazard color based on the criteria outlined in Table 2. For instance, if the rock counting software counted five 1.5 meter rocks in a particular hectare, that square would be classified as yellow and assigned the probability of failure associated with Fig. 4. Hectares in Fig. 3 where HiRISE images did not exist were assigned a constant cumulative fractional area based on the geomorphic unit within which they were located. The assigned value for a given geomorphic unit corresponded to the 75%tile rock counts from HiRISE images in that particular geomorphic unit. For example, 75 percent of the rock counts performed in the dark blue region (know as Lowland Dark) yielded as many as 18 rocks. Therefore, hectare squares in the Lowland Dark geomorphic unit where no rock counts could be performed were assumed to have 18 rocks and a corresponding hazard color of orange. The final hazard map can be seen in Fig. 5. As can be seen by comparing Fig. 3 with Fig. 5, the light blue (Lowland Bright) and light yellow (Exterior Crater Ejecta) geomorphic units both correspond to green hazards when using the 75%tile standard. These geomorphic units were considered to be safe for landing even in areas where there were no automated rock counts.

C. Phoenix's Certified Safe Zone

The hazard map, which was based on rock counts, accounted for the rocks. As mentioned above, Phoenix also had requirements on slopes. While the landing region was very flat, slopes were still accounted for when determining what would become known as Phoenix's 'Certified Safe Zone' (CSZ). Calculating slopes at the scale of the lander (~1 meter) is extremely difficult since it requires stereo image pairs be taken by the HiRISE camera. Since very little overlap existed within the HiRISE imagess, the project had to rely on extrapolation of MOLA data at the scale of a few hundred meters/pixel to estimate lander-scale slopes. While the basis for the extrapolation is not presented in this paper, a conservative analysis showed that 5 degree slopes at the MOLA-scale could correspond to the 16 degree requirement at the lander-scale. Figure 6(a) shows the Phoenix map with white areas representing regions where the MOLA slopes were less than this 5 degree limit. Black areas were assumed to violate the slope requirements.

The Certified Safe Zone (CSZ) was determined by specifying 3 'Keep Out Zones' (KOZ) maps and then placing them on top of each other. The remaining region that was not covered by at least one of the KOZs was classified as the Certified Safe Zone. The first KOZ was any area not of the geomorphic units Lowland Bright (light blue in Fig 3.) or Exterior Crater Ejecta (light yellow in Fig. 3) *and* not marked green by the auto rock counts. The second KOZ was defined as any region that violated the 5 degree MOLA slope limit. The third and final KOZ was any region that had been classified as red from the auto rock counts. The KOZ masks were laid down on top of each other correspondingly. Figure 6(b) gives a Venn diagram picturing how the KOZs were used to define Phoenix's Certified Safe Zone. Figure 6(c) depicts the map of the final Certified Safe Zone (green) prior to landing. In this picture, the large contiguous area makes up the 'Green Valley' that designated the landing site. The small red area in the CSZ located to the northwest of the target in Fig. 6(c) would become known as 'The Hill' and played a very important role in the decision process during the days and hours leading up to landing.

D. Phoenix's Landing Site

The Certified Safe Zone map is what was ultimately used in determining the specific landing site target for Phoenix. The target, labeled in Fig. 6(c), was located at 68.151 degrees areocentric latitude and 233.975 degrees east longitude. The yellow ellipse in Fig. 6(c) is a 100 km X 20 km ellipse with a 111 degree azimuth that was used in the development phase (i.e. prior to any real operations

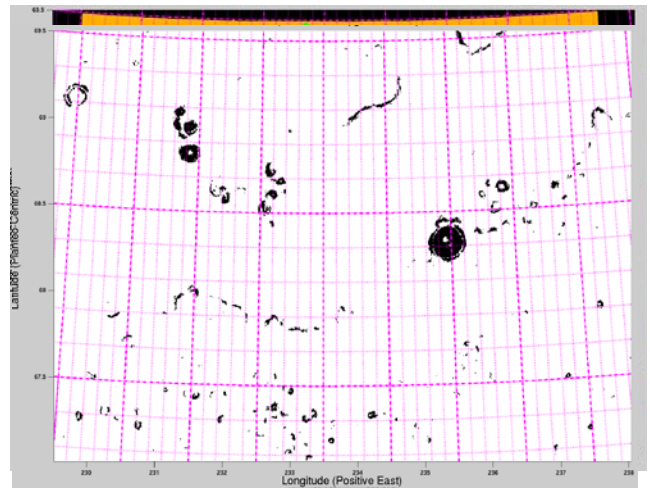

a) **Figure 5. Phoenix Hazard Map**

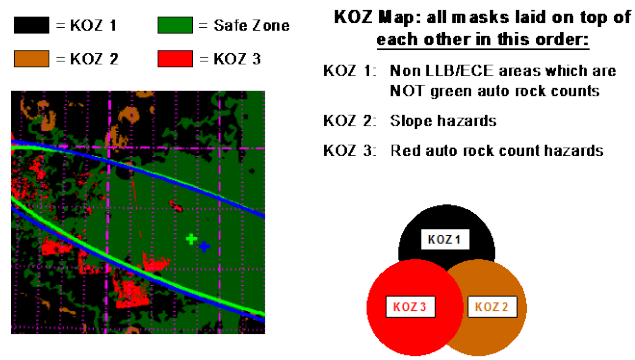

b)

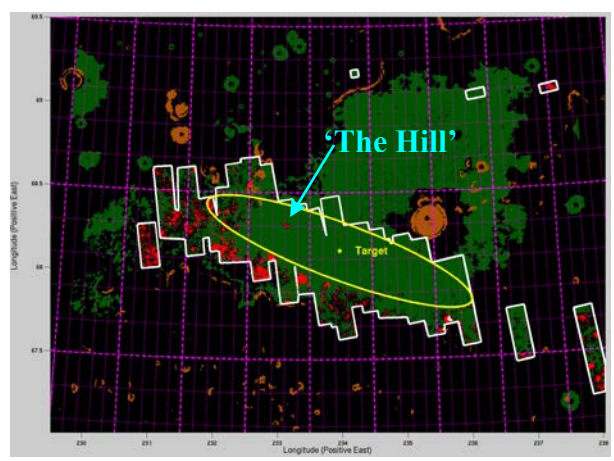

c)

Figure 6. a) MOLA Slope Hazard Mask b) Venn Diagram Describing the Keep Out Zone Layering c) Phoenix's Certified Safe Zone Map

entry estimates) as the expected delivery accuracy. This ellipse factors in the ±0.2° entry flight path angle delivery requirement placed on navigation. While the ellipse size was limited by navigation capability, the location of the ellipse was essentially a free parameter. Phoenix's specific target was selected based on a desire to maximize the probability of landing in the CSZ. Looking carefully at the map in Fig. 6(c), it's clear that the target is not exactly in the middle of the portion of the CSZ that runs parallel to the major axis of the ellipse. This is because not all of the HiRISE images had arrived at the time of the final landing site selection in April 2008. Had it been known that all of the images to the southeast of the target would have come back green, the target certainly would have been shifted in that direction.

III. Landing Dispersion Determination and Associated Applications

In order to perform landing site assessment, two items are needed. The first is the landing site characterization as described in Section II. The second is a representation of the probability of the spacecraft landing in a given region. This section will concentrate on the tools used to determine the likelihood of Phoenix landing in a particular (latitude, longitude) location. Once this probability is represented in a mathematical way, it is convolved with the landing site characterization to calculate the probability of landing in the characterized regions. This information was essential in defining criteria that the landing site safety team set for determining whether or not to recommend TCMs. It was also used to help quantify the effectiveness of the TCM-6XM options that the project had predefined for contingency scenarios as well as to calculate the probabilities of needing future maneuvers in operations.

Table 3 Comparison of DSENDS and POST for a Monte Carlo Analysis Performed During the Development Phase

Comparison Against EDL Anchor (TCM-6 Delivery ESF with Expected Dispersions)		
	DSENDS	**POST (Anchor)**
Mean Chute Time	**819.149 s**	**820.089 s**
3-Sig Chute Time	**+/-16.317 s**	**+/-16.414 s**
Mean TD Lat/Lon	**68.191/233.293**	**68.195/233.302**
Mean TD Miss Dist	**307 meters**	
Nominal Lat/Lon	**68.174/233.356**	**68.179/233.352**
Nom. TD Miss Dist	**304 meters**	
Ellipse Size (km)	**103.10 X 19.57**	**103.12 X 19.96**

A. Monte Carlo Analysis and Certification

Phoenix EDL trajectory simulations were performed independently by the EDL systems team, using POST[7,8,9,10], and the Navigations team, using a Python/C++ based tool called DSENDS. POST was mainly used for EDL system design[11], validation, and verification. The DSENDS tool, which went through a rigorous certification process by both JPL's Guidance, Navigation & Control element *and* the Phoenix project, was the main tool used for performing maneuver targeting, landing dispersion analysis and landing safety assessment.

The EDL simulation in DSENDS differed from POST in that it did not attempt to model the EDL system to the level of fidelity that was achieved in POST. While the simulation did have an accurate 6-DOF model of the trajectory through parachute deploy, it only had an engineering model of certain parts of the flight software which were relevant to accurately predicting landing location. (Namely, the parachute deploy triggers and separation of the lander from the parachute.) The parachute deploy trigger in the DSENDS simulation was very similar to that used by actual flight software, although the lander separation was commanded at a constant 1000 meters altitude in the DSENDS model whereas the trigger in flight software was actually based on the propulsive descent event. DSENDS did not model the reaction control system (RCS) on the Phoenix

Table 4 Dispersed Parameters in the Monte Carlo Simulations

Dispersed Parameters	Source for Dispersion
Entry State	PHX Navigation Team
Aerodynamic Coefficients	NASA Langley
Atmospheric Density	PHX Atmosphere Working Group
Atmospheric Winds	PHX Atmosphere Working Group
CG Location	PHX Spacecraft Team (Lockheed)
Mass	PHX Spacecraft Team (Lockheed)
Inertias	PHX Spacecraft Team (Lockheed)
Parachute Drag	NASA Langley

spacecraft. This was not necessary because the attitude and attitude rate error deadbands in the flight software had been increased to the point that there was almost no chance of the RCS thrusters firing during the hypersonic phase. Also, the simulation had no IMU or radar modeling in DSENDS and, therefore, assumed perfect knowledge for all of the major events during the EDL trajectory. Finally, the descent thrusters were not modeled either because the distance traveled over the ground during the near-vertical propulsive phase is insignificant. Table 3 is a comparison of POST and DSENDS that was presented at a certification review. This comparison is for a Monte Carlo run performed in both DSENDS and POST during the development phase. The comparison clearly shows the differences between the two tools (~300 meters) are minor compared to the size of the landing ellipse (~103 km).

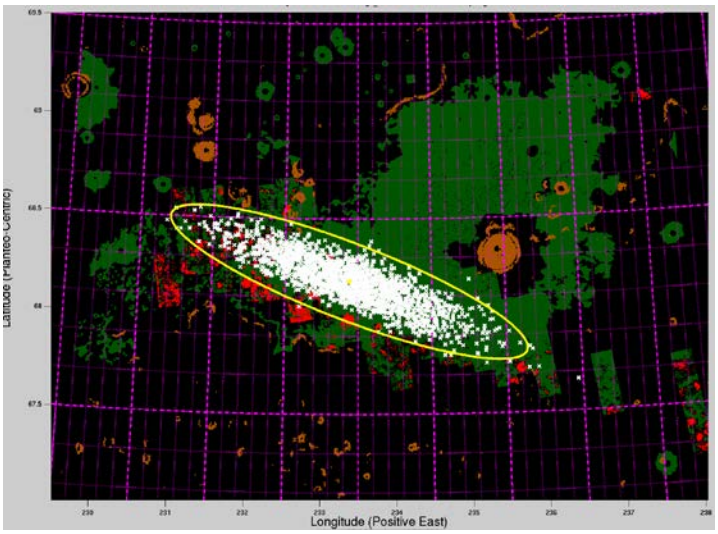

Figure 7. CSZ Map Shown With an Example of Landing Points and the Corresponding 99% Ellipse

Table 5. Relationship Between Sigma Ellipses & Sample Percentage for Bivariate Gaussian Distributions

Ellipse Sigma	Percentage of Samples in Ellipse
1-Sigma	39.35%
2-Sigma	86.47%
3-Sigma	98.89%

The inputs in the Monte Carlo process captured all of the dispersions known to have a significant impact on the landing location. The Monte Carlos consisted of a nominal run and 2000 dispersed runs, each with an independent set of inputs dispersed per the recommendation of the sources given in Table 4[12]. Figure 7 is the CSZ map with 2000 cases plotted on the ground (white) from a sample Monte Carlo run. The distribution of these points is assumed to be bivariate Gaussian. (the validity of this assumption is discussed later in Section V). A common way to describe this distribution is to give the major and minor axis of the 99% ellipse (~3.02 sigma ellipse for a bivariate distribution) and an associated azimuth which is usually measured with respect to north. The 99% ellipse, as its name implies, contains 99% of the expected cases. The 99% landing ellipse associated with the dispersed landing points in Fig. 7 has a major axis of about 100 km X 20 km and a 111 degree azimuth. Table 5 gives the relationship between ellipse size in terms of sigma and the percentage of cases that are contained within that ellipse. For example, a 1-sigma ellipse contains 39.35% of the cases.

B. A New 'Pixel-Based' Method for Calculating the Probability of Landing in a Hazard

With the landing site fully characterized and the ability to estimate the landing distribution developed as described above, the project spent a great deal of time convolving these into relevant probabilities of landing in particular hazard types (the probability of landing in the CSZ, for example). The main tool used for these calculations is called MarsLS[13]. MarsLS was also used by the Mars Exploration Rover (MER) project to calculate probabilities of landing in safe and scientifically interesting regions, although

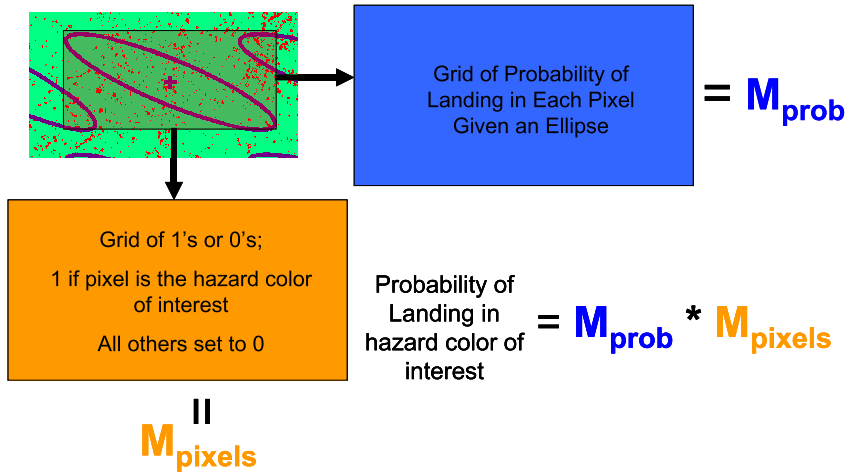

Figure 8. Illustration of Calculating Probability of Landing in a Hazard

modifications to the MER algorithm for calculating the probability had to be made due to computational run-time issues. In the case of MER, the hazard maps were usually hand drawn and the hazardous regions were generally consolidated into large, contiguous regions. This allowed the EDL analysts on MER to draw contours around the hazardous regions and then calculate the probabilities of landing in those contours.

With hand drawn maps, those hazardous regions are typically limited to a few hundred, at most. In the case of Phoenix, automated programs were used to count rocks and set individual pixels in a 100 m/pixel map their respective hazard color. This meant that a Phoenix hazard map could contain hundreds of thousands of isolated hazard areas (sometimes as small as a single pixel). Of course, because of the high resolution images that were used, there was high confidence in the classification scheme that could sandwich a single green pixel right between a pair of red pixels.

While the process used by MER to calculate probabilities could take minutes with hand drawn maps, it could take hours or even days with the Phoenix style maps. Given this computational headache, a new scheme (referred to as the Pixel-Based method) was developed to calculate the probability of landing in a particular hazard type. This new scheme is represented in Fig. 8. A grid is drawn around the ellipse for which you want to calculate probabilities of landing in a particular hazard. Each square in the grid is actually one of the pixels on the map. The grid can be used to create two different matrices. In one case, you calculate the probability of landing in each grid square (i.e. pixel on the map) assuming your bivariate Gaussian distribution. This matrix becomes M_{prob}, or the matrix of probabilities. In the other case, the grid is turned into a matrix of ones and zeros. A one is placed everywhere that the hazard color of interest exists in the grid and zero everywhere else. This matrix becomes M_{pixels} in Fig 8. Multiplying the corresponding elements of these two matrices and summing the resulting elements gives the cumulative probability of landing in the hazard of interest. Note, that this is *not* matrix multiplication, but simply an element-by-element multiplication two matrices. The summation of the elements in the resulting matrix gives the probability of landing in the hazard color of interest.

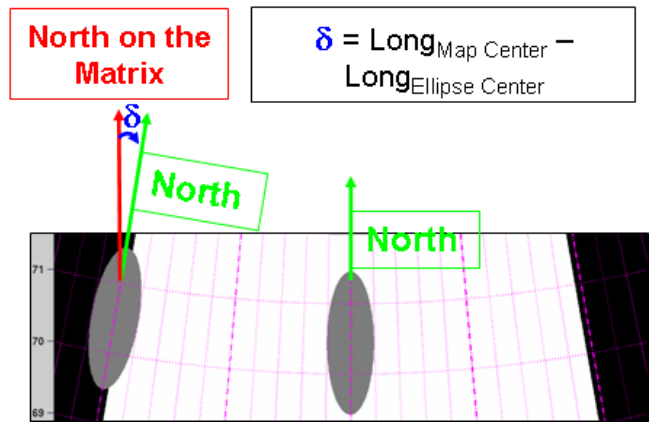

Figure 9. Illustration of the Required Azimuth Adjustment When Using the Pixel-Based Method

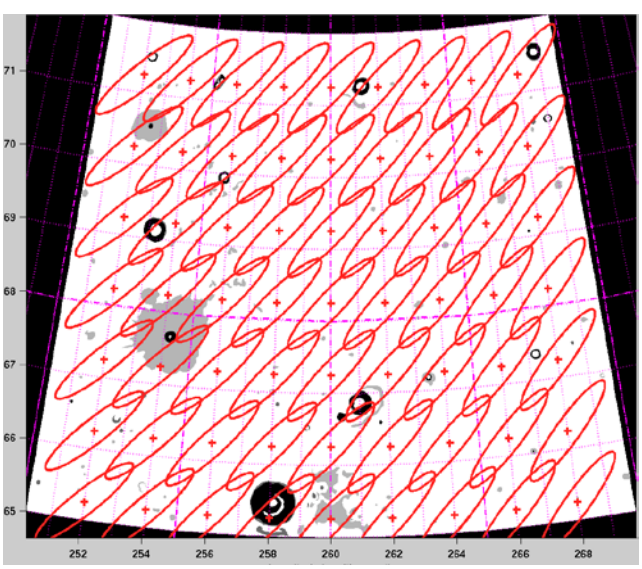

Figure 10. Illustration of Ellipses Placed on a Map in a Grid Pattern

While the Pixel-Based method of calculating probabilities is much faster than what was used on MER, there is an error source that needs to be accounted for when dealing with north polar projected maps, as was used on Phoenix. The problem arises when calculating the M_{prob} matrix. On a polar projected map, the direction to north varies with respect to the map axes depending on the longitude. At the longitude of the center of the map projection, north is 'straight up' on the map. But if you are at a longitude that is 10 degrees away from the center of projection, the direction of north is actually 10 degrees from 'straight up.' Since part of the ellipse definition is azimuth, which depends on the direction of north, this offset must be accounted for when filling in the M_{prob} matrix. That is, since the squares of the matrix are orthogonal to each other, the direction to north in the matrix is always 'straight up.' So if the ellipse of interest has an azimuth of 0 degrees, but is located 10 degrees to the west of the center longitude of projection, then when filling in the M_{prob} matrix, the azimuth of the ellipse must be adjusted to be 10 degrees. If the ellipse were 10 degrees to the east, the azimuth would be adjusted to -10 degrees when filling in the matrix. This is illustrated in Fig. 9. There are other error sources of course; most

due to the fact that a a two-dimensional, flat surface distribution is being forced onto a sphere which is being represented in a map, but the azimuth adjustment accounts for the most significant part of the error. A certification of this new algorithm, using the azimuth adjustment, showed agreement to within less than 0.1% of expected results and about 0.1% when compared to the algorithm used by MER.

C. Using Probabilities to Assess the Landing Site and Define Landing Site Criteria

In order to fully understand the hazard maps, which include the CSZ map, it was necessary to determine relevant probabilities as a function of landing site target, or ellipse center. For instance, how does the probability of landing in the CSZ (P_{CSZ}) change as you move the target on the map? Of course, it may be obvious from Fig. 7 where the 'sweet spot' on the map is, but in reality the navigation team cannot hit an entry point perfectly. Even if there are opportunities for future TCMs, it may be desirable to cancel them if the current entry condition is 'good enough.' In order to understand this function of P_{CSZ} and ellipse center, 10,000 ellipses of the same size and azimuth were placed in a grid pattern on the map, as illustrated in Fig. 10. The ellipse size (68 km X 20 km) and azimuth (111 degrees) that was used corresponded with the knowledge of the trajectory that the navigation team would have at the time of the TCM-6 decision meeting. The P_{CSZ} was calculate for each ellipse. The results were then used to create the shaded contour map in Fig. 11. The magenta curves in Fig. 11 correspond to the 80%, 90% and 99% contours of P_{CSZ}. This plot gives the P_{CSZ} for an ellipse centered at a particular latitude and longitude. For example, if the orbit determination solution of the navigation team resulted in an ellipse whose center was on the outermost magenta curve, that ellipse would have an 80% P_{CSZ} – provided that the ellipse was 68 km x 20 km in size.

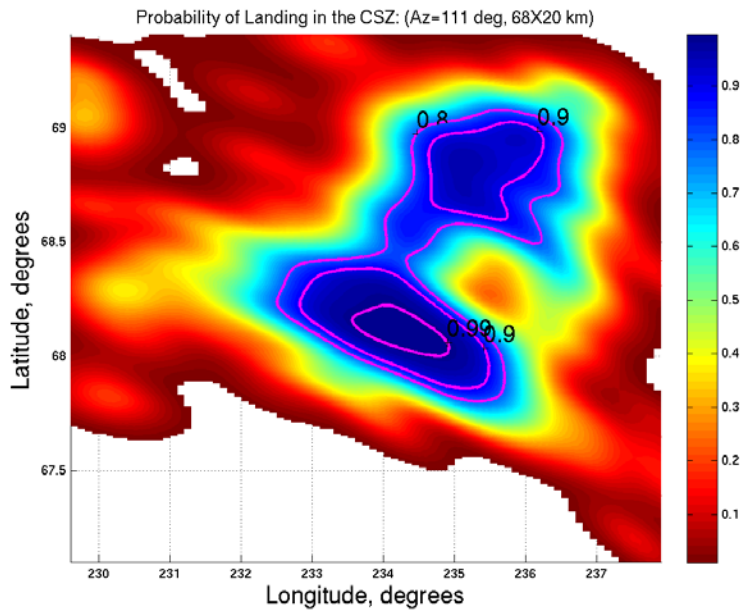

Figure 11. Shaded Contour Map of the Probability of Landing in the CSZ for the Phoenix Landing Region (Magenta Curves Outline the 80%, 90%, and 99% Contours)

Table 6. Final LSS Criteria Set by the Project

Criterion Number	Primary Decision Factor	Criterion
1	Percent of 1σ landing ellipse in CSZ (by area)	95%
2	Percent of 1σ landing ellipse in safe MOLA slope region (by area)	99%
3	*Probability* of landing in a Certified Safe Zone	80%
4*	Percent of 2σ landing ellipse in CSZ (by area)	97%
5	*Probability* of landing in a high resolution HiRISE image	Consideration only

* **Note that criteria number 4 was added *AFTER* TCM-5**

Plots similar to Fig. 11 were used to help define the landing site safety (LSS) criteria was used during the development phase. The purpose of the criteria was to determine if an ellipse on the ground fit within the risk posture of the project. If it did not, that would mean a 'Go' decision for the next planned maneuver. Because the probability as a function of landing site is nonlinear problem and deciding if something fits within a risk posture is subjective by nature, setting criteria is a very difficult task. The difficulty was compounded by the fact that many of the images covering the southeastern portion of the ellipse were note available until less than two months before landing. Also, during development phases, subsystems typically hold margin on their design uncertainties while their performance in operations can often be much better than what those uncertainties would have suggested. Consideration needed to be taken into account of scenario in which engineers had possibly overestimated their uncertainties and Phoenix's

'real' landing dispersion would be better represented by a smaller ellipse. This is similar to saying that the lander could have a very high probability of landing in the 1-sigma ellipse compared to the 39.35% probability normally associated with a 1-sigma ellipse with 'perfect knowledge' of the uncertainties. And on top of all of this, if the criteria were made too tight, the navigation team would never be able to meet them.

The final set of criteria agreed upon by the Phoenix project is enumerated in Table 6. The first and second criteria deal with the concern that overly conservative uncertainties could have guaranteed Phoenix would land in the 1-sigma ellipse. These criteria aim to keep the center of an acceptable landing ellipse as vacant of rock and slope hazards as possible while not imposing unacceptable constraints on the navigation team. Note that the first and second criteria do *not* specify a probability. They simply require that the area within the 1-sigma ellipse be covered by a specified percentage of a safe region. The third criterion addresses the issue of overall probability. This states that the resulting ellipse shall have a probability of at least 80% of landing in the CSZ. This may seem like a low number, but considering that most red rock hazards actually had a probability of success greater than 85%, the overall probability of success is typically greater than 96% for ellipses that meet the criteria in Table 6. Criterion number 4 was added after TCM-5. The reasons for this will be discussed in Section IV. Criterion number 4 turns out to be the strictest of all of them. It requires that 97% of the area in the 2-sigma ellipse be covered by the CSZ. This is an additional guarantee that the most likely landing locations are associated with the safest possible terrain. The last criterion listed did not have a specific requirement associated with it. The probability of landing in HiRISE images was reported as a matter of course because much of the area in the map had not been imaged by HiRISE. The project wanted to be made aware if there was had a high probability of landing in a region that had only been characterized on a statistical basis, rather than from actual rock counts.

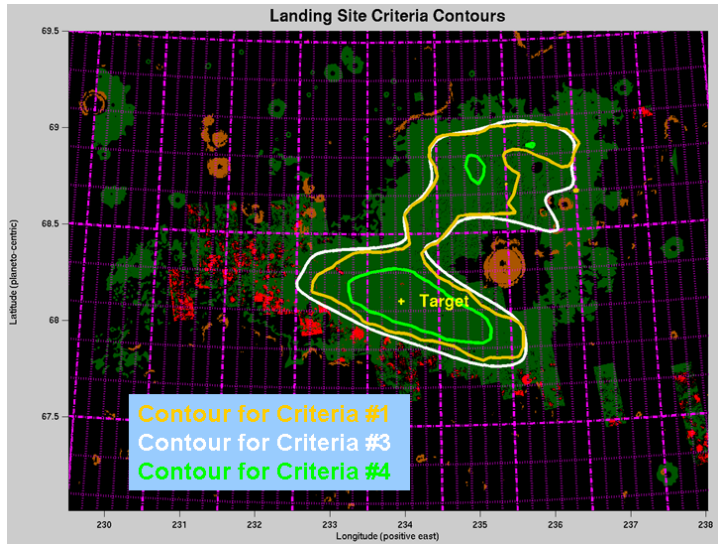

Figure 12. Landing Site Criteria Contours Used in Operations

Through the analysis used to create Fig. 11, contours that meet the landing site criteria can be created. The criteria contours were useful in operations because they allowed us to plot the nominal trajectory solution as the data came into the navigation team and quickly assess whether or not the solution would meet the criteria. (The nominal solution on the ground tends to be close to the mean of the landing points that are calculated from the EDL Monte Carlo.) Creating the criteria contours in advance requires an assumption of the ellipse shape. The criteria contours, plotted on the CSZ map in Fig. 12, assume an ellipse shape based on trajectory knowledge that the orbit determination team *expected* to have at the time of a TCM-6 Go/No-Go decision (i.e. 68 km X 20 km). As it turns out, the contour is not very sensitive to the exact ellipse size. If the major axis of the actual ellipse that is determined

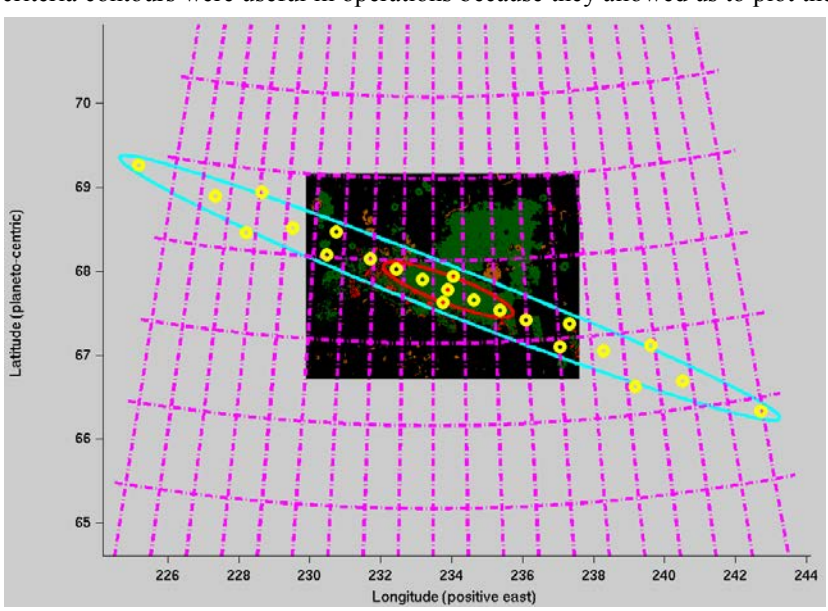

Figure 13. Illustration of the Phoenix TCM 6X Menu Options

through the Monte Carlo process is within about ±10 km of that used to create the contours, then the contours can be quite useful in performing 'first look' assessments of the situation in operations. Note that there is no contour given for criteria number 2. This contour was never calculated because, with so few high slopes in the landing region, a situation was never encountered where the slope criteria was not being met. In practice, the contours in Fig. 12 were quite useful because they represent the *locus of ellipse centers* which meet the criteria given in Table 6. Generally speaking, if a navigation solution was mapped to the ground and it fell within the contours, it was expected that once all the Monte Carlo and probability calculations were done, the solution would meet the criteria as set forth by the project.

D. TCM-6X Menu Analysis

Another application to the Phoenix project of the probability assessment that MarsLS could perform was to determine the effectiveness of the TCM-6X menu options that the navigation team defined for contingency purposes. TCM-6X was the E-8 hour contingency maneuver that Phoenix had planned in the event of an emergency

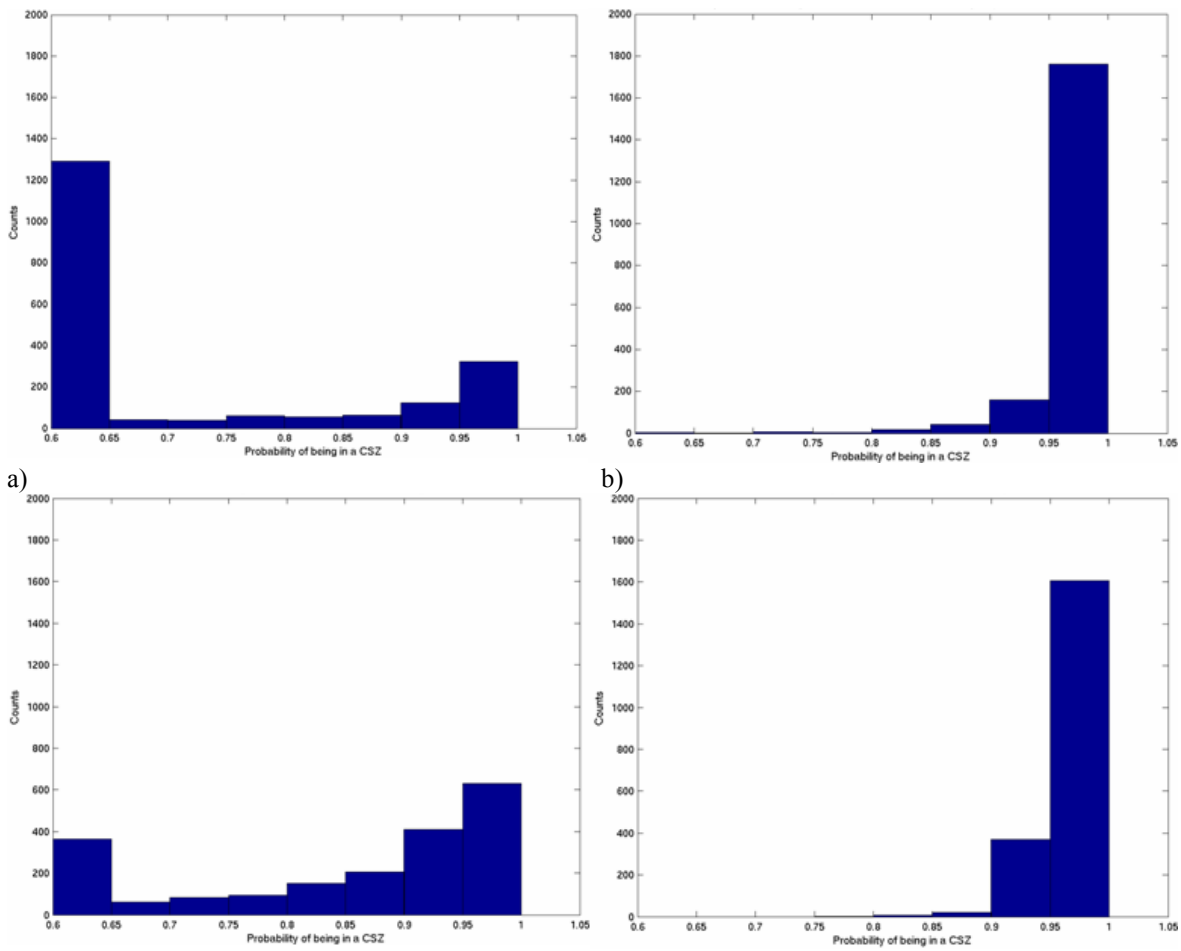

Figure 14. Histograms Comparing the Number of Samples from a 2000 Monte Carlo Simulation that Meet the CSZ Probability Criteria After Applying the TCM-6X Menu Options

(i.e. TCM-6 either did not execute as planned for some reason or the navigated solution changed to such a degree that a last minute maneuver was required to hit the landing target). In this event, menu options (maneuvers predetermined to move the landing solution on the ground by a fixed amount). There were a total of 25 menu options, 22 of which are pictured in Fig. 13. The maneuvers for the menu options were designed such that, if the entry condition solved for by the navigation team put the lander exactly on one of the 22 yellow circles, the corresponding maneuver, once performed, would move that landing point back to the target. Likewise, if the

navigation solution predicted a landing 10 km east of a given yellow circle in Fig. 13, performing the maneuver associated with that yellow circle would result in a predicted landing which was landing 10 km east of the target. The light blue ellipse in Fig. 13 corresponds to the delivery ellipse from TCM-5 that was used in the development phase. That is, the ellipse outlines the worst case of where the navigation team predicted TCM-5 would put Phoenix, mapped to the ground. The red ellipse is the navigation teams TCM-6 delivery ellipse.

An analysis was performed to assess the effectiveness of the TCM-6X menu with respect to the TCM-5 and TCM-6 delivery ellipses in Fig. 13. That is, for a sample set of landing points corresponding to the delivery ellipse, how effective was the menu at increasing the overall probability of landing in the CSZ (P_{CSZ}). For this analysis, a Monte Carlo was run where the navigation error corresponded to the TCM delivery error. This would result in landing points that could be anywhere in one of the Fig. 13 delivery ellipses – depending on which TCM delivery was being assessed. The TCM-6X menu option maneuvers were then applied to each of the 2000 landing points, creating 2000 sets of 25. Each time the menu set is applied to one of the original 2000 landing points, that point moves on the ground by an amount equal to the offset between the target and the yellow points in Fig. 13. The When the 25 maneuvers are 'applied' to the original landing point, the P_{CSZ} should improve for some of them. For each set of 25, the resulting landing point with the highest P_{CSZ} was selected. In essence, the selected point corresponds to the optimal TCM-6X menu maneuver for the original, unshifted landing point. The results are given in histogram format in Fig. 14. Figure 14(a) shows the histogram of the 2000 points before the optimal TCM-6X menu maneuver is applied with the x-axis giving the P_{CSZ}. Figure 14(b) shows the histogram of the 2000 points *after* the optimal TCM-6X menu maneuver is applied. Figures a & b correspond to the TCM-5 delivery error (which would be realized, for example, if TCM-6 failed to execute). Figures c & d provide similar plots, but for the TCM-6 delivery error. A scenario where this would be applicable would be if TCM-6 executed with high errors for some reason. In either case, the TCM-6X menu is highly effective at increasing the P_{CSZ} to at least 0.90 (or 90%), which is well above the criteria #3 limit of 80%.

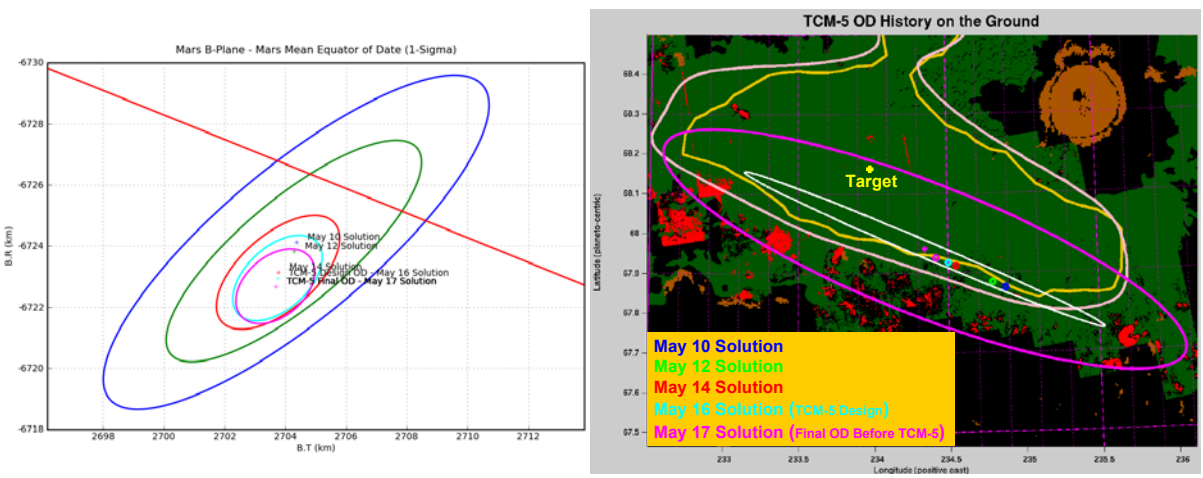

a) TCM-5 OD B-plane Plot **b) TCM-5 OD Solutions Mapped to the Ground**

Figure 15. OD Solutions Leading Up To TCM-5

IV. Landing Dispersion Evaluation Performed in Operations

Most of the work dealing with landing site safety analysis of the 'as flown' trajectory could only be begin after the end of April, ~1 month before landing. Prior to this point, the navigation error was so large that it did not make sense to do any operations analysis. In fact, prior to TCM-3 the predicted landing site was closer to one of the MER landers than the actual targeted landing site. This changed, however, with a near flawless execution of TCM-3 by the spacecraft.

A. TCM-4 and TCM-5

TCM-3, a 1.4 m/s burn put Phoenix on a target that would eventually be predicted to miss the landing site by only 15 km. Because TCM-3 was so accurate, the project decided that there was no need for TCM-4 and cancelled the maneuver. Figure 15 provides two corresponding plots of the orbit determination (OD) solutions leading up to TCM-5, which was successfully executed on May 18, 2008 04:00:00 (UTC). The plot in Fig. 15(a) gives the b-plane solutions (i.e. the solutions which relate to the entry point calculated by the navigation team). The ellipses in

Table 7. Monte Carlo Probability Results for OD Solutions Leading Up to TCM-5

Primary Decision Factor	Criterion	May 10th Solution	May 12th Solution	May 14th Solution	May 16th Solution	May 17th Solution
Percent of 1σ Landing Ellipse in CSZ (By Area)	95%	Ellipse too large	Ellipse too large	96.3 %	96.0%	96.9%
Percent of 1σ Landing Ellipse in safe MOLA slope region (By Area)	99%	Ellipse too large	Ellipse too large	100.0%	100.0%	100.0%
Probability of landing in a Certified Safe Zone	80%	Ellipse too large	Ellipse too large	89.9%	92.4%	93.4%
Probability of landing in a HiRISE image	None	Ellipse too large	Ellipse too large	96.7%	98.5%	99.2%

Fig. 15(a) depict the navigation uncertainty in the OD solution and the small plus signs in the center of the ellipses represent the nominal OD solutions. The solutions presented in the plot span from May 10th to May 17th, right before the maneuver was executed. The cyan-colored solution was used to design the maneuver and the magenta-colored solution is the best knowledge the project had of the orbit before the designed maneuver was executed.

The plot on the right (Fig. 15b) gives the OD nominal solutions mapped to the ground (colored circles). The contours for criteria #1 and #3 are given in gold and pink. Recall that in the days leading up to and including TCM-5, criteria #4 (97% coverage of the 2-sigma ellipse by the CSZ) had not yet been established by the project. There are two ellipses given in Fig. 15(b). The larger magenta colored ellipse represents the 99% landing ellipse which includes all errors modeled in the Monte Carlo and corresponds to the final OD solution before the maneuver was executed. Note that the center of this ellipse (magenta plus sign) is about 1-2 km to the northwest of the analogous nominal solution (magenta circle). This is due to the fact that the nominal case was consistently separated from the mean of the 2000 sample Monte Carlo. This magenta ellipse covers the region within which the probability of landing is 99% if no further maneuvers are performed. The center of the magenta ellipse, however, will move as new knowledge is obtained and nominal OD solution converges to the actual solution. The narrow, white ellipse is the 99% ellipse that includes navigation errors only. In a strict sense, this means that future OD nominal solutions (the small circles) could be anywhere in that narrow ellipse. That is, if no maneuver were performed, the nominal OD solution (and the center of the magenta ellipse) could migrate anywhere within that 99% ellipse over the course of the week leading up to landing.

The first solution shown on the plot, May 10th in blue, falls within the pink contour, but not the gold one. This means that criteria #3 was met, but #1 was not. Criteria #2 – 99% of the 1-sigma ellipse by safe slope regions - was met by each of the solutions presented in Fig. 15. As illustrated in the figure, the day-to-day movement in the solution was along a line that is almost congruent with the southwestern border of the gold contour. This means that on any given day, there was a chance of meeting all 3 of the original criteria set. As it turns out, on about May 14th the solution crossed a border where the probability calculations began to meet the criteria. Table 7 lists the calculated probabilities for the solutions corresponding to Fig. 15. While the criteria were technically met, the project was very concerned about the black and red regions in the southern part of the ellipse. This was terrain that had not been imaged *and* was in a statistically rocky geomorphic unit (black), or had been imaged and been classified as having significant rock abundance (red). The final decision to go ahead with TCM-5 was based partially on a consensus that the current situation did not meet the 'spirit' of the criteria, which was to put the ellipse

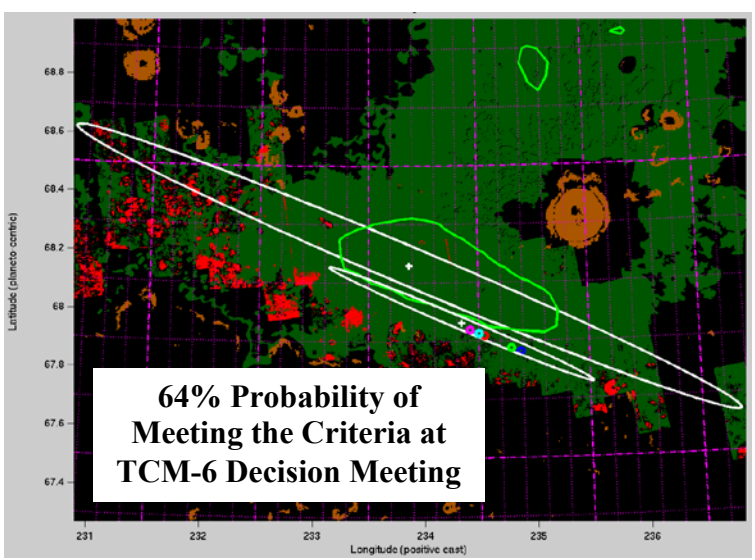

Figure 16. Nav Error Only Ellipses for the May 17th OD Solution

in a region which was as benign as possible with respect to hazards. At this point in time, in an attempt to better meet the spirit of the criteria, the project decided to add criterion #4 which was related to the 2-sigma ellipse.

Part of the decision to go forward with the maneuver was based on confidence in the spacecraft and navigation team's ability to correct the landing site error which was causing us to violate the new criterion. Most of the correction in landing site required to move the ellipse away from the red/black regions was in the cross-track

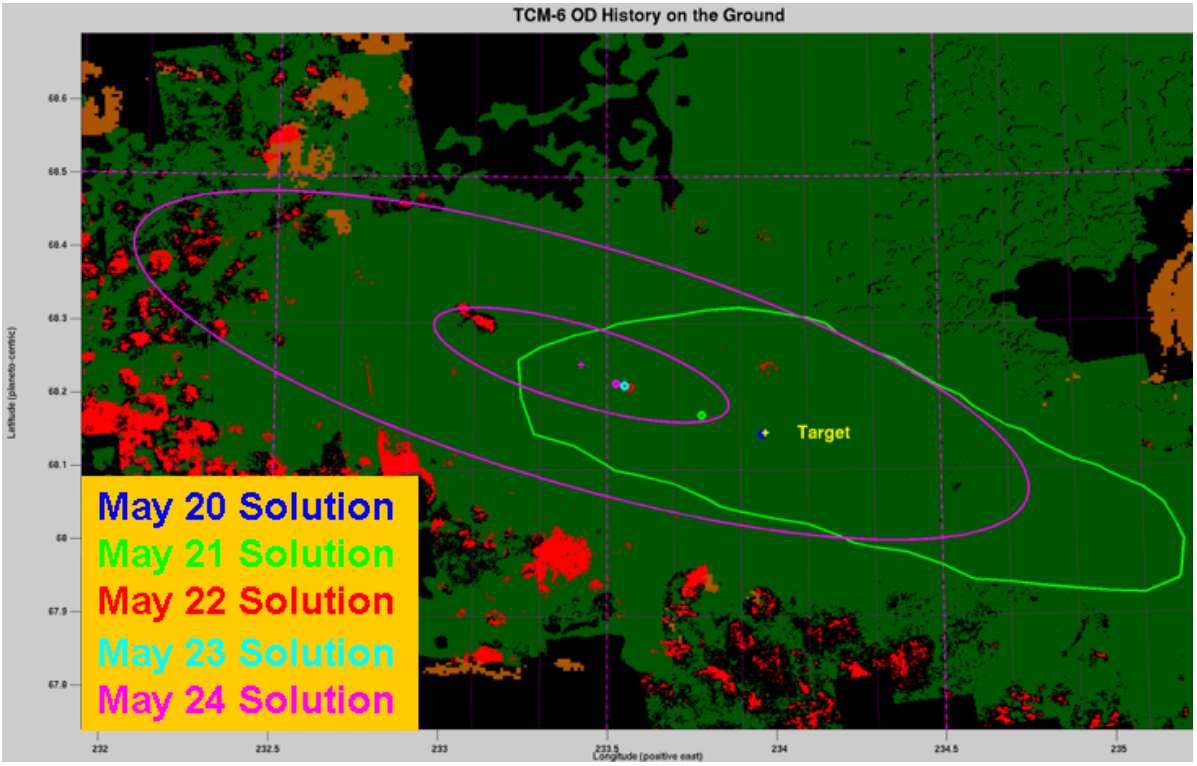

Figure 17. TCM-6 OD History on the Ground

direction. Cross-track is the direction which is perpendicular to the long axis of the ellipse. Considering that the cross-track error was on the order of 7 km and the semi-minor axis of the navigation error ellipse was only 1 km, the knowledge required to correct the error was great enough to justify performing the maneuver from a navigation perspective. That is, enough knowledge existed to say that if TCM-5 was executed, the probability of the cross-track error after the maneuver being worse than 7 km was infinitesimal (barring a spacecraft anomaly).

Additionally, through the use of the criteria contours and the navigation-only ellipse in Fig. 16, it was possible to quantify the probability of meeting the criteria at the TCM-6 Go/No-Go meeting if TCM-5 were performed. A navigation-only ellipse was calculated based on the knowledge at the time of TCM-5 and an estimate of TCM-5 execution errors. This ellipse is pictured as the large white ellipse in Fig. 16. The small white ellipse is the same as shown in Fig. 15 and is included to give an idea of how much uncertainty is introduced due *only* to the maneuver execution errors. The large white ellipse represents the 99%-confidence location of the nominal OD solution (and of the center of the ellipse with EDL errors) at the time of the TCM-6 decision meeting (about 1 day prior to entry, after execution of TCM-5 and subsequent collection of additional tracking data). The probability that the large white ellipse lies within the green contour representing criteria #4 (which turned out to be the strictest of all the criteria), is the probability is that all of the criteria will be met at the time of the TCM-6 decision. The assessment that the available knowledge favored performing TCM-5 and that the subsequent probability of not needing TCM-6 was at least a 60% gave additional credibility to the decision to go ahead with TCM-5.

B. TCM-6, TCM-6X and 'The Hill'

TCM-5 was performed on May 18th and was, yet again, another near flawless execution by the spacecraft team. Of course, no maneuver can be perfect and the project once again found itself in the position of meeting criteria while being in an uncomfortable situation with respect to hazards. The OD solutions following TCM-5 and leading up to the TCM-6 decision are mapped to the ground in Fig. 17. The first solution in this figure, from May 20th, is

Table 8. Monte Carlo Probability Results for OD Solutions Leading Up to TCM-6

Primary Decision Factor	Criterion	May 20th Solution	May 21st Solution	May 22nd Solution	May 23rd Solution	May 24th Solution	Final Solution
Percent of 1σ Landing Ellipse in CSZ (By Area)	95%	99.9%	99.8%	97.3%	97.2%	97.1%	96.1%
Percent of 2σ Landing Ellipse in CSZ (By Area)	97%	99.0%	99.8 %	99.0%	99.0%	99.0%	98.9%
Percent of 1σ Landing Ellipse in safe MOLA slope region (By Area)	99%	99.9%	99.1%	99.9%	99.9%	100.0%	99.6%
Probability of landing in a Certified Safe Zone	80%	99%	98.9%	98.1%	98.1%	98.0%	97.3%
Probability of landing on 'The Hill'	None	N/A	N/A	N/A	1.1%	1.2%	1.6%
Probability of landing in a HiRISE image	None	99.7%	99.6%	99.4%	99.4%	99.4%	99.3%

nearly on top of the target. The nominal point continued to move 'up-track' as the solution converged following the TCM-5 event. The ellipses shown in the figure are the 1-sigma and 99% ellipses for the final solution before the TCM-6 decision. While the new criteria were met at every point leading up to that decision, as seen in Table 8, something was occurring on the ground that was grabbing the project's interest. The solution was continuing to move towards what became known as 'The Hill' – the red hazard in the northwestern portion of the 1-sigma ellipse. Eventually, the project even started tracking the probability of landing on the hill, reaching a peak probability of 1.6% in the *hours* prior to entry.

All of this amounted to a difficult decision that had to be made – with the hill firmly entrenched in the 1-sigma ellipse, should the spacecraft perform TCM-6? It was decided that if TCM-6 was executed, the ground target would be moved to the center of the green contour (the criteria #4 contour). The original target (seen in Fig. 17) was not exactly centered for reasons mentioned earlier. At first glance, the probability numbers are so high that there seemed to be little need to worry about the hill. In fact, there was, at worst, a 20% probability of encountering a dangerous rock if Phoenix landed on the hill - meaning that the probability of landing *on* the hill *and* encountering a dangerous rock was approximately 0.3%. However, based on an analysis similar to that performed prior to TCM-5, the landing site safety team determined that there was a 99% chance of meeting the criteria if TCM-6 was retargeted. Essentially this meant that if the maneuver was performed, the navigation team and landing site safety team could almost guarantee that 'The Hill' would no longer be an issue for the project. Ultimately, the project decided that the risk of landing on any rocks was so low that they could not justify taking the risk of attempting a

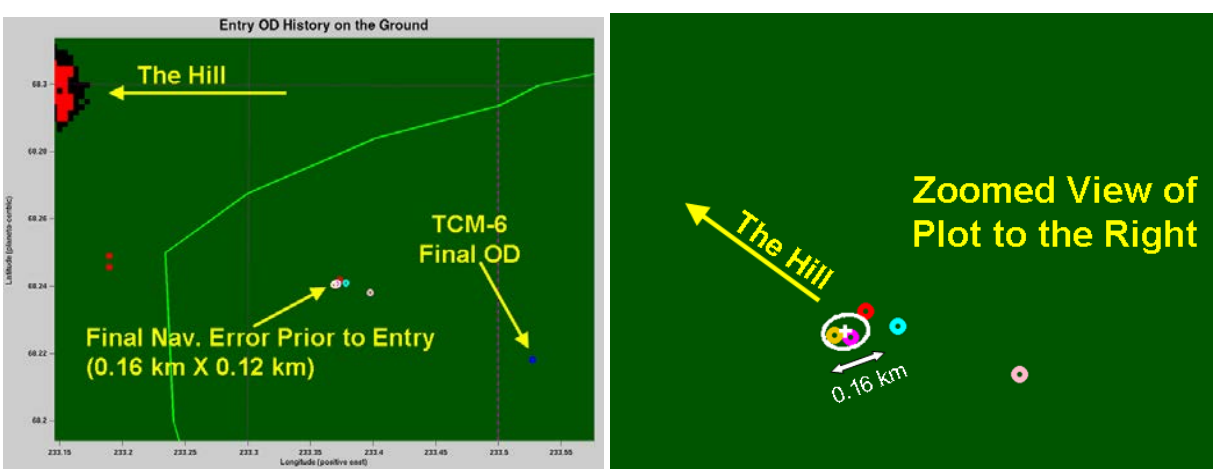

Figure 18. OD Solution Convergence From TCM-6 to Entry

maneuver and having something go wrong with the spacecraft, and TCM-6 was cancelled.

At this point, TCM-6XM was still a possibility, and the situation with respect to the hill did not get any easier. The orbit determination solution continued to move closer to the hill as seen in Fig. 18. In total, the solution moved another 3-4 km towards the hill as new OD data came in, increasing the probability of landing on the Hill to the

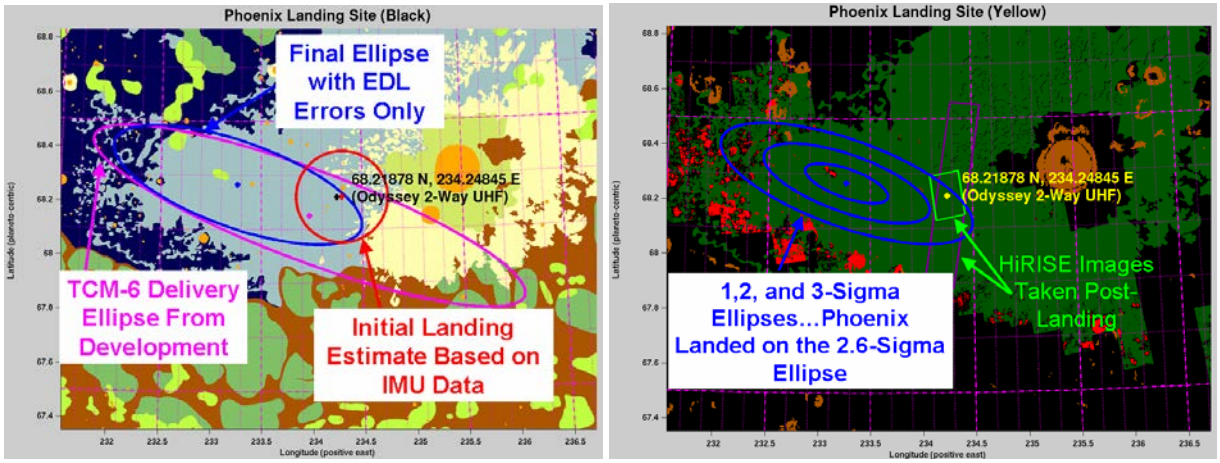

Figure 19. The Actual Phoenix Landing Site

1.6% mentioned above. Because overall risk had not changed much from the time of the TCM-6 decision, TCM-6XM was cancelled as well and the project prepared for landing. By the time the lander was ready to enter the atmosphere, the OD team had determined the entry state to such accuracy that the navigation errors mapped to the ground only accounted for 0.16 km. Any miss on the ground from the nominal entry solution would be due to EDL uncertainty (including atmosphere modeling, vehicle aerodynamics, etc.) as the entry state knowledge was nearly perfect at entry.

C. Final Entry Ellipse and Landing Position

On May 25th, 2008 at 23:32:09 UTC, the Mars Phoenix Lander entered the atmosphere of Mars and descended to the ground for a near perfect landing. The lander now resides at 68.219 degrees north-latitude (areocentric) and 234.248 degrees east-longitude. This landing site is shown in two different pictures in Fig. 19. The first picture gives the development phase TCM-6 delivery ellipse (magenta), the landing ellipse – with EDL errors – for the final OD solution (blue), and the spacecraft team's initial estimate of landed position based on IMU data (red circle). As it turns out, the IMU data was pretty accurate – the center of the red circle was within 1 km of the actual landing site. The other picture in this figure shows the 1,2 and 3-sigma landing ellipses and the HiRISE swaths of the images that were taken by MRO after landing.

Table 9. Relevant Phoenix Miss Distances

	Miss Distance in km	Miss Distance in 'Sigma'
Along-Track	~21.5 km	~2.3 Sigma
Cross-Track	~4 km	~1.3 Sigma
Total	~21.9 km	~2.6 Sigma
Miss From Target	7 km	Not Meaningful to Report

The lander ended up about 22 km "long" relative to the nominal OD solution. As mentioned earlier, this 22 km is entirely due to EDL uncertainties and not OD uncertainties. The location on the ground corresponds to a landing ellipse that is 2.6-sigma. This would be equivalent to the 97% ellipse. EDL reconstruction is currently ongoing to see if any possible explanations can be found in the 200Hz IMU data taken during EDL that Phoenix returned to Earth. Statistics for the miss on the ground are provided in Table 9. Phoenix landing only 7 km from the original intended target of TCM-5 was somewhat fortuitous. The miss from the intended target could easily have been 25-30 km had the TCM-5 execution left us long relative to the target rather than short of it.

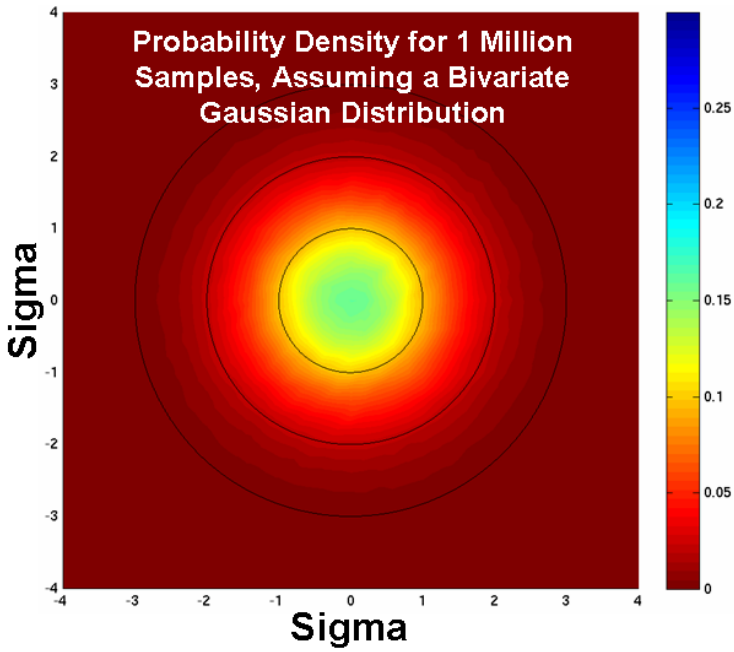

Figure 20. Example of the PDF using 1 Million Samples

V. Validity of the Bivariate Gaussian Distribution Assumption for Phoenix

While some of the Monte Carlo inputs had uniform distributions associated with them, the major contributors to landing dispersion (entry state, aerodynamic coefficients, and atmospheric density) assumed to have Gaussian distributions. Given this and the fact that a landing dispersion is a two-dimensional distribution, the ground dispersion is modeled as a bivariate Gaussian distribution. Using this model assumes that the 2000 landing points take on a Gaussian distribution in the major and minor axis of the ellipse used to represent the distribution. While this is a common assumption to make, it begs the question of whether or not it is a valid assumption – especially in light of the Phoenix lander landing on the 2.6-sigma ellipse.

In an attempt answer this question, the probability density was analyzed for a few specific cases. The first test was to create 1 million samples in MATLAB assuming a perfect bivariate Gaussian distribution. The probability density was calculated and plotted in Fig. 20. The purpose of this plot is to show what the probability density should look like in a perfect scenario where under-sampling is not an issue. The same analysis was performed for a 2000 sample case (again assuming a perfect distribution) in Fig. 21(a). The purpose of the 2000 sample case with a perfect distribution was to see how the probability density would look with under-sampling issues (as is the case in a 2000 run Monte Carlo). Clearly, this density plot is not nearly as continuous as the 1 million sample case. Fig. 21(b) is the probability density for a Phoenix Monte Carlo. Again, the density is not as continuous but, also of note, is that the center of the distribution clearly shows a dip in the density. This means that for Phoenix, there was a lower probability of landing in the 1-sigma ellipse than was predicted by the bivariate assumption.

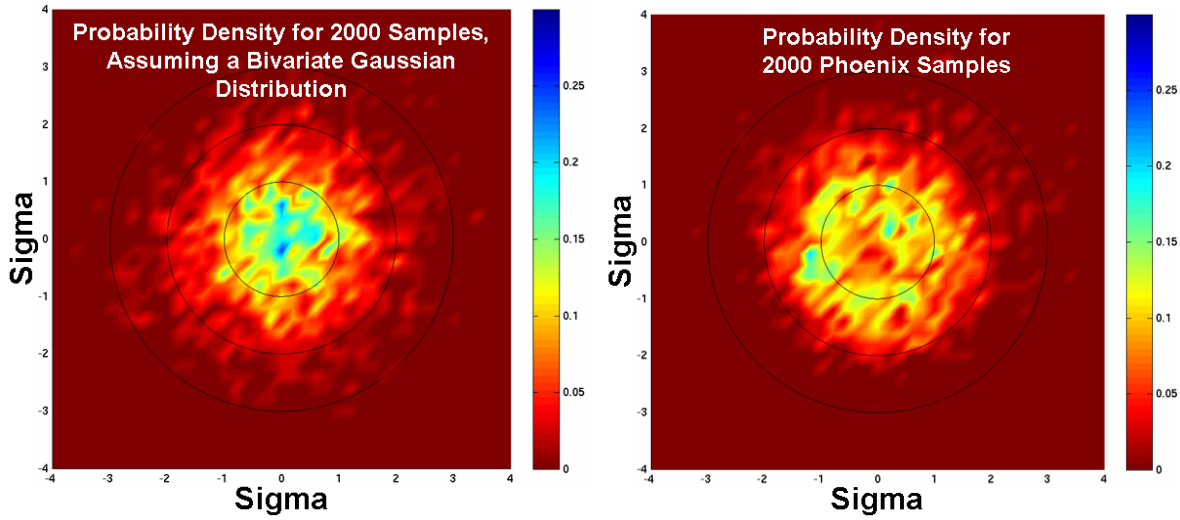

a) b)

Figure 21. Example Probability Density for a 2000 Samples from MATLAB (a) and Phoenix (b)

The figures above each contain the 1, 2, and 3-sigma ellipses (which appear as circles in these plots). A more quantitative measurement of how well each of them matches the actual probability density function for the bivariate Gaussian is the percentage of cases that fell within the circular borders. The first column in Table 10 gives the expected percentage of cases that fell within the 1-sigma ellipse, between the 1 & 2-sigma ellipse, and the 2 & 3-sigma ellipse. The second and third columns, which assume a perfect Gaussian distribution, agree to within expected levels (±2% for 2000 samples). The last column, which is the Phoenix distribution, is similar to the Gaussian distribution but definitely shows a bias towards the region between the 1 & 2-sigma ellipses. The Monte Carlo analysis predicts that nearly 60% of the cases fall within this region. While this analysis says that the Gaussian distribution is not exactly perfect, it is a fairly good approximation (and does not point towards a tendency to land outside of the 2-sigma ellipse, which is where Phoenix landed).

Table 10. Percentage of Cases Within the Ellipse Contours

	Bivariate Gaussian PDF (Expected Result)	1 Million Bivariate Gaussian Samples	2000 Bivariate Gaussian Samples	2000 Phoenix Samples
% Cases between 0 & 1-Sigma	39.3%	39.3%	40.2%	30.9%
% Cases between 1 & 2-Sigma	47.1%	47.2%	45.8%	59.3%
% Cases between 2 & 3-Sigma	13.5%	13.5%	14.1%	9.9%

IV. Conclusion

On May 25, 2008, the hard work and dedication of many engineers and scientists was realized when Phoenix made the first successful soft landing on Mars since the Viking landers. The landing site work was a constant battle to move Phoenix 'away from the rocks.' It involved many firsts for lander missions. The most notable being the effort to count individual rocks on the surface. Additionally, the landing site team was constantly trying to stay one step ahead by using the available data to understand the probability of future events. This was very helpful with putting critical decisions in context. Of course, there is no way to know what would have happened if different courses of action had been taken. But the landing site team had one driving thought in mind all along – the knowledge that if the EDL team could successfully navigate the spacecraft to the ground, one of the worst scenarios would be for a rock to cause a mission failure. When the first pictures to come back from the spacecraft confirmed what the landing site team knew to be true – the landing site selected for Phoenix was among the safest locations in the northern plains of Mars.

Acknowledgments

Putting Phoenix down in an area on Mars where absolutely no rocks existed was the triumph of many more people than the authors of this paper. The work performed by Matt Golumbek, Tim Parker, Yang Cheng and Andres Huertas of the Jet Propulsion Laboratory was instrumental in setting what will be the gold standard for landing site safety analysis. The team at Lockheed Martin in Littleton Colorado was a tremendous help in this regard as well. Not only did they build an amazing spacecraft, they helped determine what the true risks were with respect to landing on the rocks. The research described in this publication was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[1] Ryne, M.S., Graat, E.J., Lau E., et al. "Orbit Determination for the 2007 Mars Phoenix Lander," Proceedings of the *AIAA Guidance, Navigation, and Control Conference*, AIAA-2008-7215, Honolulu, HI, Aug. 2008.

[2] Raofi, B., Bhat R. Helfrich C., et al. "Propulsive Maneuver Design for the 2007 Mars Phoenix Lander Mission," Proceedings of the *AIAA Guidance, Navigation, and Control Conference*, AIAA-2008-7216, Honolulu, HI, Aug. 2008.

[3] Portock, B.M., et al., "Navigation Challenges of the Mars Phoenix Lander Mission," Proceedings of the *AIAA Guidance, Navigation, and Control Conference*, AIAA-2008-7214, Honolulu, HI, Aug. 2008.

4 Arvidson, R. E., D. Adams, G. Bonfiglio, P. Christensen, S. Cull, M. Golombek, J Guinn, E. Guinness, T. Heet, R. Kirk, A. Knudson, M. Malin, M. Mellon, A. McEwen, A. Muskin, T. Parker, F. Seelos IV, K. Seelos, P. Smith, D. Spencer, T. Stein, L.

Tamppari, 2008, "Mars Exploration Program 2008, Phoenix Landing Site Selection and Characteristics", *J. Geophys. Res.*, 113, E00A03 doi: 10.10292077JE003021.

[5] Seelos, K. D., R. Arvidson, S. Cull, C. Hash, T. Heet, E. Guinness, P. McGuire, R. Morris, S. Murchie, T. Parker, T. Roush, F. Seelos, M. Wolff, 2008, Geomorphologic and Mineralogic Characterization of the Northern Plains of Mars at the Phoenix Mission Candidate Landing Sites, *J. Geophys. Res*., in press

[6] Golombek, M. P., A. Huertas, J. Marlow, B. McGrane, C. Klein, M. Martinez, R. E. Arvidson, T. Heet, L. Barry, K. Seelos, D. S. Adams, W. Li, J. R. Matijevic, T. Parker, H. Sizemore, M. T. Mellon, A. S. McEwen, L. Tamppari, Y. Cheng 2008, "Size-Frequency Distributions of Rocks on the Northern Plains of Mars with Special Reference to Phoenix Landing Surfaces", *J. Geophys. Res.*, 113, E00A09 doi: 10.1029/2007JE003065.

[7] Prince, J, Queen, E. and Desai, P, "Mars Phoenix Entry, Descent, and Landing Simulation and Modeling Analysis," Proceedings of the *AIAA Guidance, Navigation, and Control Conference*, AIAA-2008-7507, Honolulu, HI, Aug. 2008.

[8] Prince, J, Queen, E., Desai, P, and Grover, M.R., "Entry, Descent, and Landing Operations Analysis for the Mars Phoenix Lander," Proceedings of the *AIAA Guidance, Navigation, and Control Conference*, AIAA-2008-7349, Honolulu, HI, Aug. 2008.

[9] Queen, E, Prince, J. and Desai, P, "Multibody Modeling and Simulation for Mars Phoenix Lander Entry, Descent and Landing," Proceedings of the *AIAA Guidance, Navigation, and Control Conference*, AIAA-2008-7347, Honolulu, HI, Aug. 2008.

[10] Desai, P., Queen, E., Prince, J. and Grover, M.R., "Entry, Descent, and Landing Performance of the Mars Phoenix Lander," Proceedings of the *AIAA Guidance, Navigation, and Control Conference*, AIAA-2008-7346, Honolulu, HI, Aug. 2008.

[11] Grover, M.R. and Desai, P., "Overview of the Phoenix Entry, Descent and Landing System Architecture," Proceedings of the *AIAA Guidance, Navigation, and Control Conference*, AIAA-2008-7218, Honolulu, HI, Aug. 2008.

[12] Tamparri, L. et. Al, 2008, "The Atmospheric Environment Expected for the Phoenix Landed Season and Location", *J. Geophys. Res.*, doi:10.1029/2007JE003034, in press.

[13] Knocke, P. C., Wawrzyniak, G. G., Kennedy, B. M., Parker, T. J., Golombek, M. P., Duxbury, T. C., Kass, D. M., and Desai, P. N., "Mars Exploration Rovers Landing Dispersion Analysis," *AIAA/AAS Astrodynamics Specialist Conference*, AIAA-2004-5093, Providence, AIAA, Washington, DC, August 2004