PERISCOPE: PERIapsis Subsurface Cave OPtical Explorer

NASA Innovative Advanced Concepts Program

Program Manager Jay Falker

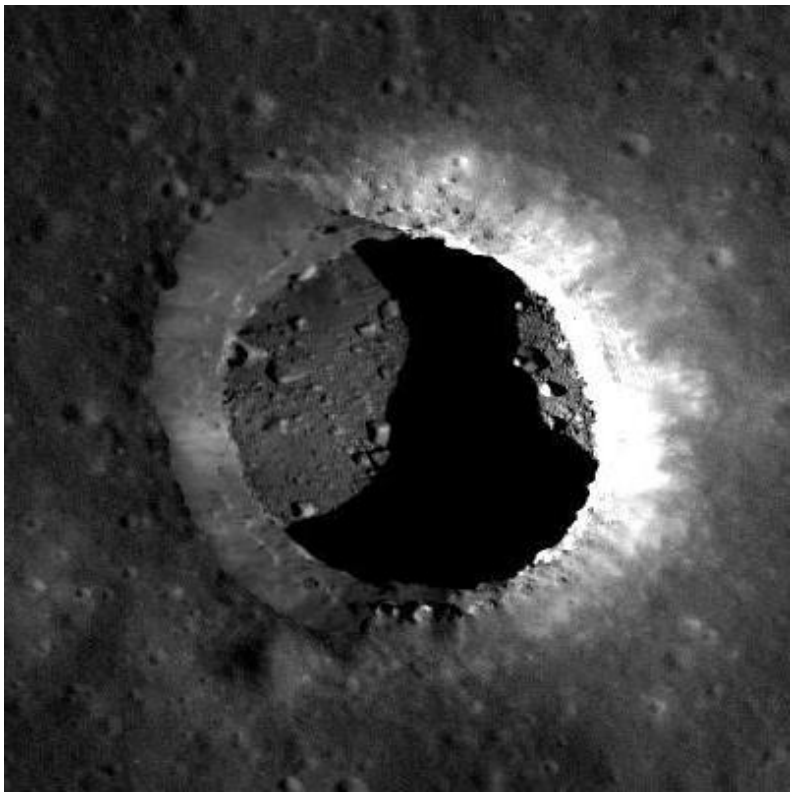

NIAC Phase 1 Final Report

Nosanov Consulting LLC, University of Wisconsin, NASA Jet Propulsion Laboratory

Jeffrey Nosanov, Andreas Velten (U of W), Karl Mitchell and Nitin Arora (NASA JPL)

This report is dedicated to Penelope Jean Nosanov, born February 15th, 2014. May the mysteries of the moon inspire her and future generations.

Table of Contents

Background

The possibility of lunar caves has fascinated people in both the scientific community and popular culture since before the start of the space age[i],[ii]. Their presence was inferred based on Apollo data[iii], but it's only in the last few years that glimpses into these caves, via skylights, has made them a truly compelling target for exploration[iv]. This was perhaps not well captured by the most recent Planetary Science Decadal Survey [v], in part due to the limited time in which the community was able to consider the recent direct discovery. However, scientists and engineers are now starting to develop the reasons and technologies for planetary exploration, and in-depth contributions are anticipated for the next Decadal Survey.

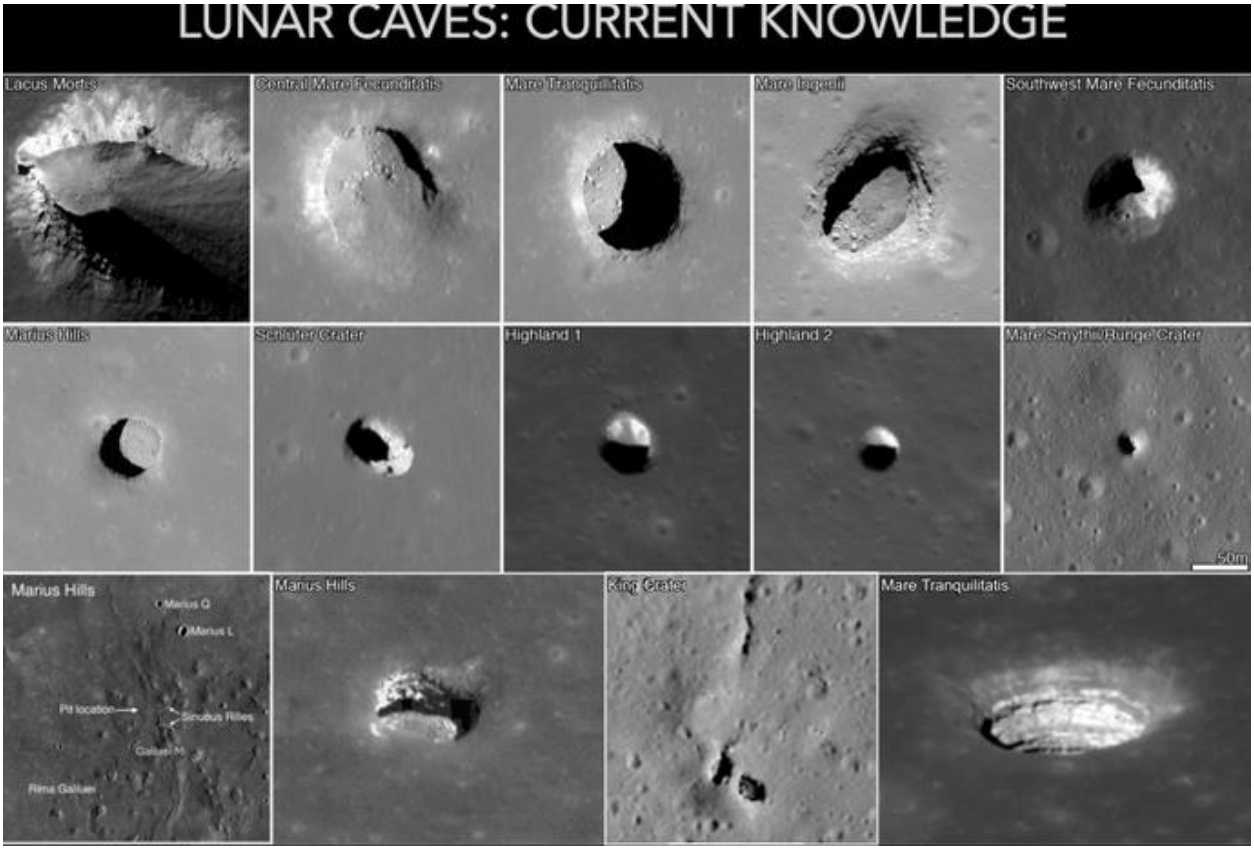

Figure 1: Pits/skylights observed on the Moon. (top) Ten example pits to scale (upper and middle rows) as well as other illustrative observations (bottom) including oblique views showing structure in the skylight walls.

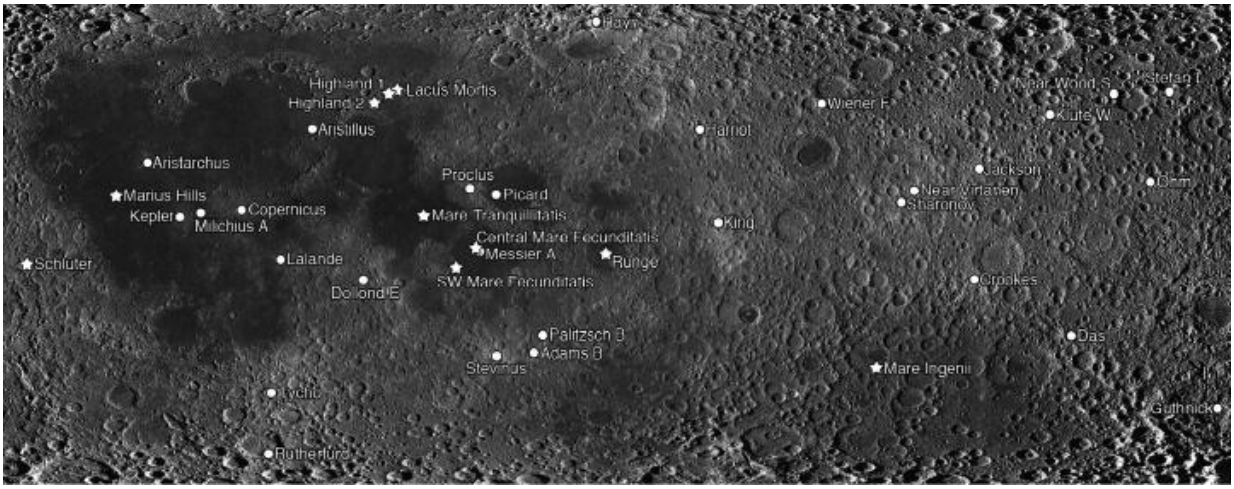

Figure 2: Distribution of lunar pits shown on simple cylindrical projection, from Robinson et al. (2014)[vi], including clusters of impact crater melt pits (totaling 221) indicated by circles, and the additional eight mare and two highland pits indicated by stars.

Photon Time-of-Flight Imaging

In this NIAC report we present a new technique enabling a practical option for mapping these structures at a relatively low cost: Photon Time-of-Flight (PTOF) imaging. Consider a scene with a surface that is in a camera's line of sight, with unknown geometry beyond the line of sight. PTOF works by directing laser pulses onto the visible surface and detecting the returned light after it reflects off the visible surface, onto the hidden surfaces, and back to the detector. An algorithm then reconstructs a 3d model of the hidden surface. The information gained has a wide variety of applications across NASA, the scientific community, the burgeoning commercial space industry, and in commercial and industrial use here on Earth.

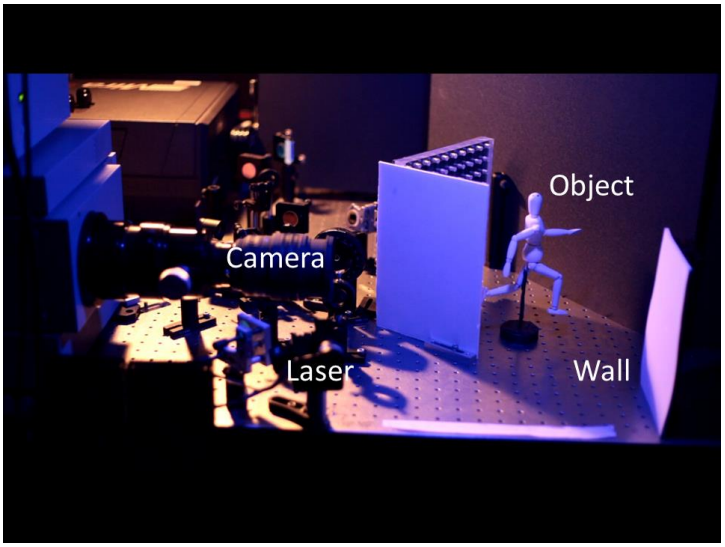

Figure 3: Dr. Velten's initial experiment, shown above, inspired this mission concept.

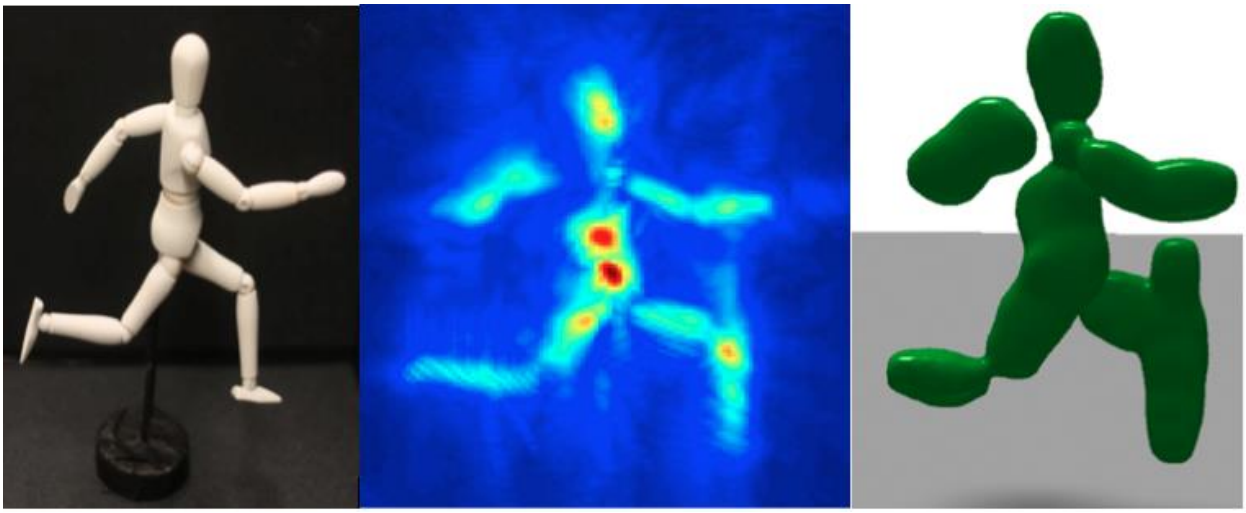

Figure 4: The setup includes a wooden artist's mannequin (left) positioned beyond the line of sight of the camera behind an obstacle. The laser shines on the surface labeled "wall", reflects in all directions including onto the mannequin. Some of the light incident upon the mannequin reflects back to the wall, and then back to the camera (middle). An algorithm then reconstructs a 3d model (right) of the mannequin based on the time of flight of the small number of photons that completed all 3 bounces and reached the detector.

Current Status of Global Lunar Ambitions

It seems likely that humanity will ultimately develop a permanent presence on the Moon, and lunar caves provide a uniquely protected environment for this. As of early 2015 NASA has no concrete, funded plans to return to the moon in a long term, sustainable way. However multiple robotic missions to the moon have been extremely successful (GRAIL, LRO) and several more remain on the Planetary Science Decadal Survey (Lunar Geophysical Network, Lunar South Pole Aitken Basin.) China may have such plans but they are not public and NASA is not permitted to work with the Chinese space agency. However, the moon remains by far the nearest astronomical body to the Earth, the easiest to reach, and will likely be a major part of the overall worldwide space exploration effort for the coming centuries.

Human Explorations and Operations Mission Directorate (HEOMD) Applications

Human beings will eventually have a continuous or permanent presence on the moon. Caves provide multiple benefits for future crewed missions: protection from the extreme temperature swings on the surface, the possibility of resources, and simplified local habitat construction using "ready-made" structure. PERISCOPE offers HEOMD a way to investigate both questions with one low-cost mission. Photon time of flight imaging can map the interior structures of lunar skylights and potentially characterize elemental distribution within the caves. Knowledge of ice distribution in various caves would be valuable for future crewed visits as both a possible source of consumables as well as fuel.

Science Mission Directorate (SMD) Applications

The PERISCOPE concept can contribute significantly to the NASA SMD long-term goals of understanding our solar system and the bodies within it. The Earth's Moon is a particularly valuable target, because it preserves a record of exogenic processes in the Earth-Moon system, and thus the history of impacts (which wreaked havoc on early Earth.) Volatile deposition and space

weather on the Moon has direct relevance to that on Earth. Due to its proximity, it is also a cost-effective target, and a likely staging post for human exploration of the Solar System.

Three lunar missions featured strongly in the Planetary Science Decadal Survey: the Lunar Polar Volatiles Explorer (a rover for studying polar ice deposits), Lunar Geophysical Network (multiple long-term landed ground stations) and the Lunar South Pole-Aitken Basin Sample Return (a mid-sized sample return mission). The goals of the latter of these were to:

- Determine the chronology of basin-forming impacts and constrain the period of late heavy bombardment in the inner solar system and thus address fundamental questions of inner solar system impact processes and chronology;
- Elucidate the nature of the Moon's lower crust and mantle by direct measurements of its composition and of sample ages;
- Characterize a large lunar impact basin through "ground truth" validation of global, regional, and local remotely sensed data of the sampled site;
- Elucidate the sources of thorium and other heat-producing elements in order to understand lunar differentiation and thermal evolution; and
- Determine ages and compositions of farside basalts to determine how mantle source regions on the far side of the Moon differ from regions sampled by Apollo and Luna.

Lunar pits and caves provide a convenient, ready-made entrance to the subsurface, allowing deep sampling of basins and impact melts without drilling. By studying them we can gain insights into basin forming impacts and related processes. Additionally, they are an environment protected from space weather and micro-impacts, and therefore provide a more pristine record of volatile deposition. The PERISCOPE concept would be an important pre-cursor that would improve target selection and provide critical constraints for enabling in situ robotic exploration.

The dark materials that fill the mare basins are basaltic rock that is the direct result of massive flood eruptions comparable to some pre-historic but geologically recent (possibly a few Ma) events on the Earth. A causal relationship between the lunar mare volcanism and the major impacts has long been suspected, due to the high degree of spatial correlation at least on the near side of the moon. However, the timing and extended duration of lunar volcanism from crater counting, which seems to begin ~100 Ma following basin formation, casts doubt on the possibility of very directly impact-triggered eruptions[vii]. Importantly, oblique observations of lunar pits/skylights appear to reveal stratified layers, even in smaller non-mare impact melt sheets, implying a more complex history of multiple emplacement events[viii,ix]. Thus, not only would the dating and sampling of materials in pits, skylights and caves allow a more detailed chronology of major impact and post-impact events to be revealed, a substantial enhancement to the goals stated above, but also it would give insight into a range of potentially catastrophic processes that have affected other worlds in the past, including Earth. In a broader sense, study of lunar pits/skylights in impact basins contributes directly to understanding of the roles that bombardment by large projectiles play in the formation of new worlds [v pp. 19] as well as how planetary surfaces are modified by geologic processes [v ch. 5].

Another reason to go to lunar caves is to understand the evolution of volatiles, which may be uniquely protected in cave environments. The lunar surface and interior are now known to be

less dry than thought, 10-1000 ppm globally[x],[xi], and a growing body of evidence suggests vast concentrations of water-ice present in polar permanently shadowed cold craters. Although exogenic water may be the cause of much of this, endogenic water may also be of importance. Caves provide a ready-made path into the subsurface to permanently shadowed sites. They are, however, considerably warmer (up to ~-20º C) than permanently shadowed polar craters where ice deposits have been observed, meaning that thermal sublimation effects will be greater[xii]; Thus, water is unlikely to have survived within open cave systems over geological timescales, but if the water were supplied by crustal diffusion into modern times, or if caves were closed or mostly-closed, then water may persist. Thus, the detection or lack thereof of water in caves would provide an important constraint on lunar volatile history. Furthermore, other volatile compounds as well as hydrated minerals may be more stable. PERISCOPE is uniquely suited to detecting albedo anomalies within shadowed caves, and can be used to direct future missions to putative ice deposits. Study of volatiles in lunar caves would complement other studies of polar volatiles such as the proposed Lunar Polar Volatiles Explorer concept mission, providing more complete insight into their presence and genesis.

Taken together, these studies contribute to one of the key questions of the Decadal Survey[v]: What governed the accretion, supply of water, chemistry, and internal differentiation of the inner planets and the evolution of their atmospheres, and what roles did bombardment by large projectiles play?

Planetary Use

Using orbital photon time of flight imaging at a planet with an atmosphere thicker than Earth's would introduce new challenges. These including the atmospheric scattering effect on the laser pulses and the higher orbit and larger standoff imaging distance required. This makes the moon a compelling target for this technique, as it has no atmosphere. However the technique can perform well through relatively diffuse atmosphere or at low altitude through significant atmosphere. The technique may be applicable on Mars from an airborne platform. A thorough analysis of using the PERISCOPE system through an atmosphere was beyond the scope of our Phase 1 study but we extrapolate that an airborne system on Mars at a very low altitude (~thousands of feet) could potentially work. This is a powerful avenue for future work.

SMD and HEOMD overlap

Any crewed return to the moon is likely to require significant partnership between HEOMD and SMD, much as it occurred during the Apollo program. Indeed, there are two specific, impactful statements concerning this in the Planetary Science Decadal Survey[v] (pp 62 and 63):

"Scientifically useful investigations should still be developed to augment human missions to the Moon or NEOs. The committee urges the human exploration program to examine this decadal survey and identify—in close coordination and negotiation with the SMD—objectives whereby human-tended science can advance fundamental knowledge. Finding and collecting the most scientifically valuable samples for return to Earth may become, as they were in the Apollo program, the most important functions of a human explorer on the Moon or an asteroid" and "The robotic and human exploration of space should be synergistic, both at the program level (e.g., science probes to Mars and humans to Mars) and at the operational level (e.g., humans with robotic assistants)."

Commercial Space Applications

A growing number of companies are pursuing commercially - or privately - funded lunar missions (Golden Spike, Google Lunar X-prize, Bigelow Aerospace). It is conceivable that a future

commercial company could mine and sell water resources to NASA or another agency, or that other elements found within these structures could have commercial value. Water found in these structures could be sold to a space agency maintaining a permanent base. PERISCOPE could also be used to map caves and voids (natural and artificial) on Earth. PERISCOPE may have commercial appeal.

Industrial Applications

A growing number of industries rely on systems and machines performing work that would extremely dangerous to human beings. Mining, refining, power generation, large-scale manufacturing, construction, are just a few industries in which it would be useful to look around corners and into voids and spaces from as far a distance as practical. Photon time-of-flight imaging can greatly extend this "minimum practical distance" by enabling the operator to be "around the corner" from the dangerous condition. The PERISCOPE team is very interested in exploring this avenue in the future.

Military Applications

The photon time-of-flight imaging technique that enables the PERISCOPE lunar mission has a wide variety of military applications. Strategic uses include mapping caves from aircraft, identifying whether a suspected structure is a cave or not, and getting information about what is underneath various structures by shining the laser at the adjacent ground. Hangars, storage areas, and other industrial spaces could be probed in this way. Tactical uses include individuals looking around corners in buildings or other structures to identify threats or access routes before the operator is "face to face" with the threat, as in Figure 5.

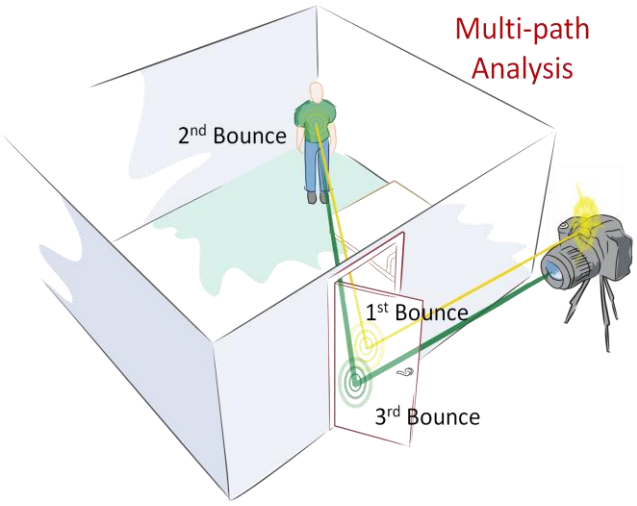

Figure 5: PTOF imaging could allow a soldier or special forces unit to identify possible threats inside a room from a distance and around a corner; far from the field of view of any individual in the room.

Current and Proposed Methods of Planetary Cave Exploration

Human

The most recent program to aspire to human lunar landings was Constellation. The program was cancelled by the Obama administration, which found its total cost to reach $150B over its lifetime[xii]. As of March 2015 NASA's current human exploration program follows the Augustine Commission's "Flexible Path to Mars" and does not include a human presence on the moon.

Unique knowledge of lunar geography, geology and resource abundance will play a strategic role in 21st century geopolitics. China appears to be pursuing the capability for crewed lunar

landings and has made several steps towards that goal including crewed missions, docking, advanced maneuvers, lunar orbiters returned to Earth.

PERISCOPE can provide NASA and other United States agencies with privileged knowledge about the interiors of lunar skylights. This knowledge will be useful to prepare for the geopolitical challenges of the 21st century.

Small Robots

Recent advances in microspine technology done at the Jet Propulsion Laboratory have given small mobile robots the ability to climb and drive on varied natural rock surfaces. Robots, such as LEMUR IIb, have been demonstrate and tested climbing vertical and inverted horizontal rock surfaces in earth gravity, a harder than low gravity case.

While microspine cave exploration robots are able to cross varied and difficult terrains, they still suffer from the power limitations and communications challenges of mobile robotics[xiii]. Solar panels work well for mobile robots, but will likely be occluded for cave and lava tube exploration. Larger batteries can be used to increase the run time of small climbing robots, but this has the detrimental effect of increasing the climbing loads on the grippers.

Large Robots

A tethered two-robot system such as duAXEL or ATHLETE can allow a topside rover to act as a power generating system and send tethered power to the underground explorer. Tethers can also lower robots past difficult obstacles such as overhangs or to save robots in the case of a slip.

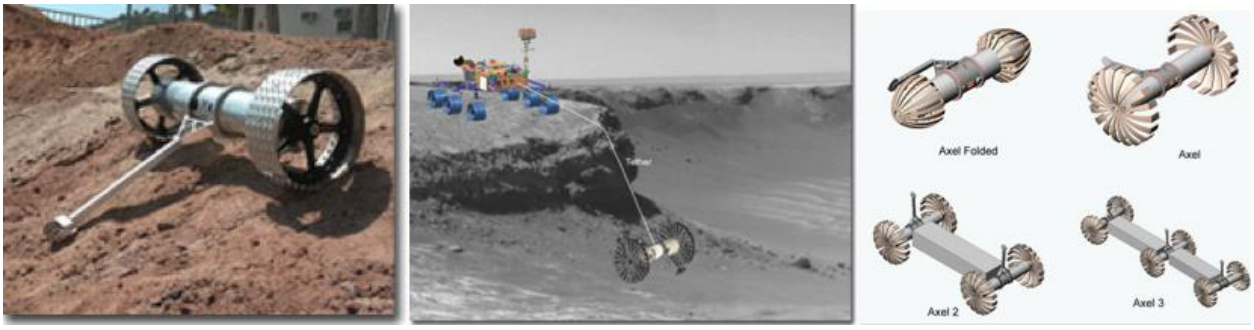

Figure 6: A variety of AXEL rover concepts from JPL. AXEL is a rover including scientific instruments built entirely within and around a single axle (left). This axel can serve as the rear axel or towed payload of a more traditional rover, and can be lowered over an edge (middle and right.)

The All Terrain Hex Legged Extra-Terrestrial Explorer (ATHLETE) is a six-legged vehicle commissioned during the Constellation program for lunar surface operations that could also deploy a smaller rover over an edge.

PERISCOPE can serve as a scouting mission for future rovers by characterizing the structures with a much lower cost orbiting mission.

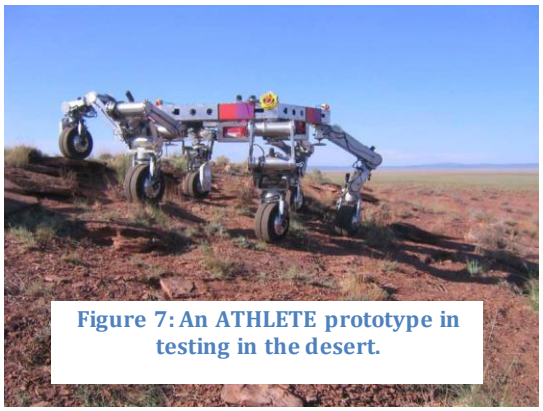

Figure 7: An ATHLETE prototype in testing in the desert.

Radar Sounding

One possible orbital solution to cave mapping comes from Radar Sounding. The 15-25 MHz Shallow Radar (SHARAD) on Mars Reconnaissance Orbiter (MRO) is capable of resolving features of ~15 m scale at up to 1 km depth[xiv]; This would almost certainly give some spatial information about caves associated with skylights on the Moon. However, this technology has limitations. Sub-wavelength features are not observed, and depth penetration is limited to at most ~50 wavelengths, and probably less in lava-tube-rich environments due to the high likelihood of voids, at wavelength scales and greater, causing clutter. Thus, the technology must be tuned for the scales involved, and is unlikely to give any structural information at scales of better than a few meters, even if wavelengths were ideally tuned.

We propose, therefore, that radar sounding can be complementary to the PERISCOPE solution for mapping lunar lava tubes, but that it is inferior in terms of designing robotic access and mobility solutions that require knowledge of wider range of structural and roughness scales down to centimeters.

Gravitational Mapping

A few groups have begun to explore the possibility of using GRAIL gravity map data of the moon to detect the largest subsurface voids such as caves[xv]. This promising but limited technique would be a useful asset to the PERISCOPE mission in order to compare our results to their predictions.

Introduction to Photo Time-of-Flight Imaging

The steps required to reconstruct images and 3d models of hidden scenes from orbit with a time of flight capture system are as follows. There are critical trade-offs and limitations in each of the described steps.

1. **Visible geometry acquisition** (mapping the bottom of the skylight as directly visible from orbit), probably with standard LiDAR techniques;

2. **Illumination** of a spot in the visible part of the target scene with a train of pulses;

3. **Collection** of returned *third* bounce light off the visible surfaces in the target scene (the second bounce occurs deep within the structure, beyond the line of sight from orbit);

4. **Iteration** of steps 2 and 3 for different illuminated spots or patterns in the scene; and

5. Computational **reconstruction** of the hidden geometry from the collected data.

Steps: (1) - **Visible geometry acquisition**: Knowledge of the 3D geometry of the *visible* part of the scene is required for the reconstruction of the *hidden* geometry. It can be collected by the same hardware used for the hidden geometry data collection or by a different independent spacecraft. This data can be collected on a separate passage over the cave before or after the actual hidden geometry detection. It can also be provided from a separate mission. The time resolution and sensitivity requirements are much lower when acquiring visible geometry.

Step (2) – **Illumination**: A laser pulse train must be focused on a small spot on a visible relay surface in the target scene. Our simulations suggest that an average laser output power of 1 kW is appropriate for the range of materials (unweathered basaltic rock) expected inside lunar caves, based on expected 2-4% reflectance[xvi]. Considering a range of laser efficiencies of between ~30%

(available today) and 50% (available within a decade), the power generation of the Lunar Reconnaissance Orbiter mission (order of magnitude ~1850 W) suggests that powering the necessary laser will be achievable within the NIAC time horizon of 10+years. The power is only required during the short illumination period of less than one second.

Step (3) – **Collection**: This step requires the collection of light from the visible relay surface in the target scene, scattered back to the spacecraft after multiple bounces within the cave, while maintaining a good registration of the visible scene geometry. The detected third-bounce light will be collected while the sensor is moving. Since the satellite position relative to the cave is changing rapidly, we will provide a way to track the satellite position and attitude in three dimensions using a self-calibration system that is integrated into the imaging process itself. We propose collecting data from individual illumination pulses, eliminating motion blur. The system will illuminate a grid of points on the surface and use two detection systems, a standard Single Photon Avalanche Detector (SPAD) LiDAR system operating on a small fraction of the returned power to detect the first bounce, and a gated avalanche photodiode detector that detects the third bounce from the same pulse. The first bounce provides depth information for a grid of points in the visible scene. Along with our knowledge of the visible scene geometry (acquired separately as described above) and approximate knowledge of the satellite rotation and acceleration, this allows us to co-register the visible and hidden geometries. The satellite motion during the time between the arrival of the first and third bounce is negligible for detection purposes due to the speed of light and the relatively short overall path length (~20-25 km.) This first bounce LiDAR data grid is also used for targeting. Using this self-calibrating data collection approach relaxes the requirements to satellite aiming and stability. Pointing accuracy can be kept low; one only needs to keep the cave in view as long as possible, in order to maximize the number of pulses into the cave.

Step (4) – **Iteration**: This provides crucial information to the reconstruction algorithm. The term "spatial diversity" refers to the number of different illuminated spots, their associated number of detected pixels, and their relative distances. Spatial diversity should be high, to provide many imaged pixels, and many different laser positions that are as far apart from each other as possible. The quality of the reconstruction depends on the signal to noise ratio of the data, the spatial diversity of the capture configuration, and other scene parameters. See the mission design section for further analysis on this.

Step (5) – **Reconstruction**: This should be performed by a computer on Earth. The spacecraft will buffer and transmit the data. Based on preliminary calculations and results from the proof-of-concept experiment, we estimate that localization of a hidden object patch will require at least 200 detected photons from that patch on each pixel of our camera in each time bin. If the subsurface material reflectivity is lower than expected, the entire process can be repeated during each overflight to slowly assemble enough information to reconstruct the hidden scene.

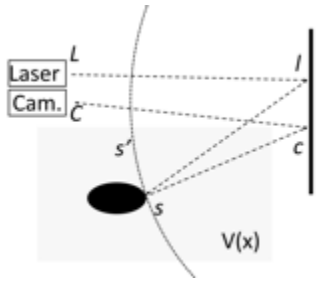

Figure 8: Each captured data sample can be projected to an ellipsoid in lab space. The ellipsoid s' is the set of all points s for which the distance L→l→s→c→C is constant for a given l and c. The voxel grid V(x) is the grid considered in the reconstruction.

The reconstruction of information from scattering photons is based on the analysis of round-trip-times and on the triangulation of signals to its origins. Here, the position of known objects has to be carefully recorded and analyzed to enable high precision triangulation. To illustrate the reconstruction algorithm we follow the light path of an individual ray as depicted in Figure 8. Further details can be

found in reference [xvii].

Referring again to Figure 8: light probing the scene leaves the position of the laser L at a time $t=t_0$ and propagates to a point l on the visible wall. Scattered light from the wall strikes a set of points s_n on the hidden object and returns to the wall at a set of points c_n that are observed by the pixels of our detection sensor. Finally, light reflecting off the points c_n is detected by the sensor with a center of projection C at a time $t_n>t0$. For a stationary camera and laser, our data therefore can be described in general as a seven dimensional transient dataset, where each detected sample intensity I(l,c,t) has the three coordinates of the projected illumination point l, as well as the coordinates of the projected camera pixel c, and the time t. A complete set T of these seven transient coordinates can be mapped to a set of potential hidden object points s_n', which includes all points in Cartesian lab space at which the scatterer causing light to be detected at T could be located. It is clear from basic geometry, that this set of points s_n' forms an ellipsoid (see Figure 8).

The goal of the reconstruction algorithm is to find the scattered positions s_n in Cartesian space giving rise to the detected transient data. This is achieved by back-projecting the transient data into the Cartesian lab space of the hidden object. This operation can be seen as taking a vote among a set of potential points s_n'. Each transient data point with coordinates Tn "votes" for the s_n' that lie on the corresponding ellipsoid in the hidden space. After all Tn are evaluated, the s_n' receiving the most "votes" are taken to be the locations of the hidden object. As a result, we obtain a confidence map of voxels V(x) in Cartesian space where each V represents the confidence that a scatterer s_n is located at position x.

Test System at UW Madison

We have completed a test system at the University of Wisconsin-Madison based on a Single Photon Avalanche Detector (SPAD). We plan to use similar technology for detection on the PERISCOPE system. The system is described in detail in Figure 9.

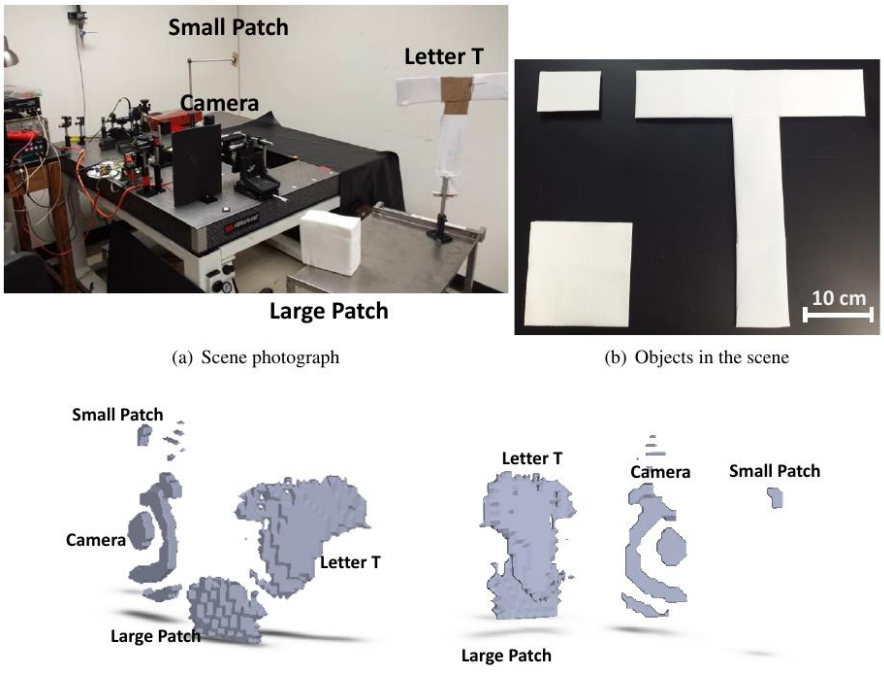

Figure 9: Scene with multiple objects reconstructed by our lab setup. The reconstruction shows the white patches to the left and right and the letter T. It also shows a reconstruction of the filter on the camera itself resulting form a bright specular reflection.

Our light source is a Amplitude Systems Mikan Laser producing 250 fs duration pulses at 1030 nm with a repetition rate of 55 MHz. We frequency double the laser pulse train to obtain a power of 50 mW at 515 nm. The laser light is directed towards one of the side walls of the laboratory with galvanometer-actuated mirrors. The SPAD detector is focused on a single spot on the same wall using a 1" diameter lens with a 1" focal length. We filter the detected light with an interference filter with a peak transmission at 515 nm and a full width at half maximum bandwidth of 10 nm.

The detector is a fast-gated SPAD module[xviii], a compact single-photon counting system capable of time-gating a silicon SPAD with ON and OFF transition times down to 110 p s, with ON-time adjustable between 2 ns and 500 ns with a diameter of 20 µm and a photon detection efficiency of 35% at 515 nm. The detector has less than 10 dark counts per second at 273 K and an after pulsing probability lower than 1% (with a 50 ns hold-off time). The time resolution is better than 30 ps (FWHM).

The detector module outputs a NIM (Nuclear Instrumentation Module) impulse synchronously with the detection of each incoming photon (Geiger-mode operation). The laser power is kept low to ensure less than 0.1 photons per pulse reach the detector. In order to block first bounce light directly reflected from the wall we use the time gating feature of the SPAD to disable it during the arrival of the first bounce. In our experiments, the detector ON-time window has a duration of 9.5 ns.

The NIM impulses corresponding to detected photons are sent to a PicoQuant HydraHarp Time Correlated Single Photon Counting (TCSPC) unit to produce a histogram of counts versus time after the illumination pulse. After a sufficient number of photons are collected, a different spot on the wall is illuminated and another dataset is collected. As a result we obtain a dataset of 185 time series.

Reconstruction Method

Our reconstruction algorithm has been modified and improved as part of this project. We rely on a filtered backprojection as originally proposed. Other methods have shown significantly superior resolution and reconstruction quality[xix], but employ regularization methods that make assumptions about the imaged scene. A filtered backprojection is also a more direct approach that allows us to assess the capabilities of the hardware.

Backprojection

Our algorithm projects every photon count $N(t,x_l, y_l, x_c, y_c)$ from the five dimensional space spanned by the photon time of arrival t, the coordinates of the laser spot on the relay surface x_l and y_l, and the coordinates of the patch on the relay surface observed by the detector x_c and y_c into the three dimensional Cartesian space V(x,y,z) close to the relay surface containing the scene to be reconstructed. Each photon count N is projected onto an ellipsoid in the reconstruction volume V. This backprojection process results in a confidence map of the reconstruction volume that describes the confidence that light was reflected by at the different voxels in the map. The algorithm can be summarized as follows:

- Create a grid of voxels V(x,y,z) referring to points in the reconstruction volume.
- For each collected photon count $N(t, x_l, y_l, x_c, y_c)$ compute the set of voxels V where the scatterer reflecting those photons could have been located. Increment the confidence value for those voxels by N.

Laplacian Filter

We then apply two filter methods to the data. The first filter is a Laplacian as used in [xx],[xxi]. This Laplacian enhances surfaces in the reconstruction volume.

Thresholding

To remove noise and obtain surfaces suitable for rendering and display, we apply a thresholding algorithm that favors continuous regions over individual disconnected voxels. Each voxel is considered above threshold if the confidence at the voxel coordinate and the confidence of at least 4 neighboring voxels is above threshold. We use a threshold between 0.2 and 0.5 for all our reconstructions. The threshold is currently adjusted manually to remove visible uniform noise from the reconstruction. This process could be automated in future applications if required, but would take at least several minutes of computation time in the current MATLAB implementation.

The result of the reconstruction is a voxel space that is converted to a 3D object using a graphical visualization tool (UCSF Chimera[xxii]) and rendered using SolidWorks (Dassault Systemes SOLIDWORKS Corp., Waltham, MA).

Including Angular Dependence and Simple Occlusion

The original backprojection algorithm completely ignored lambertian shading and occlusion. In previous work it was found that corrections for shading did not significantly affect the reconstructed geometry[xvii]. It would however affect the surface brightness of reconstructed surfaces and will become more important for more challenging geometry reconstructions. Past scenes were set up with flat relay surfaces and hidden scenes that did not contain occlusion. To obtain results on a cave scenario we had to implement a method to handle occlusions by the relay surface which will also allow us to add lambertian shading effects at the relay surface. The main reason is the complex geometry of the relay surface that occludes light send to one side of the relay surface (i.e. the rubble pile/mound at the bottom of the cave) from a significant part of the hidden volume. To address this we modified our algorithm to consider the surface normals of each illuminated spot and only draw ellipsoids in the half space that light would actually be reflected in given the orientation of the patch on the relay surface. We use the depth gradient provided by the geometry data of the visible scene to compute surface normals.

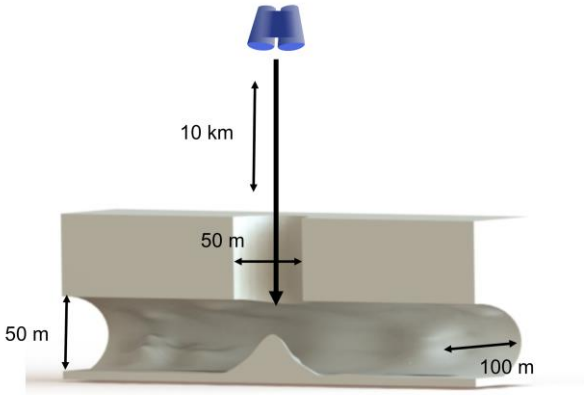

Figure 10: A model of the simulated scene representing an average skylight and (potential) lateral cave component.

DIRSIG Simulations

The Digital Imaging and Remote Sensing Image Generation model (DIRSIG) is a rendering engine developed by at the Rochester Institute for Technology. It generates physically realistic images and datasets for the design of air and spaceborne remote imaging systems. DIRSIG has been validated in multiple published studies and has been used in NSF and NASA funded research, for example for modeling Landsat [xxiii], [xxiv],[xxv].

The scene modeled is shown in Figure 10. The laser and detection system are positioned 10 kilometers above the lunar surface. The cave opening is circular and 50 meters in diameter and 100

meters deep. At the base it has a mound of debris and two 50m high and 100m wide tunnels leading in opposite directions. The tunnels walls have a wave shape that would be expected for tulles formed by a lava flow. In addition to this empty cave model we simulate caves that have obstacles placed in the tunnels at various distances. The reflectance of the cave materials is assumed to be 2.5%[xvi].

The laser is projected onto 25 different points at the base of the cave. The detection array images a square of 64 by 64 pixels with a side length of 32 meters. We assume a laser pulse width of 400 picoseconds and a detector time bin size of 200 picoseconds.

An example of simulated data for an empty lava tube is shown in Figure 11. Even without a reconstruction, the data reveals the location of the cave openings and allows an estimate of their size.

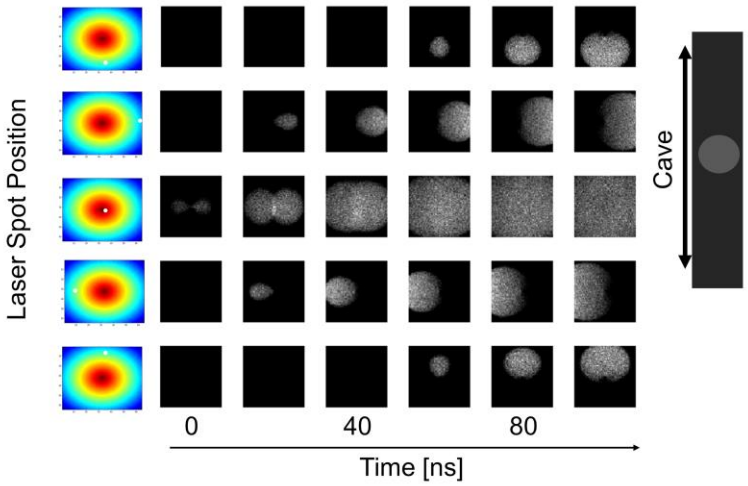

Figure 10: Frames from a DIRSIG simulation. The column on the left shows the position of the laser spot on the base of the cave as a white dot. To the right of that column are frames of the simulated data seen at different time points. Cave openings are to the top and bottom of the relay surface and it takes longer for light to return from those positions.

Link Budget

From the simulated data we can generate a link budget to estimate the power requirements of the system. To simplify this process, we simulated a scene of a 10 m diameter sphere, located in free space, 50 meters away from the center of the mound (see Figure 12). The brightest return signals are collected by SPAD pixels imaging surfaces close to the sphere after a laser illumination close to the sphere. Many detector pixels do not receive any light from the sphere due to occlusions. To estimate a representative number of photons per pixel collected we compute the average photon counts for pixels on the side of the mound facing the sphere for different illumination positions. We find about 0.56 to 60 photons per pixel for each of the 25 illuminated laser positions will be detected over a 0.056 second total time span with a 1 kW illumination power from an altitude of 2 km. From an altitude of 15 km only 0.01 to 1.1 photons are detected.

Figure 11: The simulated scene for the link budget.

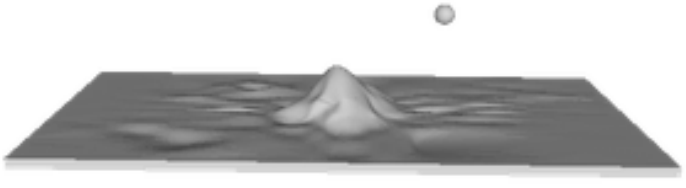

Alternatively we can calculate a link budget by estimating the light reflected by three surfaces in the scene as was shown in our Phase 1 proposal. We estimate about 7

photons per pixel for 1 kW over 0.058 seconds from 2 km altitude and 0.12 photons per pixel for 1kW over 0.56 seconds from 15 km. This is in agreement with the result generated by DIRSIG.

To operate in geiger mode we require each individual pulse to return less than 0.1 photons per pixel. To achieve this the pulse energy has to be approximately 100 mJ. For imaging altitudes above 2 km, increased laser power or a larger light collection aperture may be necessary.

Detector Noise and Background Light

For noise estimates we consider two different scenarios. One with a laser and detector in the green (532 nm) where detection noise is low, and one in the near IR at 1064nm, with higher detection noise, but lower solar background light and more available laser power.

SPAD detectors can achieve dark count rates below 10 Hz for green light and around 200 Hz for near infrared light. The light source would have to be converted to the green by frequency doubling with an energy efficiency of below 50%. Since the doubled photons are twice as energetic, the number of photons after doubling is however only ¼ of the number of near infrared photons.

Background light can come from three sources: If the cave is directly illuminated by the sun, direct solar radiation will dominate. If the sun is below the moons horizon the background is dominated by earth's reflection of the sun on the side of the moon facing the earth and by the milky way on the "dark" side of the moon, i.e. the side facing away from earth.

Assuming a solar constant of 0.8 $W/m^2/nm$ at 1 micron, a lunar reflectance of 0.025%, a challenging but possible 0.02 nm detection bandwidth and a 10 kHz pulse repetition rate our imager positioned 10 km above the lunar surface with a 1 m aperture diameter collects 30 photons per pulse of background light on each pixel. Since the actual signal is on the order of 0.1 photon per pulse assuming 1 kW of illumination, imaging in direct sunlight is thus not possible. Earthshine created by the reflection of the sun off the earth is about 1/10000 times as bright as direct sunlight and would create a background of 0.003 photons per second. This would amount to a signal to noise ratio of about 30:1 and is better than the ratio achieved in the current table top experiment. On the far side of the moon, only starlight would lead to a background light level far below 0.003 photons per pulse.

Instrument computer data volume, memory, and instrument computer challenges

Collected photon counts are preprocessed on the spacecraft to generate one histogram describing the time response per camera pixel and laser position. For 2500 time bins this would amount to maximally 25 * 64 * 64 * 2500 data samples of 16 bit or a total of 4.1 GBits. DIRSIG uses a lossless compression algorithm on the simulated data that cuts the file size approximately in half. This data needs to be transmitted to earth during the remainder of the orbit.

Reconstructions from Simulated Data

We developed an improved backprojection algorithm to create a reconstruction of the cave model from DIRSIG data. A reconstruction of the cave is shown in Figures 13, 14 and 15. The cave walls are at an oblique angle to the entrance and are very challenging to reconstruct. Future improvements and new reconstruction algorithms will alleviate this problem to some extent. Surfaces facing the relay surface can be reconstructed clearer and much deeper into the cave. A second model has an obstacle placed inside the cave at 75 meters distance. The obstacle is clearly visible in the reconstruction and we are able to estimate its size.

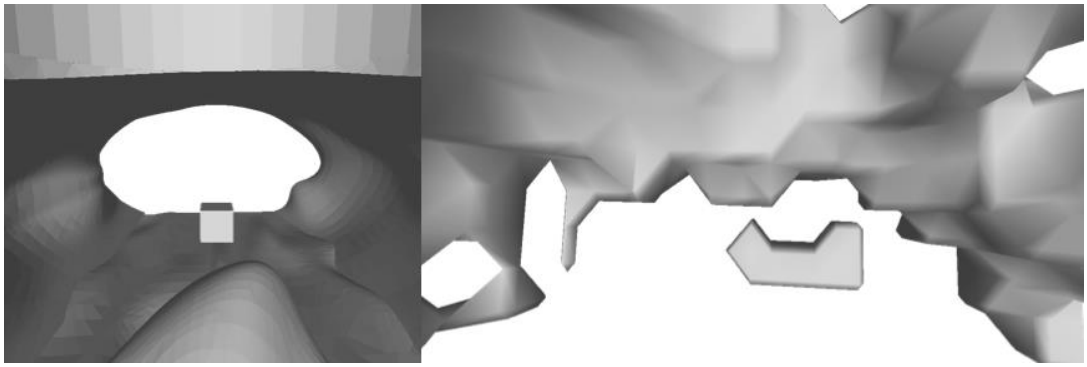

Figure 12: Ground truth (left) and reconstruction (right) viewed from a position below the skylight and looking towards one cave opening. The opening is obstructed by a 10 by 10 meter block positioned 75 meters into the cave. The mound itself is not part of the reconstruction.

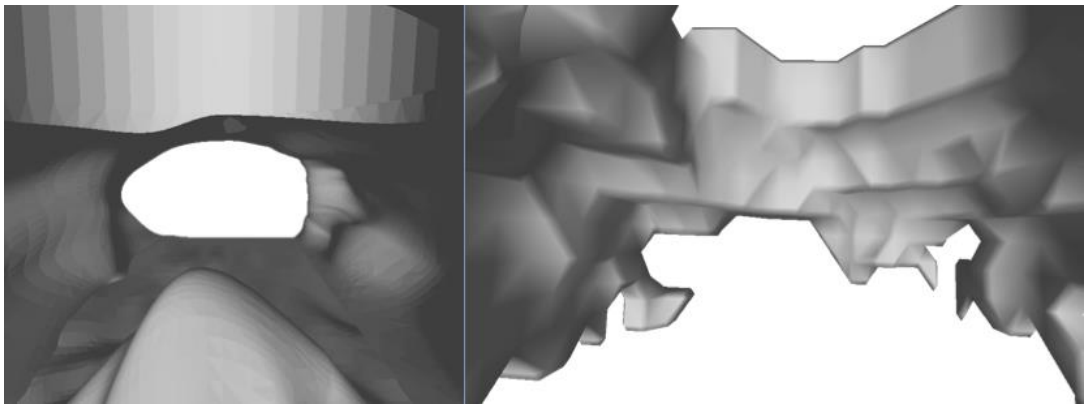

Figure 13: Ground truth (left) and reconstruction (right) viewed from a position below the skylight and looking towards one cave opening. The mound itself is not part of the reconstruction.

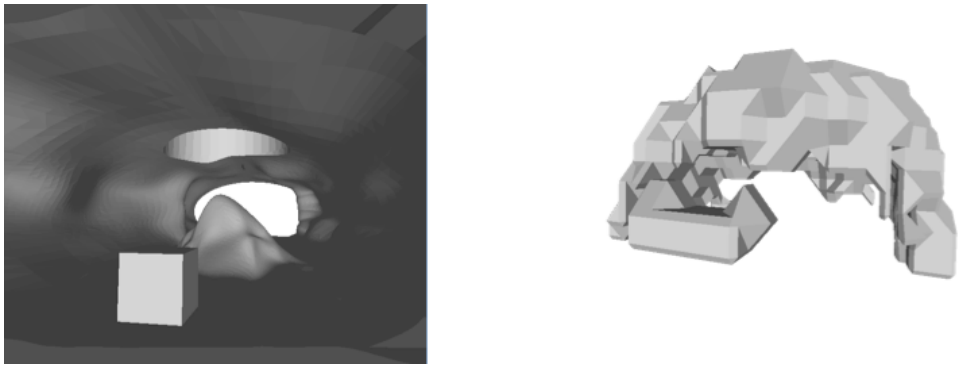

Figure 14: A view of the entire reconstructed object with a view of the ground truth from a similar perspective. The block in the cave is seen in the foreground. We can reconstruct the elliptical cave openings and the side walls below the skylight. Deeper into the caves, surfaces at oblique angles like the floor and the side walls are not visible due to the small amount of returned light.

Alternatives for Reconstruction

Our backprojection algorithm cannot incorporate certain aspects of the physics of light transport in the scene. Some challenges to be overcome include:

- Lambertian shading on surfaces in the hidden scene.
- Non lambertian reflectance models for surfaces in the hidden scene.
- Constraints on the hidden scene. Most importantly, the constraint to reconstruct surfaces in free space rather than volumes.
- Two or more than four bounce scattering.

Besides the filtered backprojection, different reconstruction algorithms have been demonstrated based on feature extraction and convex optimization approaches. An optimization-based approach using expectation maximization along with a fast and detailed forward model could incorporate all of these effects at the cost of added complexity and reconstruction time.

Alternatives for Simulated Data

Besides DIRSIG we also are able to render simple geometries directly in matlab. A further rendering tool we have available is a modified version of the physically based ray tracer (PBRT) that has been used in previous publications[xxvi]. We are also planning to use a rendering engine created by Jarabo et al[xxvii] when the code becomes available.

Visible Geometry Acquisition

The visible geometry of the relay surface needs to be known with an accuracy of about 0.5 meters. We have the option of capturing visible geometry simultaneously with the second bounce data at a slightly different wavelength. We will use a pushbroom LiDAR setup that uses a line of detector pixels to scan the scene. The power of the LiDAR illumination will be around 100 Watt. Direct geometry could also be collected separately, by a separate pass over the system, from a different instrument, or in an entirely separate mission.

Projected Capabilities

From a single pass over the cave at 10 km we expect to be able to detect the cave openings and detect obstacles to about 100 meters into the cave. We also will detect the cave ceiling and some features on the side wall and floor of the cave. In the future we hope to explore the ability to collect spectra of these surface using lasers at different wavelength or tunable laser systems.

Requirements for Photon Time-of-Flight Imaging from an Orbital Platform

To obtain a workable signal to noise ratio the orbital platform should include:

- A transmission mirror with a diameter of 25 cm
- A receiver mirror with a diameter of at least 1 m
- A seed laser/amplifier system capable of producing pulses of sub nanosecond width and of an average power of at least 1 kW. 10 kW or even 100 kW are technically feasible if additional power and space are available on the spacecraft.
- A laser linewidth of below 0.1 nm, ideally 0.02 nm combined with an equally narrow filter for the detected light to remove earthshine.
- A power system capable of powering the laser system with about 10 to 20 times the optical output power over 50 milliseconds. For a 10 kW laser system this would amount to 100 kW to 200 kW. The power requirements of the detection system are below 100 W and small compared to the requirements of the laser.
- A computer capable of compressing and storing about 4.4 GBit of collected data per observation and transmitting it back to earth. Further details about this area in the mission design section.

A key enabling technology for the mission, besides a SPAD array detector, is the availability of an efficient space qualified high power light amplifier. We are in close contact with engineers at Amphos GmbH (http://www.amphos-usa.com/) that provide a commercial 1 kW picosecond amplifier that is in the process of certification for space applications such as the destruction of orbital debris. We believe the slab amplifier technology used in their systems to be the ideal candidate for a PERISCOPE mission. Estimates for weight, power consumption and beam quality of the laser system are based on the properties of their systems (AMPHOS 400 with additional amplifier option).

Mission Design

The Periscope mission design takes advantage of the recent success of the Grail mission design campaign. The current mission trajectory relies on small high thrust impulsive maneuvers though a future study would look into a possibility of a low thrust version. The complete mission can be broken down into five phases, as follows:

Launch Phase:

The PERISCOPE mission will be launched on a low energy trajectory (similar to the GRAIL mission) taking approximately 4 months to reach the Moon. In contrast a more traditional direct trajectory would take approximately 3-7 days to reach the Moon, but at a cost of higher ΔV requirements. Using a low energy transfer trajectory to the Moon results in significant fuel savings during lunar orbit insertion (LOI), plus it also gives time for instrument and navigation checkout procedures after launch. The required C3 for a low energy trajectory like the one used by Grail mission is ~ -0.69 km^2/s^2.

Earth to Lunar Orbit Insertion (LOI):

The trajectory from Earth to LOI follows a lower energy manifold path through space. The current mission design anticipates three trajectory correction maneuvers (up to 15 m/s each). After its ~3.5-4 months of journey to the Moon, a ~190 m/s LOI ΔV puts the spacecraft in a highly eccentric (periapsis altitude of 25 km) lunar orbit with a period of ~12-14 hrs. The LOI maneuver is accomplished via a Bi-Prop engine (ISP=325 seconds), similar to one used on Mars orbiter missions.

Period Reduction Phase:

Even though the spacecraft is now captured around the Moon, the current orbit (due to its eccentric nature) is not very useful for the PERISCOPE mission. A series of period reduction maneuvers (PRMs) to lower the apoapsis are then required to lower the apoapsis of the spacecraft and reduce its eccentricity. This in turn also reduces the lunar surface velocity of the spacecraft at its periapsis (lowest altitude). The total number of period reduction maneuvers can be varied but result in total ΔV of ~450 m/s. The GRAIL mission used a similar strategy.

Targeted Science Observation Phase:

The Periscope instrument requires a lunar surface relative velocity of < 1.8 km/s to perform telescopic measurements. This in turn limits the spacecraft apoapsis to < 500 km altitude over the lunar surface.

After performing multiple PRMs, the spacecraft starts a targeted science campaign with periapsis orbit over different lunar cave locations. The obit requires little maintenance due to its

apoapsis being at ~300-500 km. The periapsis is maintained at < 10 km altitude to allow for useful targeted science investigations at most important cave locations. Multiple passes over a single site with a slow longitudinal drift is used to look inside the caves at different angles. The surface relative velocities are in the range of 1.65 to 1.8 km/s.

Global Mapping Phase:

Following a targeted science over a set of cave locations (preferably close to each other) the spacecraft then goes into a low altitude mapping phase where it maps the whole lunar surface at very low altitudes (5-15 km). The inclination of the orbit is mainlined between +- 65 degrees as most of the interesting cave locations are within that latitude band.

Our analysis showed that one ΔV maneuver (<10 m/s) after every 3 days was sufficient to maintain the periapsis above 4 km and apoapsis below 15 km. The ΔV requirement arises after taking into the higher order non-spherical gravity field effects on the moon, which tend to dominate at these low lunar altitudes. Lower periapsis are possible and would lower the ΔV requirements but could result in increased risk of the spacecraft hitting the crater.

Periscope's orbit also takes advantage of the natural longitudinal drift along with higher order gravity terms (such as J2) to complete a global map of the lunar surface with the specified 65 degree latitude band in ~15 days. Hence if we assume a mission period of 2 months we can achieve multiple passes over a large number of cave locations on the lunar surface. Furthermore, the lunar surface velocity is in the range of 1.65-1.7 km/s and this would serve as a design point for optimizing the instrument capabilities for that velocity.

Figure 16 shows an example science orbit strategy. Both the targeted orbits and a near global map of the lunar surface are shown. The four targeted lunar cave sites are Aristillus, Highland 1, Highland 2 and Lancus Mortis. The example trajectory has three passes over each of these cave sites, depicted by orbits in green, red, blue and yellow. The periapsis altitude is ~ 5km and apoapsis altitude is ~ 450 km. This results in a surface relative velocity of ~ 1.78 km/s over each of these sites. The total time of flight for doing 12 passes over the four cave sites is ~10 days. Following the targeted science phases, the spacecraft enters a global mapping orbit phase within a 65 degree latitude band. The cyan orbit in Figure 15 highlights this phase of the mission. The latitude longitude map for this mapping orbit is shown in Figure 17. It takes ~15 days for the spacecraft to complete this map.

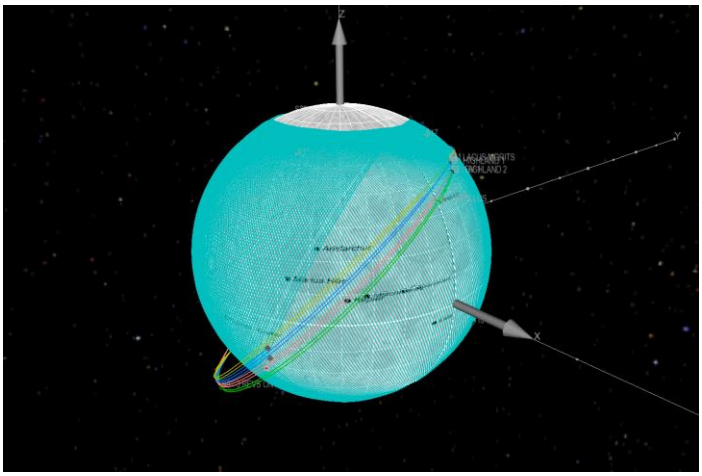

Figure 15: Example science orbit strategy both science phases of the mission.

The science phase will be designed to last for at least 60 days resulting in 4 such global maps with multiple passes over various cave sites.

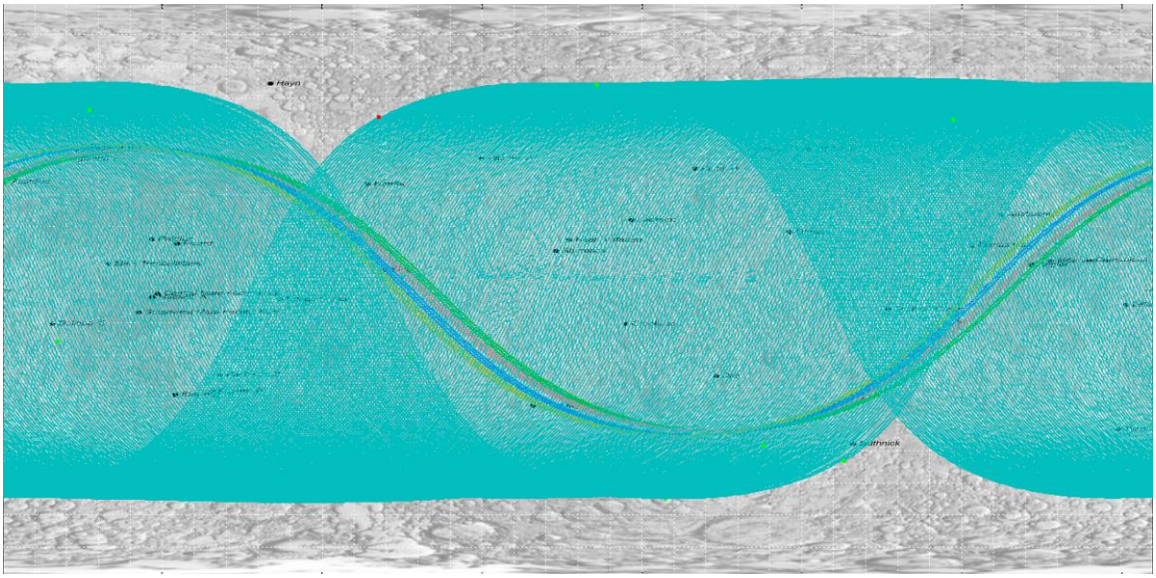

Figure 16: Lunar Lat-Lon Map (with names of cave sites) showing the 12 targeted science orbits (four cave sites) and the 15 days global mapping.

Though similar in nature, our orbital design is more flexible than GRAIL because of less stringent requirements than maintaining a fixed separation of the two GRAIL spacecraft flying in formation.

The preliminary orbital analysis has shown that apart from lunar surface relative velocity being around 1.68 km/s (on average), there is flexibility in other parameters for the mission orbit. Future trades studies are required to define the most optimized science campaign, but one could envision a series of "Targeted + Global" (as shown in the example orbit above) mapping phase campaign resulting in a global coverage with multiple pass datasets over each lunar cave site.

Preliminary Total DV requirements table:

Description	ΔV (m/s)
TCMs	45
Lunar Orbit Insertion (LOI)	190
Period Raise Maneuvers (PRMs)	450
Low Altitude Orbit Maintenance	150
Total	**835**

Table 1: Mission ΔV requirements

Spacecraft Design

The PERISCOPE spacecraft in this preliminary conceptual design phase is envisioned to be a 3-axis stabilized solar powered spacecraft. The total spacecraft system mass including 35 % margin is 590 kg. This includes the 115 kg for the science payload (the PERISCOPE instrument).

One of main mass drivers of the system is the power subsystem. Periscope carries three solar arrays (UltraFlex from SpectroLab), which are together capable of producing ~ 2kW of power at 1 AU. To provide the necessary peak power (~15 kW) for the Periscope instrument, approximately 23 kg of rechargeable batteries are stored onboard. This results in a low mass solution for providing the instrument with the required power for a short amount of time. The mission operations would have to be designed keeping the battery charging and discharging cycle in mind. This would be more challenging during the global mapping phase of the mission.

The spacecraft also enjoys a highly capable attitude control and determination system with 3 reaction wheels and 4 1-N thrusters for providing required control for the doing remote sensing science with the PERISCOPE instrument.

The telecommunication system in the current design consist of 1-m KA band high gain reflector capable of data rate of 5 Mb/s or more from the Moon. The data rate is also a function of required power for charging the batteries and needs further studies. Finally all the major maneuvers are performed by a 450 N Bi-prop stage using the Aerojet HiPAT engine. The main engine will also be responsible for performing the PRMs and the some of the orbit maintenance maneuvers if necessary. The total main engine (Bi-Prop) ΔV from this system is ~850 m/s.

Table 2 summarizes the mass breakdown for the various sub-systems of the spacecraft.

Table 2: Summary of Periscope spacecraft sub-systems masses (kg)

Attitude Control	23 kg
Avionics	7 kg
Power	40 kg
Propulsion	35 kg
Structure	60 kg
Thermal	26 kg
Telecom	14 kg
Science Payload	115 kg
Spacecraft Bus Mass (CBE)	320 kg
System Contingency	112 kg
Spacecraft Dry Mass	432 kg
Propellant	159 kg
Total Launch Mass	**590 kg**

Unique Power System Challenges and Solutions

The power system for the spacecraft has not been fully defined so the following assumptions have been made to facilitate the design of the laser power subsystem:

- Spacecraft bus voltage range of 34 to 24 volts
- Solar array for the spacecraft will be sized to meet the spacecraft load power without the laser
- Meaning that the solar array will not be able to support the 15 kilowatts pulse by itself

The laser requires 15 kilowatts pulse for one second; which includes 0.9 seconds of warm up and 0.1 seconds of laser firing. The laser fires once every two hours. Assuming the pulse is delivered at the minimum power bus voltage of 24 volts the pulse current will be 625 amps. This is a conservative assumption since the batteries will be at 100% state of charge before the pulse which is a bus voltage of 32.8 volts.

The Panasonic 18650 NCR-B battery cell has been tested at JPL for use on CubeSat projects and the Europa Clipper mission. It is similar to the ABSL battery flown on the SMAP (Soil Moisture Active Passive) mission. It has 3.1 amp-hours of capacity per cell with a full state of charge voltage of 4.1 volts. The battery pack will be made of 51 parallel strings of eight series cells for a 32.8 volt, 158 amp-hour battery. The capacity in watt-hours is 4400 with a mass of 22.5 kilograms.

Figure 17: Existing technology, such as our design based on parallel strings Panasonic 18650 battery cells as illustrated, can be adapted to meet the unusual power requirements of PERISCOPE.

A second, lower mass option combines Li-ion batteries with new higher voltage supercapacitors. The supercapacitors will require flight qualification.

This power system would use a ten parallel string eight series cell battery pack with a capacity of 870 watt-hours, with a mass of 4.4 kilograms and a bank of Skeleton supercapacitors:

- 12 cell string of 1300 F super capacitors (total string capacitance 108 F) (Model#SCHE1300)
- Cell ESR: 0.27 mohm (DC, 10 ms rating)
- Cell rated max 1 sec peak current: 1370 A (peak current here is 625 A, 547 A average)

- Specific Energy: 8.2 Wh/kg
- Maximum cell voltage: 2.85V
- Capacitor bank mass: 2.2 kilograms

Total mass of the power system with this option is 6.6 kilograms.

Figure 19: Skeleton Supercapcitors catalog image.

Analogy to Lunar Reconnaissance Orbiter (LRO) Concept of Operations

We consider a barebones mission carrying only a PERISCOPE system in order to assess whether it is realistic to carry it on a sub-Flagship class mission. For initial feasibility, we considered the Lunar Reconnaissance Orbiter (LRO), a ~$500M Discovery-class mission.

LRO operated with a maximum power requirement of 485 W, and 461 Gb/day, employing multiple instruments many of which operated almost constantly. It required a ~10.7 m^2 solar panel, providing 1850 W end-of-life power, or an average (accounting for geometry) of 800 W each orbit, thus in excess of the power draw ~485 W, of which ~125 W was for instruments. The Lunar Reconnaissance Orbiter Camera (NAC and WAV) produced most of the data, ~550 Gb/day (uncompressed), operating at 30 W peak, 22 W average. LRO's Ka-band transmitter power was 40 W and operated 100-300 Mbps, so 900 Gb/day. It also has 28 Gb onboard storage, and a very standard, rad-hard, single-board RAD750 computer operating at <20 W (usually 10 W). Thermal, mechanical, and command communications systems drew the remaining power.

Assuming that PERISCOPE requires ~4.4 Gb per skylight, thus ~100 Gb per day to transmit, it fits into the LRO envelope. With a mean power (assuming 1 observation per orbit) of only about 4 W, a mission with only the PERISCOPE instrument should come in over 100 W below that of LRO. Even 2-3 observations per orbit could be accommodated without too much difficulty, assuming that the battery power were sufficient. In other words, it is reasonable and conservative to assume LRO power and data requirements to be an upper limit for a PERISCOPE-centric mission, and thus anticipate a comparable or lesser cost.

Cost

We are able to be more precise with our costing by considering multiple previous mission designs.

The CML1/2 Cost Model is a Microsoft Excel based model that uses a small number of inputs that are known at early concept maturity levels (CMLs). The model is meant to establish the feasibility of a mission within a cost range. The model was developed for JPL's A-Team. The tool provides statistics on analogous missions from a database of Team X studies, JPL proposals, and historical actuals. It also generates a cost estimate according to the JPL Standard Work Breakdown Structure (WBS) based on simple rules of thumb and cost estimating relationships derived from Team X data.

Missions analogous to PERISCOPE were found by searching the database for Lunar Orbiters in the Medium (Discovery) or Small (Explorer) mission categories. There were 5 data points in the analogy database that met this criteria with total mission cost between $307M and $467M in FY15$

(without launch vehicles). Payload costs for these missions ranged from $21M to $58M and the flight system cost ranged from $113M to $207M. The table below provides additional statistics about these 5 analogous missions.

	Element Cost $M FY15	Total Mission Cost without LV FY15
Average	$151M	$383M
St. Dev.	$39M	$61M

	MinMission Cost FY15	Max Mission Cost FY15	Missions in Cost Range	Total Mission Cost without LV FY15		Phase E duration (months)		Payload cost $M FY15		Selected Element(s) Cost $M FY15	
				Max	$467M	Max	27	Max	$58M	Max	$207M
				Min	$307M	Min	9	Min	$21M	Min	$113M
				Average	St. Dev.	Average	St. Dev.	Average	St. Dev.	Average	St. Dev.
Bin 4	$350M	$500M	4	$402M	$51M	15	8	$48M	$10M	$160M	$38M
Bin 5	$0M	$350M	1	$307M	$0M	16	0	$21M	$0M	$113M	$0M

	MinMission Cost FY15	Max Mission Cost FY15	Missions in Cost Range	Wet Mass (kg)		Element Dry Mass with payload		Payload mass (kg)		# Instruments	
				Max	1292	Max	634	Max	57	Max	5
				Min	170	Min	78	Min	16	Min	2
				Average	St. Dev.	Average	St. Dev.	Average	St. Dev.	Average	St. Dev.
Bin 4	350	500	4	1092	245	486	144	37	16	3	1
Bin 5	0	350	1	203	0	0	0	16	0	4	0

The following inputs were used to derive a cost estimate by standard WBS for the PERISCOPE mission:

#instruments on element	1
Phase E duration (months)	3
Payload cost	$50 M
Power Source	Solar
Mission Risk Class	B
Prop System Type	Monoprop
Primary Telecom Band	Ka
Radiation Dose (krad)	<100

These inputs are conservative for PERISCOPE which may in fact be smaller than a Discovery class mission. The phase E duration may be reduced as well. The Payload cost is very uncertain at this time and there are no suitable analogies in the NICM database. The PERISCOPE team provided this estimate as a first cut in order to establish feasibility of the mission.

The cost for the mission is estimated to be $370M +/- 30%. The breakout by standard WBS is as follows:

Costs $M FY15	-30%	Nominal	+30%
WBS 1,2,3 Proj Mgmt, Proj SE, MA	$10M	$20M	$30M
WBS 4 Science	$10M	$10M	$10M
WBS 5 Payload	$40M	$50M	$70M
WBS 6 Flight System	$90M	$130M	$170M
WBS 7 and 9 MOS/GDS	$40M	$50M	$70M
WBS 10 ATLO	$10M	$10M	$10M
WBS 11 EPO	$0M	$0M	$0M
WBS 12 Mission Design	$10M	$10M	$10M
Reserves	$60M	$80M	$100M
TOTAL PROJECT COST	$260M	$370M	$480M

The following graphics shows how the PERISCOPE mission's payload, flight system, and total mission costs vs expected payload and flight system masses compare with data from the analogy database. The data point in green is GRAIL which would be a poor analogy for this mission.

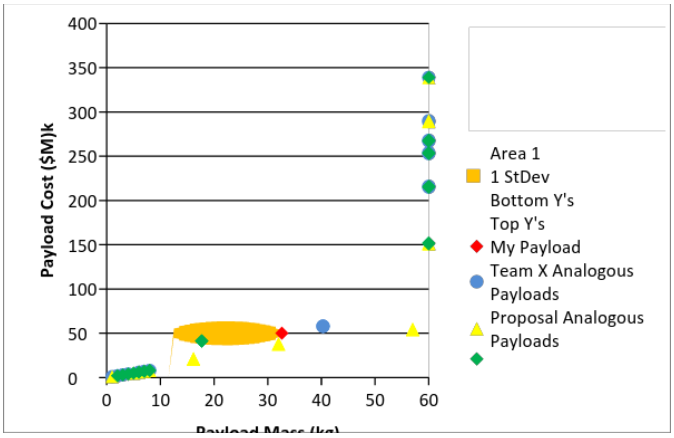

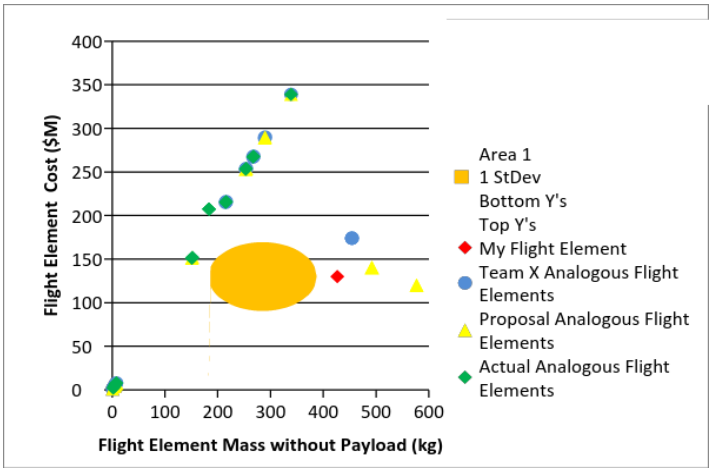

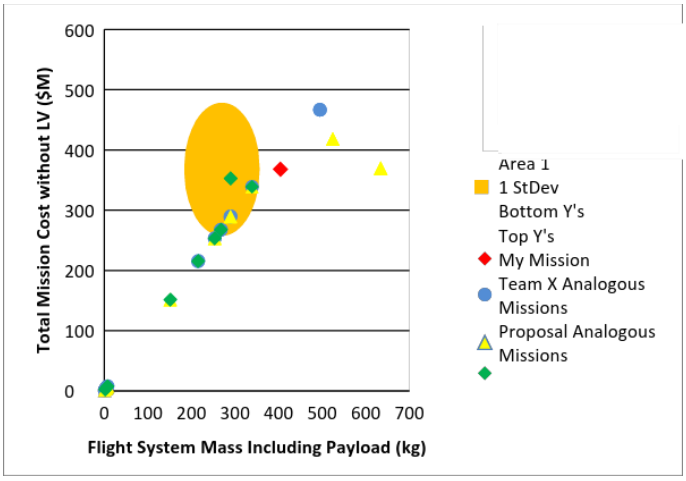

Figure 18: JPL's costing models demonstrate that the PERISCOPE mission would come in below the Discovery program cost cap.

Validity of the Estimate

The CML1/2 Cost Model has been validated by running the model for several Team X studies not in the database and for actual missions. Results are shown in the figures below.

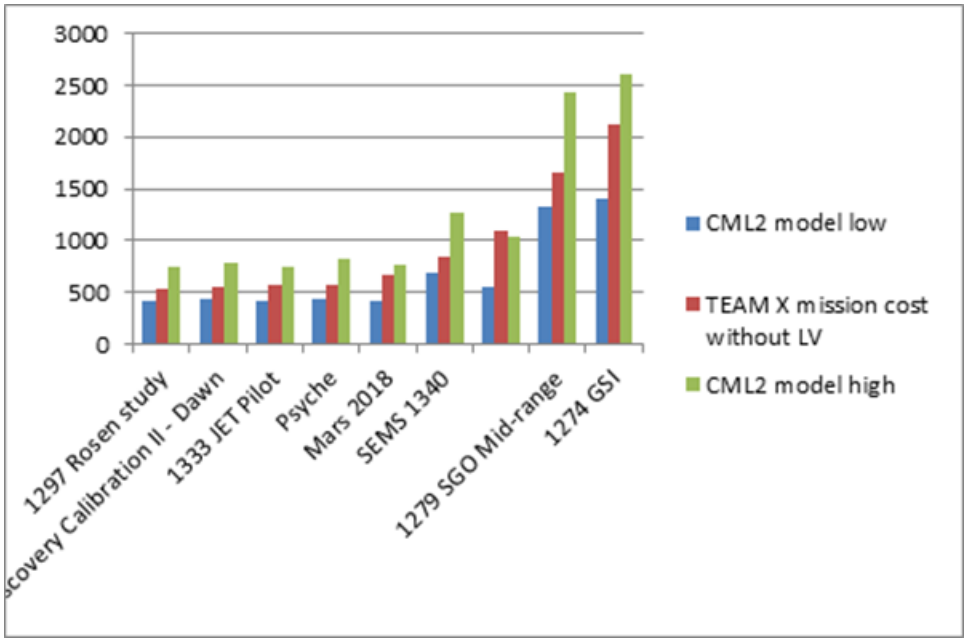

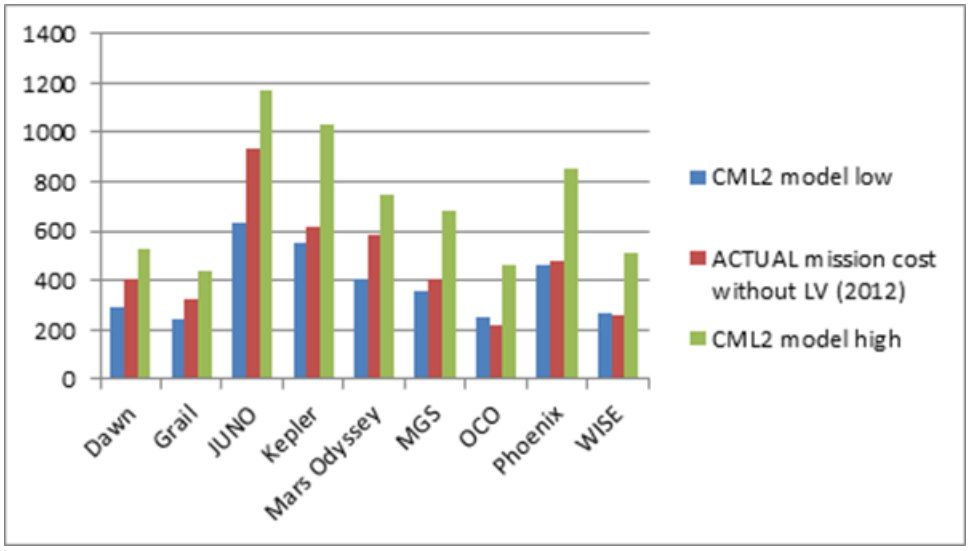

Figure 19: The CML2 cost model has proven exceptionally robust for competitive missions, especially in recent years.

To further establish the credibility of the model's estimate, the PERISCOPE estimate without payload was compared to the LRO actual cost without payload (from the Launch CADRe, converted to FY15$). Results were within 10% of each other. Note that the LRO payload costs were quite a bit higher than what we are assuming for PERISCOPE.

When more detailed information about the mission becomes available, other cost models can be run to produce a higher fidelity estimate.

Future Work

The next phase of work will focus on applying our Phase 1 results to real world situations as well as continuing to develop the instrumentation concept and refine the mission design. The tasks roughly separate into two categories: instrument and simulation work, and spacecraft and mission work.

Instrument/Simulation Work

- Providing inputs to simulations of PERISCOPE performance.
- Monte Carlo simulations of a range of different hypothetical cave geometries and materials in order to fully understand what we are likely to be able to map.
- Full 3-D laser scanning mapping of an Earth analog skylight environment so that we can provide a realistic cave and skylight terrain for inputting into simulation models.
- Explore question of whether scattering from an irregular, rough and rubbly surface will produce sufficient 3-D cover.
- Catalogue the known 3D extent of existing skylights/caves as inputs into cave models.
- Provide optical and spatial properties of analog materials for input into simulation models
- Catalogue properties of lunar samples, and terrestrial analog materials, including unweathered basalts.
- Liaise with A-Team and Team-X for systems study definitions.
- A-Team activity will explore parameter space in order to assess optimal solutions for PERISCOPE.
- Team-X activity will cost and detail a PERISCOPE-based mission in full detail.
- Detailed CAD/ZEMAX designs of the optical components of the spacecraft

Spacecraft/Mission Work

- Prepare for A-Team and Team-X studies mission and spacecraft design studies
- Explore value-added activities that could be part of a PERISCOPE caves-focused lunar mission.
- Explore possibilities of what might be present on lunar cave walls and determine whether PERISCOPE can be used to test for these.
- Develop specific lunar transfer trajectories and science phase orbits.

Conclusion

We are proud to present this final report to the NASA Innovative Advanced Concepts program and the world-wide community. We are grateful for the opportunity and the resources to explore this exciting mission concept. We are excited to continue work on this mission concept and contribute to the world-wide shared knowledge of our closest celestial neighbor.

Ad lunam in pace.

The PERISCOPE team

Citations

---

[i] Nasmyth, J. H. and Carpenter, J., 1874. *The Moon: Considered as a Planet, a World, and a Satellite*, John Murray, Albermarle St., London.

[ii] Wells, H. G., 1901. *The First Men in the Moon*, George Newnes Ltd., London.

[iii] Oberbeck, V. R., Quaide, W. L., & Greeley, R., 1969. On the Origin of Lunar Sinuous Rilles, *Mod. Geol.* 1, 75-80.

[iv] Haruyama, J. et al., 2009. Possible lunar lava tube skylight observed by SELENE cameras. *Geophys. Res. Lett.* 36, L21206.

[v] Squyres, S. et al., 2011. *Visions and Voyages for Planetary Science in the Decade 2013-2022.* National Research Council.

[vi] Wagner, R. V. and Robinson, M. S., 2014. Distribution, formation mechanisms and significance of lunar pits. *Icarus* 237, 52-60.

[vii] Heisinger, H., Jaumann, R., Neukum, G. and Head, J. W., 2000. Ages of mare basalts on the lunar nearside. *J. Geophys. Res.* 105, 29,239–29.275.

[viii] Ashley, J. W. et al., 2011. Lunar caves in mare deposits imaged by the LROC narrow angle cameras. *1st Intl. Planet. Cave Res. Workshop*, Lunar & Planetary Institute, Houston, TX, Abstract #8008.

[ix] Ashley, J. W. et al., 2012. Geology of the King crater region: New insights into impactmelt dynamics on the Moon. *J. Geophys. Res.* 117, E00H29.

[x] Akhmanova, M., Dement'ev, B. and Markov, M., 1978. Possible Water in Luna 24 Regolith from the Sea of Crises. *Geochem. Intl.* 15, 166.

[xi] Clark, R. N. et al., 2009. Detection of Adsorbed Water and Hydroxyl on the Moon. *Science* 326 (5952), 562–564.

[xii] Augustine Commission Report
http://www.nasa.gov/pdf/617036main_396093main_HSF_Cmte_FinalReport.pdf

[xiii] Personal communication at JPL

[xiv] Seu, R. et al., 2007. SHARAD sounding radar on the Mars Reconnaissance Orbiter. *J. Geophys. Res.* 112, E05S05.

[xv] http://www.hou.usra.edu/meetings/lpsc2014/eposter/1746.pdf

[xvi] Evans, D. L., T. G. Farr and J. B. Adams (1981) Spectral reflectance of weathered terrestrial and martian surfaces. Proc. Lunar Planet. Sci., 12 B, 1473-1479

[xvii] Laurentzis, Velten, Optical Engineering 2014; Velten et al., Nature Communications 2012

[xviii] F. Villa, D. Bronzi, Y. Zou, C. Scarcella, G. Boso, S. Tisa, A. Tosi, F. Zappa, D. Durini, S. Weyers, U. Paschen, and W. Brockherde, "CMOS SPADs with up to 500 μm diameter and 55% detection efficiency at 420 nm," *J. Mod. Opt.*, vol. 61, no. 2, pp. 102–115, Jan. 2014.

[xix] F. Heide, L. Xiao, W. Heidrich, and M. B. Hullin, "Diffuse Mirrors: 3D Reconstruction from Diffuse Indirect Illumination Using Inexpensive Time-of-Flight Sensors," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, p. to appear.

[xx] O. Gupta, T. Willwacher, A. Velten, A. Veeraraghavan, and R. Raskar, "Reconstruction of hidden 3D shapes using diffuse reflections," *Opt. Express*, vol. 20, no. 17, pp. 19096–19108, Aug. 2012.

[xxi] A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. G. Bawendi, and R. Raskar, "Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging," *Nat. Commun.*, vol. 3, p. 745, Mar. 2012.

[xxii] E. F. Pettersen, T. D. Goddard, C. C. Huang, G. S. Couch, D. M. Greenblatt, E. C. Meng, and T. E. Ferrin, "UCSF Chimera--a visualization system for exploratory research and analysis," *J. Comput. Chem.*, vol. 25, no. 13, pp. 1605–1612, Oct. 2004.

[xxiii] J. Schott, A. Gerace, S. Brown, M. Gartley, M. Montanaro, and D. C. Reuter, "Simulation of Image Performance Characteristics of the Landsat Data Continuity Mission (LDCM) Thermal Infrared Sensor (TIRS)," *Remote Sens.*, vol. 4, no. 8, pp. 2477–2491, Aug. 2012.

[xxiv] S. D. Brown, D. D. Blevins, and J. R. Schott, "Time-gated topographic LIDAR scene simulation," in *Defense and Security*, 2005, pp. 342–353.

[xxv] E. J. Ientilucci and S. D. Brown, "Advances in wide-area hyperspectral image simulation," 2003, vol. 5075, pp. 110–121.

[xxvi] D. Wu, A. Velten, M. O'Toole, B. Masia, A. Agrawal, Q. Dai, and R. Raskar, "Decomposing Global Light Transport Using Time of Flight Imaging," *Int. J. Comput. Vis.*, pp. 1–16.

[xxvii] A. Jarabo, J. Marco, A. Muñoz, R. Buisan, W. Jarosz, and D. Gutierrez, "A Framework for Transient Rendering," *ACM Trans Graph*, vol. 33, no. 6, pp. 177:1–177:10, Nov. 2014.