2005 IEEE International Workshop on Robots and Human Interactive Communication

Observations of a Science Team During an Advanced Planetary Rover Prototype Field Test*

Justin Glasgow, Erin Pudenz, Geb Thomas
Mechanical and Industrial Engineering
The University of Iowa
Iowa City, IA, 52242, USA
geb-thomas@uiowa.edu

Peter Coppin
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA, 15213, USA
coppin@cmu.edu

Nathalie Cabrol
NASA Ames Research Center/ SETI Institute
Moffett Field, CA, 94035-1000, USA
ncabrol@mail.arc.nasa.gov

David Wettergreen
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA, 15213, USA
dsw@cs.cmu.edu

Abstract – **Robotic exploration of remote planetary environments requires scientists and robots to interact effectively in order to generate accurate scientific conclusions. Optimizing the quality and reliability of these interpretations will increase the effectiveness and productivity of missions. During the 2004 Life in the Atacama (LITA) field test, scientists in Pittsburgh, PA interacted with a rover in Chile exploring the Atacama Desert for signs of life. Recordings of the scientists' actions and conversations revealed patterns meriting further study. For example, the stereoscopic panoramic images consumed a generous proportion of communication bandwidth (an average of 67% during week one and 72% during week two). Although the scientists effectively sampled all this data while viewing monocular reduced-resolution panoramas, they did not choose to look at the high-resolution version of every image. Assuming that stereo image pairs and triplets provide redundant information, the scientists studied 52% of the available information at full resolution. However, the information the scientists felt was necessary to form their impressions about the environment, represented only 18% of the total bytes of returned panoramic data. Stereoscopic depth information played an important role in navigating during past Mars missions, but results from this field test suggest that it will play a different role in the search for life. As new data sources are identified that assist in finding signs of life there will be a need to balance bandwidth so as to provide enough science information while still being able to safely operate the long-autonomous traverse capable rovers. The development of rovers capable of autonomous navigation and powerful onboard processing may reduce the need for downloading large amounts of three-dimensional data providing more bandwidth for science products**

Index Terms – **Efficiency, Remote Rover Exploration, Stereoscopic Imaging, Astrobiology**

I. INTRODUCTION

In September and October of 2004, scientists in Pittsburgh, PA used an autonomous rover, named Zoë, to search for life in the Chilean Atacama Desert. Understanding how the scientists used the rover, processed and analyzed the data, and eventually reached conclusions about the remote environment will help mission planners design future astrobiology robotic missions to maintain or even improve on the high level of operational efficiency already observed during the Mars Exploration Rover mission (MER).

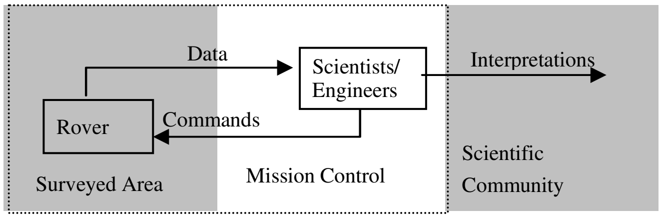

Fig. 1 Simplified system present during planetary exploration missions. It shows the loop between the mission team and the rover that leads to scientific interpretations.

A standard model, illustrated in Fig. 1, depicts the structure and process used by scientists when operating a rover to study remote regions. In this model, the rover collects pictures and takes measurements of the environment. The scientists analyze the returned data and form hypotheses. Based on these hypotheses the scientists generate new rover commands that result in the gathering of data that tests their hypotheses and/or generates new information. The rover receives the commands and collects the requested data. This cycle repeats as necessary until the science team develops a satisfactory understanding of the environment. The final output of the system is the formal interpretation the scientists' make which is reported to the scientific community and general public.

By analyzing the interaction between the rover and the scientists, the objective is to develop recommendations to maintain and optimize the quality, reliability, and productivity of rover missions. The interactions of specific interest are how the scientists process and analyze the returned data and then generate new rover command sequences. During a mission, the system describing the interaction between the science team and the rover is complex and includes a multitude of factors including team size, collective or distributed operations, and areas of scientific expertise represented. In breaking down the system to better understand the individual components before working with the whole system, the constraints we focus on are the limitations on available energy and communication bandwidth.

In the Life in the Atacama field test as well as previous rover field experiments, panoramas constructed from a mosaic of images taken by the rover at a single location have been very effective in conveying a visual impression of the rover's environment [1, 2, 3]. In many of these experiments, the scientists

---

* This work is partially supported by NASA's Applied Information Systems Research Program under NASA grant NAGW5-11981 and NNG05GA51G and the Astrobiology Science and Technology for Exploring Planets under grant NAG5-12890

also have access to three-dimensional panoramic images that present stereoscopic depth as part of a three-dimensional model. Generating the three-dimensional representation requires returning trinocular image sets which results in the rover returning essentially redundant tiles. This represents a significant demand on the limited mission bandwidth. The purpose of this three-dimensional data is typically to help in developing an understanding about the scale and geometry of the environment.

During the MER mission and past Mars missions, this scale and geometry information was important. During these missions a team of engineers are responsible for the complete navigation of the rover. For the LITA field test the rover is equipped with an onboard autonomous navigation and hazard avoidance system that allows it to traverse long distances [4]. The average daily traverse request during the mission was 3300 m. Consequently, the understanding of the scale and geometry of the local environment by the science team is primarily used to generate an initial understanding of the trafficability allowing the science team to develop a plan that follows a path of interesting science targets. The stereoscopic information was not used as extensively as rover drivers would during a full mission. One reason was that the science operation schedule required immediate development of the next rover command sequence leaving little free time to immediately analyze in detail all the received data. A more thorough assessment of the site occurred after operations had completed and time was available to assess the site.

During the mission, we investigated how or if the scientists still made scale and geometry assessments from the panoramas or whether those decisions were primarily left to the engineers. The scientists' efforts to determine scale and geometry are of interest because in general people's perception of the size and height of distant objects along with the slope of terrain is generally inaccurate [5]. The rover's decoupling of an individual's natural action/perception relationship may in fact worsen the ability to perceive size and height. Reference [6] showed that in natural settings participants overestimated, by up to 225%, the height of unfamiliar objects when viewing these objects in an open field with no outside visual clues. In a separate study participants, when asked to compare the height of a 42-inch stimulus with targets spaced 100-4000 feet away on a runway, overestimated the size of the stimuli by 6.41% to 34.04% [7].

II. MISSION DESIGN

During September and October 2004 the rover explored two distinct regions of the Atacama, one relatively "humid" (1 cm precipitation/year from fog) and the other dry. Each area was investigated for one week, with the objective of identifying signs of life. Throughout the mission, scientists located in Pittsburgh, Pennsylvania remotely sent commands to the rover. We monitored the scientists' actions and processes during both weeks of exploration.

A. The Rover – Zoë

During the 2004 campaign, the scientists were concerned with identifying potential biological habitats and finding evidence for life in an Earth-based Mars analog. The rover carried a payload designed for detecting life in extremely arid environments [4]. Table 1 details the average returned data product size and image characteristics for the payload instruments. The fluorescence imager (FI) is a camera mounted on the belly of the rover designed to measure the amount of fluorescence in of an area treated with fluorescing dyes. Mounted behind the FI is a pair of cameras, the workspace imager (Wks), designed to capture a field of view containing portions of the FI capture and the bottom of the accompanying panorama. Additionally, the rover is equipped with a Visual/Near-Infrared (VNIR) and a Thermal Infrared (TIR) spectrometer useful to determining the chemical composition of selected rocks targeted by the scientists. The stereoscopic panoramic imager (SPI), positioned on a mast extending slightly forward from the rover, provides high-resolution images that are tiled together to create the monoscopic (2D) and stereoscopic (3D) panoramas. A monoscopic panorama contains 218 tiles. Lastly the scientists could request a weather report from a weather station. The weather report included information on the temperature, the relative humidity, the wind speed, wind direction, and the sun intensity.

Table 1: Details on the data size and data characteristics of the various payload instruments available to the science team

Instrument Type	Data Size	Data Characteristics
Fluorescence Imager	0.2 MB – 1.0 MB	10 cm x 10 cm field of view (FOV)
Workspace Imager	0.5 MB	
Spectrometers	0.01 MB	Spot w/ d = 1cm
Stereoscopic Panoramic Imager	.05 MB / Tile	15.9° x 21.1° FOV
	33 MB / Set	1280 x 960 pixel resolution
Weather Station	1MB	

B. Participants

Over the two weeks of the mission, nine different scientists participated. There were a total of six scientists present for week one, and seven during week two, with an overlap of four individuals between the two weeks. During week two an eighth team member participated via conference calls and emails. The team consisted of four geologists, four biologists and one spectroscopist. The training and experience of each of the scientists was broad and varied to provide a variety of perspectives as would occur during an actual mission. The expertise varied from two still working on their graduate degrees to others that have spent 20+ years working in their respective field. Most of the team members had a mixture of planetary, including mission, and field science in their backgrounds.

C. Daily Activities

The science operations simulated a Mars mission and shared many components of MER's operational scenarios. The objective to LITA was to run a genuine astrobiology exploration of the most extreme and one of the least documented deserts on Earth. The goal was to map the distribution of life and habitats of the Atacama. At the beginning of each week's science operations, the scientists received a collection of orbital images of the region including the simulated landing ellipse. The scientists spent the first day analyzing this satellite data developing a mapping of the region. This mapping identified areas that due to morphology and/or spectral readings had a higher probability than surrounding areas of being habitable by microbial life. Around 6:00 pm the first evening, the data download from the rover began. The first data product received by the science team was a 360° panoramic image of the landing site. The scientists then took the next six hours to analyze the data and generate a command sequence for the rover. Using EventScope, a software program developed for the LITA operation, the scientists entered in the completed command sequence. EventScope provides an interface for identifying points of interest and entering in English command sequences, which it converts to a machine language understood by the rover. The rover receives the up-linked command sequence and collects the requested data the next day.

The following days the science team would arrive at ~12:00 pm. They would spend the next two to three hours further analyzing data from the previous evenings download and working on documents summarizing their observations. During this time, the work was silent with only short sporadic conversations between scientists, often between those specializing in the same area. As the scientists wrapped up their individual observations, they came together for a team meeting around 4pm. During this meeting, the team would finalize any observations from the previous day and then begin to layout a preliminary command sequence for the next day. After meeting for an hour, they would break for dinner and wait for the new data uplink.

Upon uplink completion, the scientists would silently study the data for approximately an hour, and then come together as a team. During this team meeting, the individuals would discuss initial findings and their thoughts about the next day's rover command sequence. Before breaking the meeting, the team would discuss the updated status of the command sequence and break into subgroups to prepare information for the upload. One of the team members would work on beginning to enter initial information into EventScope as well as preparing any supporting documents to send to the field team monitoring Zoë's actions. Another two scientists would work on triangulating Zoë's current position as well as the positions of stops during the next day's sequence. The remaining team members would continue to process and analyze the new data in particular the FI images.

The science team came together one last time after entering in all the data to EventScope. They would review their plan for the next rover sequence identifying any changes or clarifications needed in the plan. After completion of the meeting most of the science team left the remote operation command center while one or two remained to make necessary changes to the plan. The final upload to the rover occurred between midnight and 2am.

III. METHODS

A. Data Collection

The daily download from the rover was limited to a bandwidth of 150 Megabytes. The science team accessed all the data for the mission through a password-protected website. The website displayed most of the data products in their raw form, except for the panoramic images. The panoramas were tiled together and displayed at 22% of full resolution; a full resolution panorama would be a 33 Megabyte file. The individual high-resolution tiles returned by the rover to make the panorama were readily available. An access log of hits to the website showed when a scientist viewed one of the available data products.

In addition to recording the data products viewed by the science team, a combination of audio and visual recordings monitored the team members' conversations and actions throughout the mission. Written notes taken every five minutes of each scientist's current activities supplemented these recordings.

During the mid-afternoon science team meetings, the scientists provided a list of hypotheses formed about the environment. The team members also provided interpretations about the geology and morphology of the environment. A field team in the Atacama collected samples or took measurements as needed to provide the ground truth answers for the hypotheses while keeping the remote science team blind about the results.

B. Data Analysis

Upon completion of the field test, we tabulated the total bytes of data products downloaded and available to the science team. Data was organized for analysis by the day it was downloaded and by the payload instrument that collected the data product. Next cross-referencing the website log showed which available data products the scientists viewed during the mission.

Based on notes made during the mission, a transcriptionist wrote down important conversations from the audiotapes. The transcripts are currently being processed to identify the statements the scientists make while analyzing the data and forming a hypothesis about the environment.

A comparison between the ground-truth evaluation and the hypotheses generated during team meetings shows instances where errors occurred in the system. The operational definition of an error is an instance where the consensus opinion of the science team differs from the ground truth to such a degree that it could have significant impact on the success of a mission. Errors are expected in the rover-scientist system as they do also occur when humans are directly in contact with the field. However, by identifying common errors in the human/robot system we can begin to design components that will reduce the error rate in the rover-scientist system so that it closely approaches those observed in direct field interpretations. After identifying errors an analysis of the data usage and transcripts will show the possible causes and will lead to potential solutions for the issues.

IV. RESULTS

During phase one, the rover returned just over 1 Gigabyte of data and in week 2 515 Megabytes. Despite the difference in volume between the two weeks, due to the difficulty in determining the bandwidth of the various data packages the first two days, the allocation of bandwidth to the various payload instruments was consistent, as illustrated in Fig. 2.

Throughout both weeks the SPI data consumed a sizable portion of the daily bandwidth totaling 900 Megabytes, see Fig 3. Based on the access log analysis, the science team viewed 166 Megabytes, or 18%, of the SPI data. This number is an erroneous representation as the SPI data consists of tiles returned in triplicate sets. The data is essentially redundant between the left, right and middle tiles in a set, meaning that there is only 300 Megabytes of data that is unique from the perspective of the science team. The entire science team viewed this 300 Megabytes of data in the panoramas at 22% of full resolution. About

52% of the data was viewed at full resolution (155 Megabytes). Lastly, an analysis of the azimuth and elevation of these tiles viewed by the science team shows that the upper third of tiles in the panoramas, representing primarily sky, were never viewed at high resolution. These tiles contain important astrobiology information, like the daily evolution of fog in the "humid" region, but this data is easily discerned at lower resolution. There is no reason to expect the scientists to need to view these tiles at higher resolution suggesting that there is only 200 Megabytes of actual critical high-resolution SPI data. Therefore, the 155 Megabytes they viewed at high-resolution represents 78% of the critical returned data.

V. DISCUSSION

Initial data analysis indicates that the science team devoted very different amounts of time to each of the data sources. Time spent analyzing does not directly correlate to the importance of the data, as some datasets are more complex to process and understand than others. While initially it appears that the scientists utilized a very low percentage of the SPI data, they in fact did utilize 78% of the critical scientific data. This is a reasonable result as expecting 100% utilization of the high- resolution tiles is unreasonable (e.g., some tiles do not contain rocks that require being viewed at a high resolution.

While the scientists were very thorough with what has been deemed the scientifically relevant tiles, the fact that they only viewed 18% of the returned SPI needs to be understood. The reason for returning redundant trinocular SPI data sets was to develop 3D panoramas, which in the past have served to provide detail on subtle features for engineers to guide rovers around obstacles. Since the driving was primarily autonomous (equivalent to blind drives on MER) and the engineering team was decoupled (physically in another location) from the science team, there was an apparent lesser usage of the 3D data at the remote operations center in Pittsburgh.

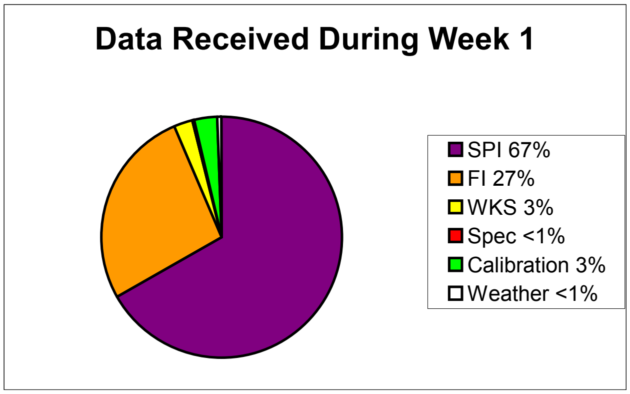

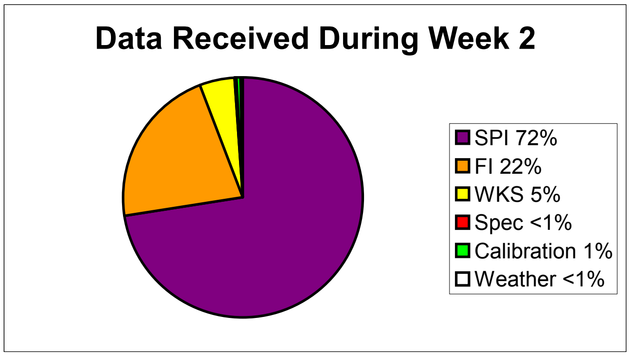

Fig. 2: Percentage of data return committed to each data product during weeks one and two of the Atacama Field Test. The component responsible for most of the data collection was the Stereoscopic Panoramic Imager (SPI), followed by the Fluorescence Imager (FI) and the Workspace Camera (WKS). The Spectrometer (Spec) data consumed relatively little bandwidth.

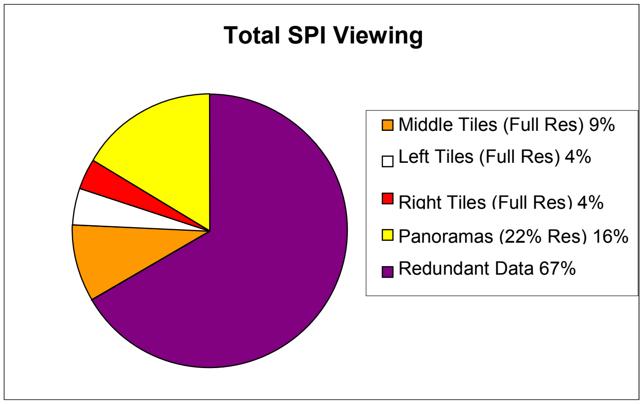

Fig 3: Breakdown of the available 900 Megabytes of SPI data. A large portion of this data is redundant information necessary to develop 3D panoramas of the scene. The other sections show the percentage of data viewed at full or less than full resolution

Essentially the conclusion here seems to be that given the mission design there were not the time resources available to analyze 150 Megabytes of data everyday. In particular there was no person or group of people specifically committed to driving the rover, so everyone on the team had to balance their scientific curiosity with assessing gross trafficability to keep Zoë safe. As the scientists become more comfortable with the payload instruments we may find them using the 3D images to assess potential habitats in rocks and soils. Very close attention will be played to the role of 3D data in astrobiology missions and allocation of bandwidth to specific datasets as we proceed with the LITA rover field campaign in 2005The reason to increase available bandwidth for other data products is that they require less bandwidth and are still quite data rich. A good example is spectral data that only requires 0.1 Megabytes per data request and once the data has been processed has a reduced margin for interpretation error. This data source is very important as knowing mineral composition helps determine the potential habitability of a given area. The relative influence of FI, SPI, and Spectrometer data on the final scientific conclusions has not yet been fully analyzed, but our initial study of the transcripts and direct observations suggests that the spectral data played a pivotal role in the scientists thinking, which is expected as composition is a very important aspect in astrobiology.

While the 3D data received little attention from the scientists there are a number of reasons why it may still have played an important role. For example, a brief observation of the 3D models may have been sufficient for the scientists to incorporate the information into their mental model of the environment. This explanation is at odds, though, with the inconsistency and inaccuracy of the scientists' estimation of the cliff height.

Past robotic field-tests have noted similar perception errors as those discussed in the introduction, so during the hypothesis generation sessions, we queried the scientists ability to perceive height in the environment [8]. To pose this question, we showed the science team a full panorama (Fig. 4) and a high-resolution inset of the panorama (Fig. 5) with a rock face of interest surrounded by a black box. The objective for the scientists was to determine the height between the two black marks on the rock face. Three of the scientists picked heights that fell in the general range of 3-4 meters, while one provided a broad range estimate of 5-10 meters. Fig. 6 is a graph of the scientists' estimates along with a horizontal bar showing the true height of 7 ± 0.5 meters.

This shows that the scientists do not assimilate the fine details about the environment but they may still be getting a general enough sense of the environment. Essentially during the compressed time scale experienced during the LITA field test the science team does not have the time to deeply analyze the data. Team members try to extract the critical information quickly for sequence planning and come back for a refined analysis at later stages. At these later stages the stereoscopic data is useful for determining scale, morphology, and shape in the environment, along with the geometry of rocks and soils.

A second issue related to increasing the efficiency of the system is determining a useful resolution at which to capture and display the 2D panoramas. The scientists did view a large percentage of the lower two thirds of the panorama at full resolution but very little of the upper portion. Only a complete ground truth assessment of the site will reveal if the scientists missed any details. The objective here is just to understand common errors so that overtime the overall reliability of the system can be improved; there is no expectation for absolute perfection in interpretations as even fieldwork has a margin of error.

There are a number of factors to consider here especially the fact that the further away from the rover an object is the less resolution it will be captured under. The goal is to understand what conclusions the scientists want to reach using the data and then ensuring that if possible sufficient data sufficient to reach the conclusion is available. In previous experiments, we have taken the scientists out to the location of the rover to determine if there were other important features they did not see in the image [8]. The 2005 field expedition will incorporate this procedure. Anecdotal evidence of various features "discovered" the day after the initial analysis suggests that there were some environmental features overlooked during the first examination of the panoramas.

If any of these instances end up being features that the science team would have liked to explore further before moving on with the rover we need to facilitate their ability to identify these features within the panoramas. A potential reason the scientists may have missed these features despite seeing them in the high-resolution tiles is that the tiles were presented out of context making it difficult to assimilate the information into a complete mental model. A proposal by the science team that may control bandwidth while not decreasing the understanding of the environment is to be able to request a panorama with a varying resolution throughout the image. This suggestion has been implemented by the development of a visualization tool that will be tested during the 2005 campaign.

Fig. 4: 2D panorama compiled for locale 8. The scientists estimated the height for a rock face located in the right hand background of the image.

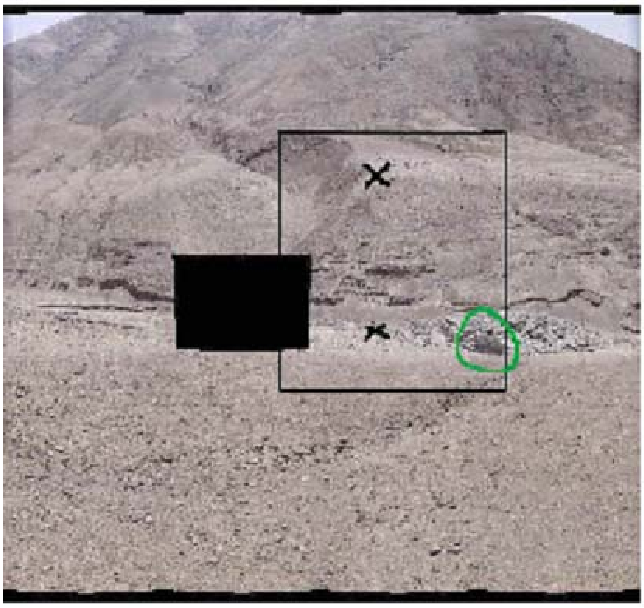

Fig. 5: Inset from the panorama showing the rock face height estimated by the scientists. The height is the distance between the two black x's. The solid black rectangle indicates the size of a single image in the panoramic mosaic.

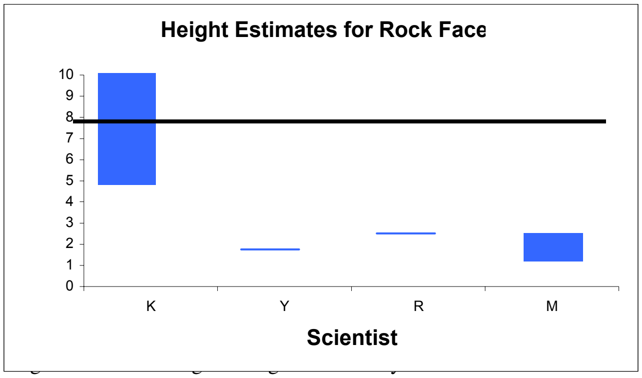

the rock face shown in figure 5. Some scientists offered a range for the height while others hypothesized a specific height. The horizontal bar shows the actual measured height of 7m. The error on the measurement is ± .5 m.

VI. CONCLUSION

The status of this research is still very preliminary and cannot indicate yet the true usefulness of any of the data products available to the science team. During the first three months of the MER mission it took the mission team 18 hours to generate a daily plan, while at its current operational level this task is completed in only a few hours. This reduction in time allows the science team to look back at the data and perform more complete analyses. It is also true that the critical data used for planning is still analyzed in a very short time frame (sometimes few hours only) and generates situations very similar to those encountered during the LITA field experiment where plans are created before all the interesting features are identified. Continued analyses such as this will help identify ways to assist the science team in efficiently identifying important features so that everything is fully analyzed before the rover moves on.

This sort of synergy where the science team can quickly develop a daily plan will be hard to develop during a field test like LITA where operations only last for one week at a time, and the capabilities of the rover and that of the interface are still being implemented. By understanding what the scientists feel is the critical data for an astrobiology investigation and what they would like available but cannot initially analyze, mission and rover designers can develop their systems to emphasize the appropriate data products helping to reach the streamlined synergy between human and rover much earlier in the full Mars mission.

REFERENCES

[1] C. Stoker, and B. Hine, “Telepresence control of mobile robots: Kilauea Marsokhod Experiment,” presented at *American Institute of Aeronautics and Astronautics*, Reno, NV, 1996.

[2] N. A. Cabrol, et al, “Science results of the Atacama Nomad rover field experiment, Chile: implications for planetary exploration,” *J. Geophys. Res.*, **106**, E4, 7664-7675, 2001.

[3] G. Thomas, A. Bettis, N. Cabrol, A. Rathe, and T. Foster, “Analysis of Science Team Activities During the 1999 Marsokhod Rover Field Experiment: Implications for Automated Planetary Surface Exploration,” J. Geophys. Res., 106, E4, 7775-7783 2001.

[4] Field Robotics Center (2004). Life in the Atacama, 2004. http://www.frc.ri.cmu.edu/atacama/, downloaded 2/1/2005.

[5] D. D. Woods, J. Tittle, M. Feil, and A. Roesler, “Envisioning Human-Robot Coordination in Future Operations,” IEEE Transactions on Syst ems, Man, and Cybernetics C., 34: 138-153, 2004.

[6] R. B. Joynson, L. J. Newson, and D. S. May, "The limits of overconstancy," Quarterly Journal of Experimental Psychology A, vol. 17, pp. 209-216, 1965.

[7] A.S. Gilinsky, "The effect of attitude upon the perception of size," American Journal of Psychology, vol. 68, pp. 173-192 1955.

[8] J..L. Wagner, “Erroneous or accurate: assessing decisions and decision-making behavior during a human-directed robotic geology field test,” PhD dissertation, the University of Iowa, 2004.