

Robotic Scouting for Human Exploration

Matthew C. Deans¹, Terrence Fong¹, Mark Allan², Xavier Bouyssounouse¹, Maria Bualat¹,
Lorenzo Flueckiger³, Linda Kobayashi¹, Susan Lee², David Lees³, Eric Park²,
Estrellina Pacis⁴, Liam Pedersen³, Debbie Schreckenghost⁵, Trey Smith³, Vinh To², Hans Utz⁶

¹NASA Ames Research Center, Moffett Field, CA. ²Stinger Ghaffarian Technologies Inc., Moffett Field, CA

³Carnegie Mellon West, Moffett Field, CA, ⁴SPAWAR Pacific, San Diego, CA,

⁵TRAC Labs, Houston, TX, ⁶Research Institute for Advanced Computer Science, Mountain View, CA

Abstract

By 2020, NASA plans to return to the Moon with a new series of regularly spaced surface missions. Crewed missions will initially be "extended sortie" (e.g., 1-2 weeks). During the first few years of the lunar campaign, humans will be on the Moon less than 10% of the time. During the 90% of time between crew visits, robots could perform tasks under ground control. This paper presents the system design for a prototype robotic recon robot and ground control approach, as well as a terrestrial analog field test designed to assess the utility of recon for augmenting and assisting human exploration of a lunar-like environment. Results are presented for recent field testing of the reconnaissance robot in northern Arizona.

I. Introduction

On June 14-26, 2009, we conducted a test of *robotic recon* at Black Point Lava Flow, Arizona, USA. The K10 Robot was used to explore a planned traverse route designed for the Lunar Exploration Rover (LER)[10]. The goal of the test was to improve our understanding of how robotic scouting can help plan EVAs, and how robots might best complement human crews. Here we summarize the objectives, describe the test setup, and present some preliminary results.

By 2020, NASA plans to return to the Moon with a new series of regularly spaced surface missions. Crewed missions will initially be "extended sortie" (e.g., 1-2 weeks). During the first few years of the lunar campaign, humans will be on the Moon less than 10% of the time. During the 90% of time between crew visits, robots could perform tasks under ground control.

Several years prior to surface missions, orbiters (e.g. LRO) will map the Moon, contributing new remote sensing data to our existing body of knowledge represented by historical archives and newly processed data from past orbiters (e.g. Clementine). However, remote sensing data may not be of sufficient resolution, lighting, or view angle to optimize mission planning. Thus, it will be important to acquire data from the surface.

Robotic recon can perform this task[7-9]. Scouting is an essential phase of field work, particularly for geology, and can be traverse-based (examining stations along a route); site-based (examining an area); survey-based (systematically collecting data) or pure reconnaissance.

During robotic recon, instruments can be used to examine the surface and subsurface at resolutions and viewpoints not achievable from orbit. The data can be used to prioritize targets of interest and improve crew productivity. Robotic recon can also help identify hazards and evaluate alternate routes to reduce operational risk [1,2,4,5].

As a practical example of how recon could have helped during the Apollo program, during Apollo 17 EVA-II, the crew drove from the landing site to the South Massif, then worked their way back to Shorty crater. At Shorty Crater, Harrison Schmitt observed orange soil (volcanic glass)—perhaps the most important discovery of the mission. If NASA had been able to conduct recon prior to the EVA, the orange soil could have been identified before the crew arrived, and the EVA plan could have been optimized to reduce time at earlier stations, making more time to work at Shorty Crater.

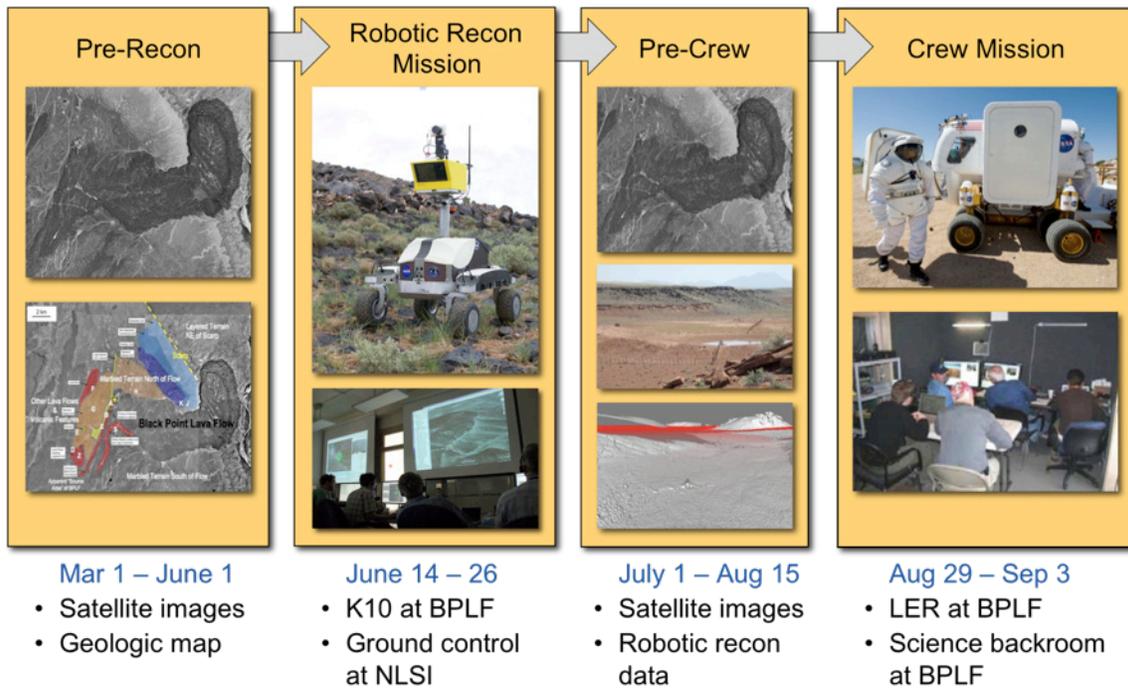


Figure 1. The approach to the recon experiment includes four phases. In our experiment we compared Crew Missions with recon data (following the flow above), and without recon (using only Phases 1 and 4).

II. Approach

In this section we describe the setup of the recon testing in detail. This includes the robot platform, the selected tasks and payload instruments, the ground control structure and the field test site.

Primary objectives of our robotic recon test were to (1) evaluate the impact that robotic scouting has on traverse planning and crew productivity, (2) test our ground control, rover systems, operations and assessment protocols, and (3) understand the requirements on instrumentation and operations for a lunar surface reconnaissance robot.

K10 scouted two areas near the Black Point Lava Flow in Northern Arizona in June, to support a later crew traverse in the LER vehicle in September. The science team started with remote sensing data, planned a traverse for the LER, and identified major risks or uncertainties. K10 was remotely controlled to gather more information about these high recon priorities. This information was used to replan the traverses. The pre- and post-recon traverse plans were compared, and each of the plans was executed and evaluated. Below we describe these steps in greater detail.

A. Robotic Recon

The approach to the recon experiment includes four phases, shown in Figure 1. Phase 1 is Pre-Recon, including photogeologic interpretation of satellite imagery and initial traverse planning. The science team started with a QuickBird satellite image of the site at about 60 cm per pixel, analogous to what will be available from LRO. Photogeologic interpretation was done, and two one-day LER traverses were designed using only remote sensing. These two traverse plans are called the pre-recon plans, and designated W1 (for the West area) and N1 (for the North area)

Phase 2 is the Robotic Recon Mission, where the site is studied telerobotically. The traverse plans made using remote sensing were used to identify areas with high potential for operational risks such as steep slopes, rough terrain, steep drop-offs, as well as science operations questions, including which strata are exposed at the surface, whether contacts between units are visible or not, accessibility of units of interest, etc. The major outstanding questions from traverse planning become the top priority recon objectives. Recon robot plans are then generated to address these objectives. Real time operations of the robot are described below.

In Phase 3, Pre-Crew, the team updates their plans using information from recon. This phase of the field test was done with the robot in the field at Black Point Lava Flow during June of 2009, and ground control located at the

NASA Lunar Science Institute (NLSI). Scientists representing NLSI, the Field Exploration Analysis Team (FEAT), the Optimizing Science and Exploration Working Group (OSEWG) and the Lunar Exploration Analysis Group (LEAG) staffed the Science Operations Team. The flight control team included robotics engineers, computer scientists and flight controllers from JSC's Mission Operations Directorate (Codes DO and DX).

Output products from the robotic recon are analyzed to reduce the uncertainty and fill in gaps in photogeologic interpretation. In many cases, Lidar was able to provide detailed analysis of cliff heights, surface roughness, and slope angles. Panoramic camera images provided measurements of slopes, surface roughness, and helped to identify matching geologic units. PanCam images also clearly showed things that were instantly informative to the geologists but not visible in remote sensing data, including jointing in the basalts and cross-bedding in the sediments. Microscopic imager images provided highly detailed views of the surface materials, showing color, texture, albedo, and whether the materials were loose or consolidated.

Recon data was used to update the traverse plans. The updated plans were modified by introducing new EVA stations to sample or more closely evaluate things that were discovered during recon, by dropping EVA stations that were adequately observed by recon or which were identified through recon to be inaccessible or lower priority, and by moving or modifying EVA stations to change the tasks to be done based on better knowledge of the site. Crew briefings and science back room traverse plan documents for post-recon plans included detailed images of each EVA station to prepare the crew and support team for what was to be expected at each EVA location.

Phase 4 is the Crew Mission and implementation of the final traverse plan. In our experiment we compare the traverse plans, and compare the execution by simulated lunar traverse with a crew in the LER. Plans with recon data follow the flow above with all four phases. Plans without recon use only Phase 1 and Phase 4 (planning with remote sensing only, followed by execution). There were 4 traverse plans in all, designated W1, W2, N1, and N2, with W/N indicating West/North, and 1/2 indicating pre- or post-recon.

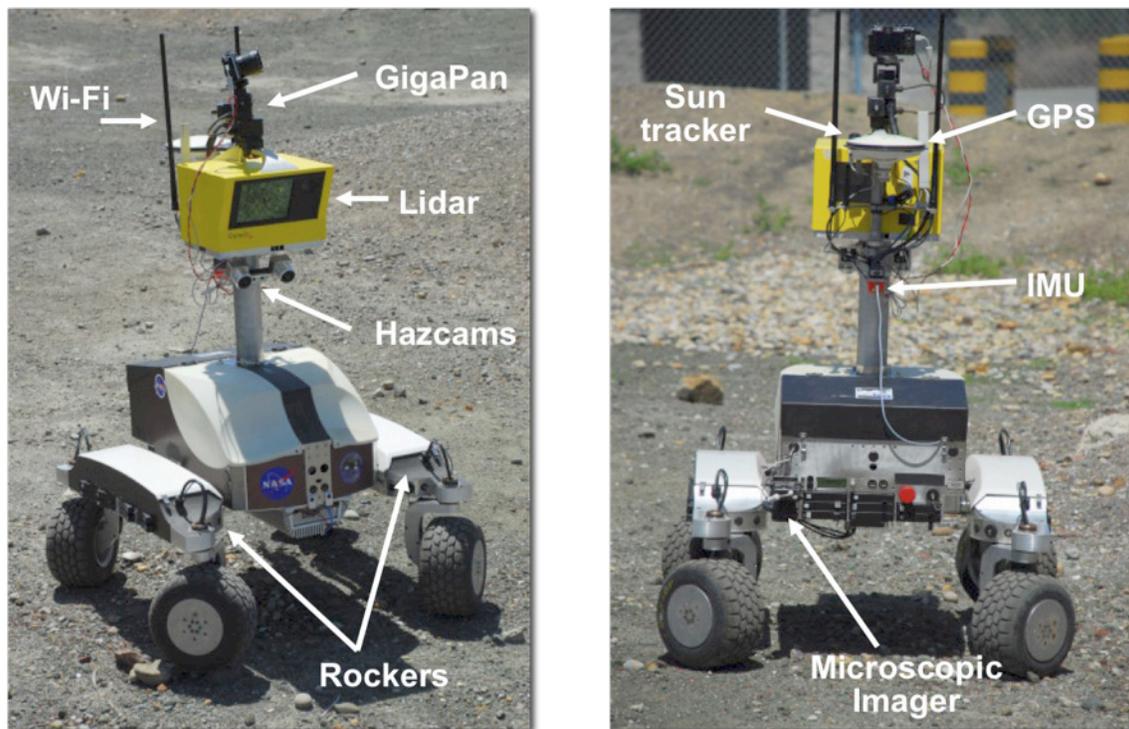


Figure 2. K10 Rover with recon instruments. **Left:** Wifi network connection, Gigapan camera for high resolution panoramic imaging, Lidar for 3D terrain scanning, HazCams for obstacle avoidance, and rockers for suspension and body averaging over rough terrain. **Right:** Sun tracker for orientation estimation, GPS for positioning, IMU for navigation, and Microscopic Imager (Mic) for high resolution terrain imaging and grain size/shape analysis.

B. The K10 robot platform

In our work, we use one of our third-generation “K10” planetary rovers (Figure 2). As a research platform, the design of the K10 emphasizes high flexibility with respect to target scenarios and payload instruments, and aims at providing an adequate operational analog fidelity as well as serviceability in a cost-effective manner. This is achieved through the use of inexpensive COTS parts and without requiring physical similarity to flight hardware.

Each K10 has four-wheel drive and all-wheel steering with a passive rocker suspension. This design allows operation on moderately rough natural terrain at human walking speeds (up to 90 cm/s). K10 has hard mounting points on the front, back, and bottom of the body as well as a 100 cm high mast. This allows attachment of additional components including antennas, navigational sensors, science instruments, and mechanical equipment.

K10’s avionics design emphasizes off-the-shelf components and modules. The robot is powered by twenty-four hot-swappable Inspired Energy 14.4V, 6.6 AH Li-Ion smart battery packs. K10’s controller runs on a Linux-based laptop and communicates via 802.11g wireless[3] or a Tropos unit for meshed wireless, respectively. K10’s standard sensor suite includes a Novatel differential GPS system, a Honeywell digital compass, an XSens accelerometer, Point Grey Research IEEE 1394 stereo cameras, DLP Design temperature sensors, a Hetronics or Omnex wireless e-stop, a suntracker, and wheel encoders.

The science payload of the robots as illustrated in Figure 2 was limited to non-contact instruments. Canon PowerShot consumer cameras were used as imaging sensors: The robot was also carrying a downward-pointing microscopic imager for surface grain size analysis, and a panoramic imager using GigaPan[3] technology to provide high-resolution panoramic views of the site. Additionally the robot was carrying an Optec 3D Lidar scanner[15], providing high resolution depth information with centimeter scale accuracy up to 500 m as well as terrain reflectance information.

C. K10 robot software

The K10 controller is based on our Service-Oriented Robotic Architecture (SORA)[6]. This architecture allows us to meet 3 critical system requirements: flexibility, scalability and reliability. Under SORA, rover controller functionalities are encapsulated as a set of services. Each service in this architecture provides specific functionality and exposes a clear high-level interface to the other services in a way that is transparent to network and programming language. The services interact using two distinct modalities depending on the requirements: remote method invocation and data distribution. Examples of major components include: locomotion, localization, navigation, and panorama acquisition. The system strongly relies on middleware that offers advanced functionalities while guaranteeing robustness.

Each subsystem is a service component that only relies on abstract (IDL-defined) interfaces to other components, which are resolved at service startup time. The SOA allows us to group the services into dynamic libraries that can be loaded and configured at run time. This reduces development time as well as memory requirements. The highly decoupled nature of the system also facilitates removal of deprecated services with minimal effort. The central components for the robotic recon activities were the payload instrument integration and the navigation system.

The payload instrument integration focused on usability, using instrument autonomy to provide a “point-and-shoot” model for recon data acquisition, while keeping sufficient parameterization accessible as user settings to satisfy scientists’ needs for trading acquisition time versus data resolution or quality. The panorama instrument was restricted to 5 different configuration with different fields of view and focal length settings, with acquisition times ranging from 2 minutes for a medium field of view, low zoom (low resolution) panorama, to 15 minutes for a wide field of view, high zoom (high resolution) panorama. The lidar images provided 2 operational modes, single shot and 360 panorama and 3 different resolutions. A single shot at low resolution taking 1 minute and a high-res 360 about 45 minutes. The microscopic imager did not provide any user settings, but would shoot a full-zoom image of the ground under the robot when triggered. The navigation system provided different operation modes, following the concept of sliding autonomy. K10 was capable of fully autonomous navigation with autonomous obstacle avoidance, driving and steering continuously at about 0.5 m/s. Different configurations of the obstacle detection system were available to drive through very rocky terrain (navigating near rocks up to 20 cm) or brushes at reduced speeds. At any point, ground operators could take over manual control of the robot to control K10 in position-mode, if the robot got stuck in rocky terrain or to bring the robot close to recon targets (e.g. the edge of a cliff) that would exceed the safety thresholds of autonomous terrain assessment.

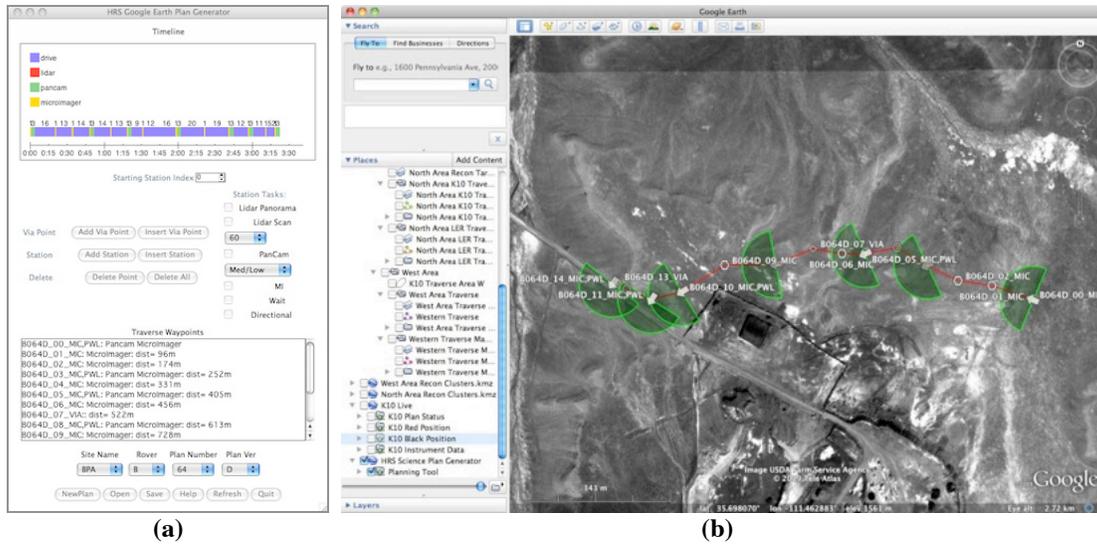


Figure 3. (a) The GEOPs control panel shows the plan timeline and task list. **(b)** Google Earth shows the rover path and directional waypoints with green wedges showing the planned instrument field of view.

D. Ground Data Systems

The ground control team has several software tools for both uplink and downlink. The uplink side is concerned with commanding the robot, essentially where to go and what to do. This is done by designating waypoints and data collection tasks. Waypoints are designated using Geops, shown in Figure 3, which is a lightweight custom user interface that uses Google Earth for its map display. Geops shows the task list as it is constructed, as well as an estimated timeline. Waypoints are designated manually and shown in the Google Earth map as they are added. The waypoints can be directional, so that a feature can be targeted by placing the rover at a good vantage point and then orienting toward the feature of interest, then capturing a panoramic image or a lidar scan. Uplink also includes low level direct teleoperation tools. These tools are used to maneuver the robot in contingencies, such as when the onboard obstacle avoidance will not allow the robot to move into an area that the operators judge to be safe.

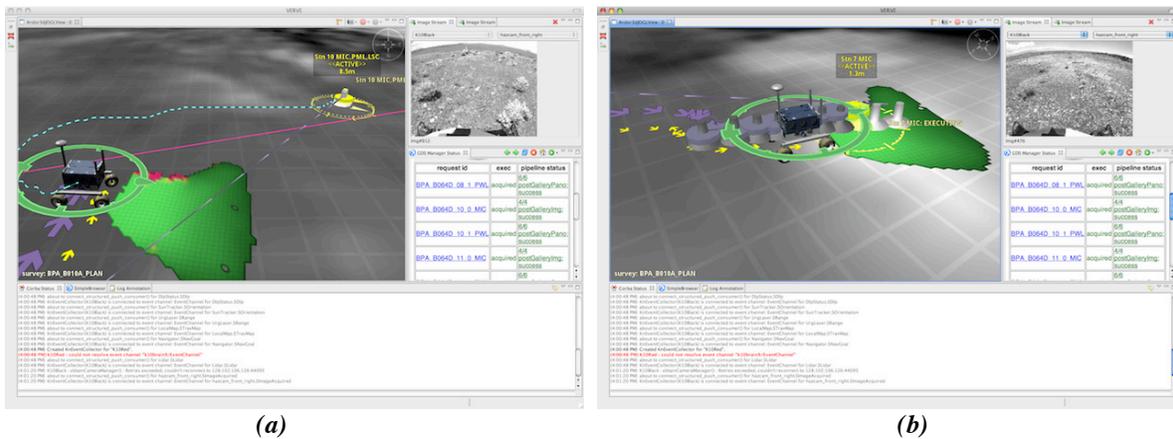


Figure 4: VERVE display. (a) The rover obstacle map, planned track (purple line), actual track (purple arrows) and computed track to target (dashed blue line, computed dynamically on rover to avoid obstacles). **(b)** Closely spaced MIC image targets (inverted “golf tee” icons), and a real-time view from the rover’s hazard camera in upper-right corner of window.

The downlink side monitors and displays rover activity, plan execution status, archives data products, and creates a live map of where the robot is, where it has been, and where it acquired data. The VERVE (Visual Environment for Robotic Virtual Exploration) virtual reality interface, shown in Figure 4, displays high-fidelity views of rover state, position, and plan status correctly registered over a satellite image basemap. It also shows 3D data from the onboard lidar as it is captured. In addition to the VERVE display, the rover track is output as a

continuously updated KML file which can be shown in Google Earth (GE). This means that the archived rover position and data product locations can be viewed in the same map that was used for planning and creating the rover track. This data can be viewed in real-time while a plan is executing. The data archive is deployed on a server running the open source

Gallery2 software, a MySQL database, and some custom software are used to index all of the image data collected by the robot. When an image is taken onboard, it is tagged with position and time, sent to the ground, and imported into the Gallery. An icon is placed in the GE map with a link to the Gallery page for the data product. The image data stored in Gallery2 is also tagged with additional metadata (e.g. rover name, instrument name, plan name, and station number) by our ground data software to aid in searching and organizing the archived data.

We use Google Earth extensively for K10 robot planning and operations. Google Earth is a very flexible map viewer, which has allowed us to put in lots of different kinds of content (image overlays, points and polygons, etc.). Google Earth provides a single operational view for looking at site data, robot plans, robot activities, and the data collected from K10's instruments.

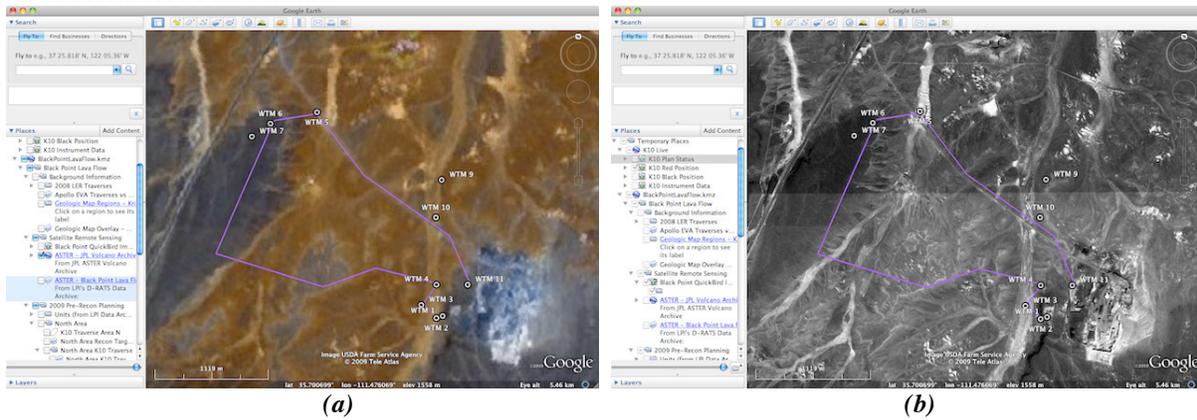


Figure 5. (a) Recon traverse path over ASTER image overlay. (b) Same path over QuickBird image.

Figure 5(a) shows recon traverse goals in the West area with an ASTER image overlay, which provides an indication of mineralogy. Figure 5(b) shows the same map with the greyscale QuickBird image (60 cm/pixel), which is similar to what LRO will provide. Google Earth allows scientists and controllers to easily switch, and blend, multiple data layers.

We also used Google Earth to develop recon traverse plans for K10: specifying waypoints and data collection activities directly on a map. During K10 operations, traverse plans varied significantly in scope, duration, and complexity. After the science operations team developed a plan, the flight control team vetted the plan (to verify operational constraints) and then executed it with the robot.



Figure 6. (a) Tracking K10 in real time, with the planned robot track in red, the actual robot track in black, and downlink data products shown with icons. Data products also include a preview balloon, with a link to the data product in Gallery. (b) Gallery view of the same data products.

While K10 was operating, we continuously tracked its location in Google Earth (Figure 6). As K10 acquired recon data, geo-registered placemarks were automatically added to the Google Earth display. Placemarks contained preview images as well as hyperlinks to the K10 Ground Data System, which allowed scientists to work directly with source data.

By the end of two weeks of operations, K10 had acquired more than 8.5 Gbytes of ground recon data. The map in Figure 6(a) shows all of the data in a single view: microscopic terrain images are shown as yellow "M" icons, 3D lidar scans are shown as pink "L" wedges, and GigaPan panoramas are indicated as green "P" wedges. This image shows that many microscopic terrain images were collected, which reflects the low-cost (in terms of operations time) of acquiring these images.

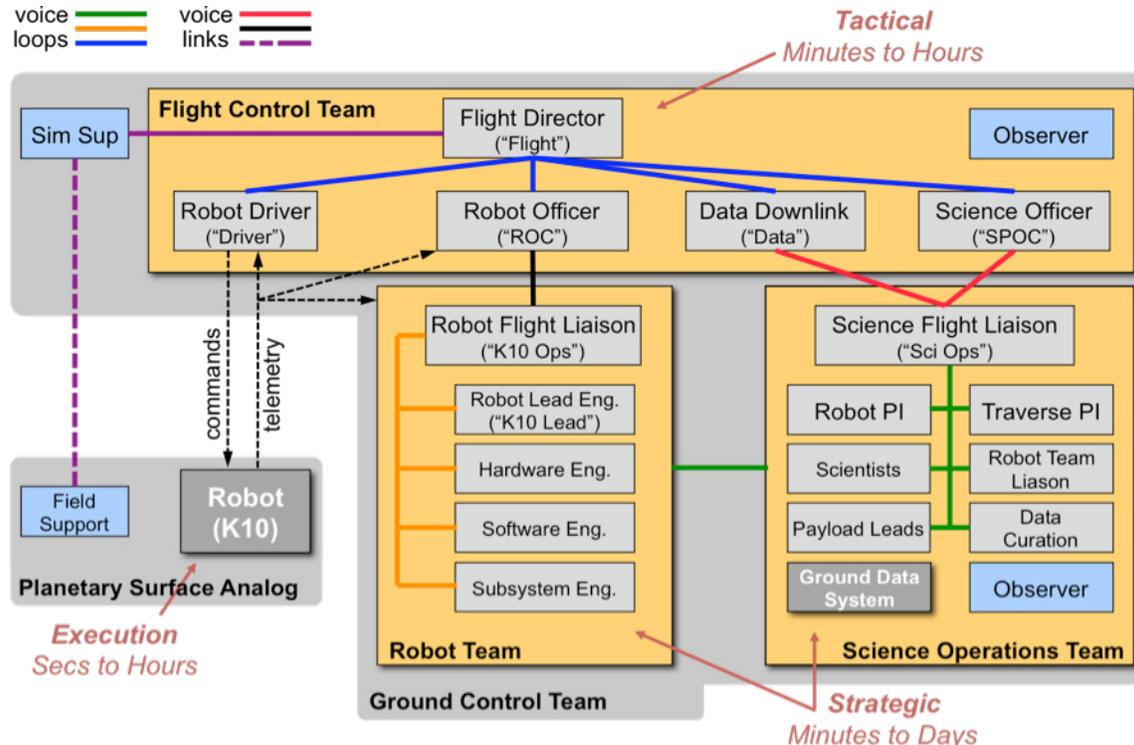


Figure 7. The ground control team structure blends Apollo, ISS, Shuttle, and MER ops concepts.

E. Ground Control Structure

The design of our ground control structure (Figure 7) draws inspiration from ground control used for Apollo, the Space Shuttle program, the ISS, the MER, and the planetary rover field tests that we have conducted during the past fifteen years. The operations team is split up into multiple "rooms". The vertical decomposition of the teams is based on the different time-scales on which these teams interact or respond to situations during operations.

Horizontally the decomposition is based on different tasks. Our recon scenario, with one robot and no crew, resembles only a small fraction of the early lunar campaign. We organized our operations into three sub-teams: *Flight Control*, *Science Backroom* and *Robot Backroom*. All real-time decision making and control was located in the Flight Control team. It conducts all commanding of and real-time response to data from the robotic assets.

During operations, the main task of the Science Backroom is to produce activity plans for the robot based on science objectives, a priori orbital imagery, and new data coming from the ongoing activity plan. The Science Backroom generates task lists for the robot spanning 30 to 60 minutes.

The Robot Backroom provides support for medium and long time-scale robotic issues. It responds to robot performance and robot health issues on an incident-oriented fashion. While the Flight Control team will address rover platform issues on an operational level, the Robot Backroom would be tasked with resolving such issues by updating or changing the software configuration of the robot.

After initial testing of the control setup during the Moses Lake field test, we refined the roles, assignments, and procedures for the Marscape test. Special emphasis was given to the communication setup. Each room shares a

communication channel (voice loop). All teams can listen-in on other teams' voice-loops. Communication between teams (especially upwards in the vertical decomposition) is only to be initiated through the designated officers/coordinators in the flight-control room. Once initiated, the conversation can be taken to the appropriate channel (e.g. instrument specialist from Science Backroom speaking on the Flight Control voice loop explaining unexpected instrument readings etc.) The goal of this scheme was to minimize distractions of the operations teams, while at the same time minimizing restrictions on the information flow, especially to and from the Flight Control team.

The Flight Control team has a hierarchical, chain of command-oriented setup. The reasoning is that it needs to be able to make critical decisions in real-time. All direct interaction with the robot has to be initiated by Flight Control. Our simulated Flight Control consists of five people. Flight Director ("Flight") is the executive in charge of ops decisions. The flight director coordinates and reviews information from the flight control team and passes decisions to robot ops team via the Robot Officer/Coordinator. Robot Driver ("Driver") directly operates robot mobility, also commands "real-time" science ops, such as panning cameras to look at specific targets as the robot is moving. Robot Officer/Coordinator ("ROC"). Is the representative of the Robot Backroom in Flight Control. He is the authority of the Robot Backroom for real-time decisions regarding robot operations. The ROC passes information from flight director to the Robot Backroom. This role—minimizing distraction of surface team—is similar to that of CapCom in human flight missions. ROC also has access to the same controls as the Driver. Downlink Lead ("Data") supervises data return from the robot. His task is the online monitoring of instrument data, ensuring that potential science instrument problems (de-calibration, parametrization) and communication problems (dropouts, etc.) are detected immediately. Science Protocol Officer/Coordinator ("SPOC") represents the Science Backroom to Flight Director. SPOC monitors science telemetry from robot, watches for interesting targets that require follow-up, and performs short-term planning. Example roles: prioritizing data products for downlink based on science team interest, suggesting follow-up observations on particular targets.

The Robot Backroom is designed primarily based on our own experiences with terrestrial robotic field testing over the past 15 years, tailored to the specific design of K10 and the expertise within our research group. Robot hardware, software, and payload specialists provide expertise on specific robot subsystems when issues are beyond the ROC's knowledge or ability to analyze or debug a situation in real time.

The Science Backroom design is influenced by the MER Science Operations Working Group[11]. The SOWG is responsible for generating plans for MER, where the planning cycle is done once per Mars Sol. Lunar near side operations are potentially continuous and low latency, so our tactical cycle is instead designed to be rapid and iterative throughout the day. The Science Backroom consists of individuals with technical expertise and workload focus on each of the recon instruments, as well as an overall Science PI and a Liaison between the Science Backroom and Flight. The Science Flight Liaison's job is to keep a real-time link between what is happening in the Science Backroom and the Flight Control Team so that the PI can focus on strategic, longer-term thinking and coordination of the team.

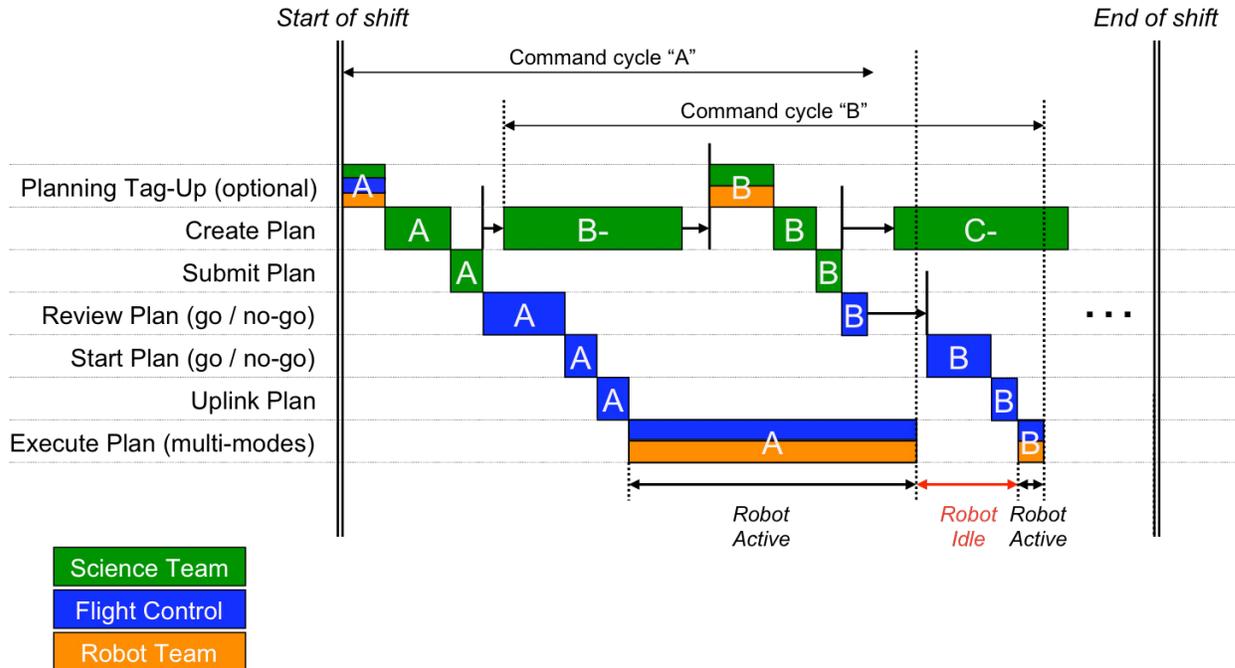


Figure 8. The goal of the tactical operations pipeline is to plan in parallel with robot activity so that the next plan is ready when the previous one completes.

F. Operational Timeline

In order to maximize robot operations time and get the most science return from our test setup, the ground control team would work on planning, execution, and maintenance in parallel, using a pipelined setup. Figure 8 illustrates the steps of the tactical operations cycle and the pipelining of planning and execution. Initially the Science Backroom creates an activity plan for the robot that balances the short-term science opportunities and the over-all science goals, consulting with the Robot Backroom (or a co-located representative) on the feasibility of their goals for robotic execution.

Afterwards the activity plan is submitted and briefed to Flight Control for approval and uploaded to the robot. While Flight Control is monitoring and managing the execution of the plan by the robot, the Science Backroom starts the next planning cycle. The Robot Backroom is assisting with robotic issues and has a time-slot at the end of the operations cycle to schedule maintenance tasks or to upload configuration changes to the robot.

This pipelined approach is designed to minimize the time during which the robot has finished executing a plan and the Science Backroom is still generating the next plan. The time between when the robot completes one plan and begins execution of the next plan is called “Robot Waiting”. Ideally this time would be zero, which maximizes the time utilization of the robot. This is not always practical. The Science Backroom is concurrently monitoring robot execution, analyzing and discussing data products that have been downloaded from the robot, and planning future activities. Often discoveries are made in new data products that influence the next set of objectives or tasks, so the planning process must take these into account.

G. Performance Monitoring Systems

Our approach to performance monitoring includes three different categories of analysis. The first is automated analysis of rover telemetry to compute performance metrics for the robot. This includes daily distances driven and the number and type of data products acquired (PanCam images, Microscopic Imager images, and Lidar scans).

The second is software-based real time analysis of operations to capture interactions between the robot and the human team, including the duration of communications drop-outs, time spent generating plans, execution time for plans vs. the expected time, etc. These metrics show us where we can improve the interactions.

The third is activity logging by human observers. This analysis captures how the remote operations team spends their time throughout the day. Activity types are logged manually by human observers sitting in each of the ground control team rooms. Results from the first two types of analysis are presented below. Results from observer teams will be presented in another paper which focuses on human factors analysis.

III. Results

Below we present the performance of the reconnaissance robot and ground control team as a system for collecting recon data from the Black Point Lava Flow test site between June 15th and June 26th 2009. Analysis of the recon data and analysis of the impact that recon has on planning and executing crew traverses is still ongoing and will be presented elsewhere.

Rover Performance Evaluation

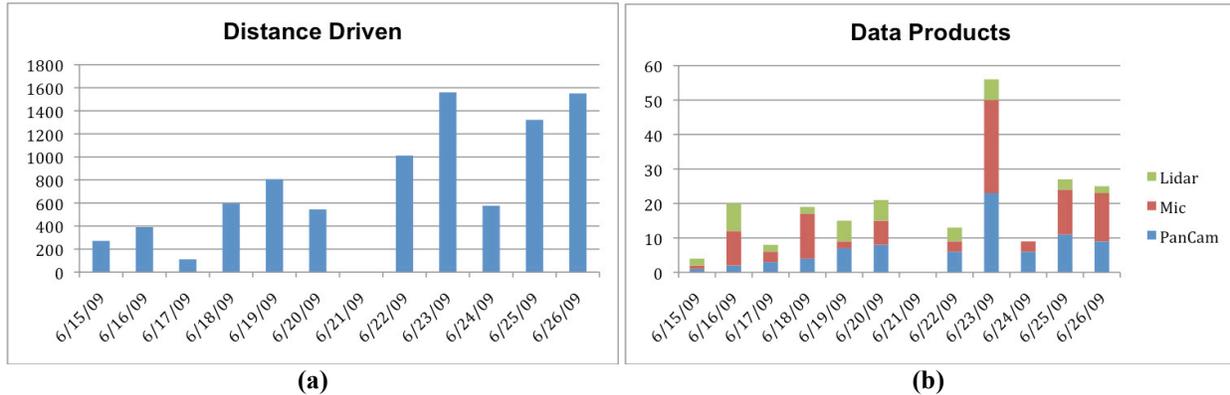


Figure 9. (a) Distances driven during each day of operations. (b) Data products captured during each day.

Figure 9(a) shows the distances that K10 covered during each day of operations. We did not operate the robot on Sunday June 21st. In general we covered greater distances during the second week of operations, but this is primarily due to operational issues including intermittent communications, logistical difficulties for the field team operating in two different regions of the field site, differences in reconnaissance objectives in the two regions, and the overall comfort level of the team operating the robot improving in the second week. Notably, K10 is capable of traversing several kilometers per day, as has been demonstrated in previous field testing [7].

Distance is not a critical metric for evaluating recon performance. Recon objectives may or may not be separated by large distances, particularly for a given day of operations. If a recon robot is to be operated for several months in order to study an area of that a crew will visit over the course of a two week surface mission, then in one day of robot operations, the robot only needs to cover a small fraction of what the crew will cover in one day of crew operations. Figure 9(b) shows the number of data products collected per day, which is an indication of how many detailed measurements were made in the area covered by the robot. This is a better indication of how much better a remote team will understand a site through the detailed analysis of recon data. Again differences from day to day, and between the first and second week of operations, are confounded by many issues beyond the scope of the intended test. Intermittent communications, logistical difficulties in the field, and the overall comfort level of the remote operations team are a few.

During robotic reconnaissance we monitored rover telemetry and computed performance metrics in real-time. Flight controllers used these metrics to gain insight into whether the rover was performing as expected [17]. These metrics also can provide insight into human performance during recon operations. Metrics for rover productivity can reveal the effectiveness of the human team in producing and supervising rover plans. Metrics computed for rover reliability indicate the time spent by the human team performing unplanned interventions in rover operations. Metrics on loss of communication with the rover indicate the availability of data for remote situational awareness as well as the opportunity to uplink new plans or tele-operate the rover. Below we describe our observations on human and robot team performance derived from rover metrics.

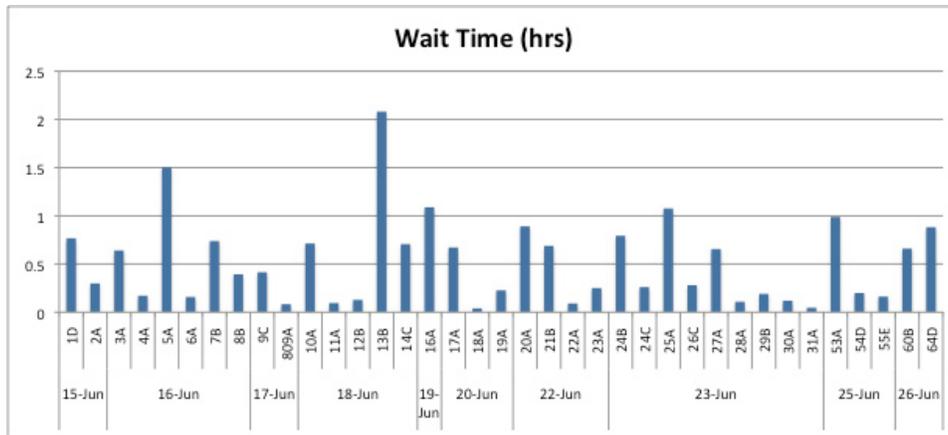


Figure 10. Time that the robot waited for each plan during the field test.

To maximize the data return, the science team tried to minimize the time the rover spent waiting for the next plan once it had completed its current plan. Thus, plan wait time indicates the efficiency of the science team in producing recon plans (see Figure 10). As mentioned above, if the tactical planning timeline is efficient, then there is zero wait time because the next robot plan is ready to execute when the current plan finishes. In practice, the average plan wait time during our experiment was 31 minutes. The team required more time - 45 minutes to 1 hr - each day to produce the first plan. Otherwise, wait times exceeding half an hour often were due either to the team eating lunch or to the rover being moved manually to a new location. The longest wait was 2 hours, for Plan 13B, during which time a flat tire on the rover was fixed before it was manually moved. The rover also was manually moved while the science team was planning 21B, 25A, and 64D. The rover also waited for a plan while the team ate lunch (typically 30-45 min), which affected the wait times for plans 5A, 13B, and 27A. Removing the six wait times that included either moving the rover or lunch, the average wait time was 24 minutes.



Figure 11. Ratio of actual plan time to estimated plan time.

The Flight Operations Team supervised the rover executing recon plans. This included uploading a plan to the robot, initiating plan execution, pausing automated task execution when the robot had problems completing a task (e.g., difficulty navigating rough terrain), tele-operating the robot out of areas where it had problems, and aborting robot operations when it was no longer possible to complete a plan. The flight team also was responsible to monitor for potential conflicts in plan execution. This included detecting when the allocated time for a plan was insufficient to complete all tasks in the plan or when adequate resources (specifically battery power) were not available to perform all tasks in the plan.

During the field test we computed whether a plan completed successfully or not, and how long the rover spent executing the plan. This gives us insight into how well the flight operations team did at keeping the rover “on plan”. A total of 37 plans were uploaded to the rover. Of these plans, 17 were partially completed and 20 were completed successfully. The large number of incomplete plans primarily reflects the policy to push the rover’s limits to acquire data from difficult sites (i.e., try to get data, abort if not possible). Second, because the science team could review

data soon after it was collected, it was possible to assess whether remaining tasks were needed or could be dropped or revised into a new plan. The ability of the science team to reactively re-plan made this protocol possible.

Figure 11 shows the ratio of actual time on plan to estimated time on plan for the plans that went to completion. Of the 20 plans that were completed, 16 plans were completed within 10% of the expected time.

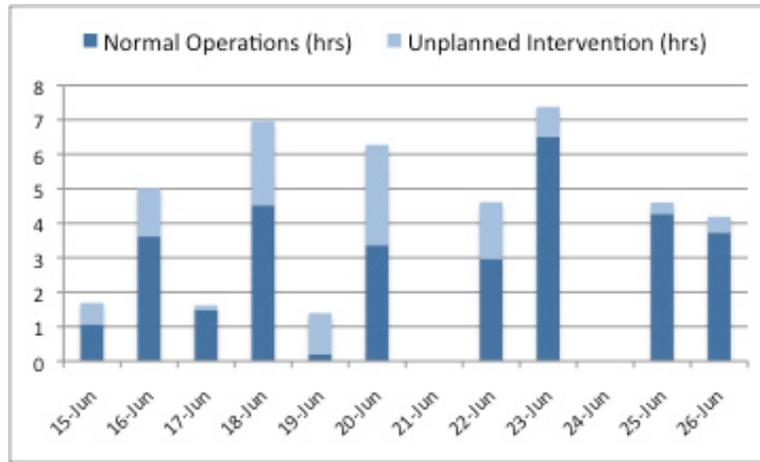


Figure 12. Time Spent in Error States Requiring Unplanned Intervention

For lunar reconnaissance operations, minimizing all human interaction time may not translate to more efficient reconnaissance operations. In fact, it may often be more time and resource efficient to tele-operate the robot in difficult terrain than to operate autonomously. Thus, our objective was to minimize the time spent on *unplanned human intervention*. For this test, we detect unplanned intervention as the time the rover spent in an anomalous state where a person typically was required to intervene. Anomalous states detected during this field test include the following: (1) emergency stop, (2) locomotor failure, (3) navigator failure, (4) failure of the plan executor, (5) position error, and (6) joint failure. Twelve hours total were spent during the test on unplanned interventions. Most of this time was spent with the rover emergency stopped, although the rover experienced position errors (often caused by a stuck wheel) for nearly 20 minutes total. On an average, 1.18 hours were spent per day on unplanned intervention, with a minimum time per day of 8 minutes on June 17 and a maximum time per day of 2.75 hours on June 20.

Figure 12 shows the time spent in unplanned intervention each day compared to the time spent in normal operations for that day. From these statistics, we see that performance improved significantly in the second week of operations. Part of this was due to better communication with the rover (discussed in the next section) and part was due to the protocols being more familiar as the team gained operational experience. On the 19th, the operations period was abbreviated because communication with the rover was significantly degraded. Significant time was spent in emergency stop waiting to see if communication quality improved. On the 20th, the rover operated for a normal time period, but a combination of poor communication quality and bad weather resulted in significant time in emergency stop.

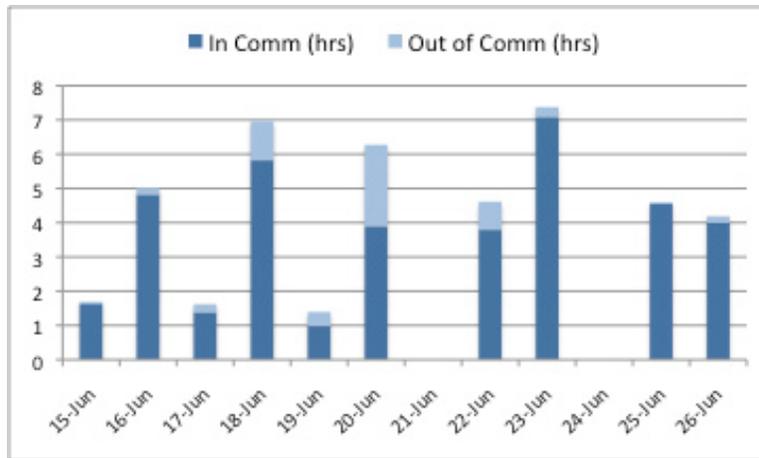


Figure 13. Time In and Out of Communication with Rover

We also measured the time spent out of communication with the rover (called Loss of Signal or LOS) as a dropout of all robot data for at least 10 seconds. Statistics on LOS provide us with insight into the human team’s ability to command the rover and to maintain current situation awareness of rover activities. The total time without rover communication for the field test was 5.77 hours or 13.2% of operations time. The quality of communication with the rover varied significantly over the course of the field test, ranging from 1% to 38% of the day in LOS. We illustrate this in Figure 13, which show the time in communication compared to time out of communication with the rover for each day of the field test. We also can see in this figure that communication was worse during the middle of the field test. This corresponds to the rover operating in a remote area of the test site. On these days with reduced quality communication, we also see worse performance on plan execution (see Figure 11). On June 19, there were no plans taken to completion. On June 20 we had 2 plans taken to completion, but it took at least twice as long as expected to complete them. With improved communication from June 23 – 26, all but one plan were either completed or aborted well within the estimated time.

IV. Discussion

When the new lunar surface campaign begins, it will be important to make use of the considerable technological advances that have been made since the end of the Apollo program in 1972. Robotic recon has the potential to significantly improve science return from the lunar surface. In particular, we have confirmed that recon can significantly aid plan prioritization, supplement and complement remote sensing, and assess hazards and routes to reduce operational risk. The concept of robotic recon is simple: the more information you have, the better you can plan. However, several important questions still need to be answered.

What should be the operations concept for lunar robots? How do the capabilities and operations of robotic rovers need to be changed from current practice (“robot as science instrument”) to be used for recon (“robot as scout”)? What ground control structure is needed to support scouting activities?

What surface mobility system should be used for scouting? The LER could be used in an unmanned mode. However, this presents a tradeoff: improving understanding of a site prior to crew activity vs. risking damage to the rover before crew arrive. Smaller robots could be used instead, but they may not have sufficient power for long-range operations.

What are the required relationships between crew and recon mobility? In particular, does the recon robot need to have the same performance as a crew rover? Recon may not need to follow the same route as crew (e.g., assessing the descent route into a crater might best be done from an opposing viewpoint), so terrain performance might not need to be the same. Also, there may be significantly greater time for robot missions. Thus, ground speed might not need to be comparable.

How should recon data be processed and presented? The design of tools, displays, and protocols all impact the efficiency of science operations. Given that analysis, decision making, and plan generation can be slow, a key question is: How can we reduce the bottleneck of viewing and analyzing recon data?

How should recon data be logged and georeferenced? Sharing position information among different exploration assets (humans, robots, orbiters) might require absolute positioning, or fixed references. Terrain relative navigation may work for individual visits, but if a sample is identified during recon, a later crew will need to be able to find the exact spot to collect it.

What is the most effective way to coordinate human-robot activity? How can robotic recon data be most rapidly and effectively incorporated into the planning (or replanning) of a crew traverse or an EVA? What scouting data need to be presented to crews in training and during a mission? How and when should this data be conveyed or made available?

Acknowledgments

We thank David Morrison and the staff of the NASA Lunar Science Institute for their support during the test. We also thank Mary Sue Bell, Mike Downs, Jim Dumoulin, Marwan Hussein, Frank Jurgens, Tim Kennedy, David Kring, Nina Lanza, Mike Miller, Ann Ollila, Jim Rice, Marc Seibert, Art Snoke, Robert Stewart, Eugene Stoccardo, Jeff Tripp, Mike Wyatt, Aileen Yingst, and Kelsey Young for supporting the robotic recon mission simulation.

This work was supported by the NASA Exploration Systems Technology Development Program under the Human-Robotic Systems Project and by the NASA ESMD Directorate Integration Office Analogs Program.

References

- ¹ M. Broxton, et al. (2008) "Traverse planning for human and robotic missions to Hadley Rille", Tech. rep. NASA Ames.
- ² M. Bualat, L. Edwards, T.W. Fong, M. Broxton, L. Flueckiger, S.Y. Lee, E. Park, V. To, H. Utz, V. Verma, C. Kunz, and M. MacMahon. "Autonomous Robotic Inspection for Lunar Surface Operations." *Field and Service Robotics Results of the 6th International Conference, volume 42 of Springer Tracts in Advanced Robotics*, pp 169–178. Springer - Verlag, Berlin / Heidelberg, June 2008.
- ³ M. Bualat, L. Kobayashi, S.Y. Lee, and E. Park., "Flexible Rover Architecture for Science Instrument Integration and Testing." *Space. AIAA*, 2006.
- ⁴ E. Heggy, T. Fong, D. Kring, M. Deans, A. Anglade, K. Mahiouz, M. Bualat, P. Lee, and F Horz., "Potential of probing the lunar regolith using rover-mounted ground penetrating radar: Moses lake dune field analog study." *Proceedings of the 40th Lunar and Planetary Science Conference*, Houston, TX, 2009.
- ⁵ R. Elphic, H. Utz, M.G. Bualat, M. Deans, T.W. Fong, L. Kobayashi, S. Lee, and V. To., "Preliminary results of hydrogen prospecting with a planetary rover." *39th Lunar and Planetary Science Conference, number 2400*. Lunar and Planetary Institute, March 2008.
- ⁶ L. Fluckiger, V. To, and H. Utz, "Service-Oriented Robotic Architecture Supporting a Lunar Analog Test." In *International Symposium on Artificial Intelligence, Robotics, and Automation in Space (iSAIRAS)*, 2008.
- ⁷ T. Fong, M. Bualat, M. Deans, M. Allan, X. Bouyssounouse, M. Broxton, L. Edwards, R. Elphic, L. Fluckiger, J. Frank, L. Keely, L. Kobayashi, S. Y. Lee, P. Lee, D. Lees, E. Pacis, E. Park, L. Pedersen, D. Schreckenghost, T. Smith, V. To, and H. Utz. "Field Testing of Utility Robots for Lunar Surface Operations." In *Proceedings of AIAA Space 2008*, Number AIAA-2008-7886, San Diego, CA, 2008. AIAA.
- ⁸ T. Fong, M. Deans, et al. "A Preliminary Examination of Science Backroom Roles and Activities for Robotic Lunar Surface Science." In *Lunar Science Conference*, Number 2142. NASA Lunar Science Institute, 2008.
- ⁹ T. Fong, M. Allan, X. Bouyssounouse, M. G. Bualat, J. Croteau, M. C. Deans, L. Flueckiger, S. Y. Lee, D. Lees, L. Keely, V. To, and H. Utz. "Robotic Site Survey at Haughton Crater." In *Proceedings of the 9th Symposium on Artificial Intelligence, Robotics, and Automation in Space (iSAIRAS 2008)*, Los Angeles, California, February 2008.
- ¹⁰ D. A. Harrison, R. Ambrose, B. Bluethman, and L. Junkin. "Next generation Rover for Lunar Exploration." In *Proceedings of the IEEE Aerospace Conference*. IEEE, March 2008.
- ¹¹ A. Mishkin, Y. Lee, D. Korth, and T. LeBlanc, "Human-Robotic Missions to the Moon and Mars: Operations Design Implications." In *Proceedings of the IEEE Aerospace Conference*, 2007.
- ¹² NASA Advisory Council. "Workshop on Science Associated with the Lunar Exploration Architecture, Final Report and Recommendations." *Technical Report NP-2008-08-542-HQ*, NASA, 2008.
- ¹³ National Research Council. "The Scientific Context for Exploration of the Moon: Final Report." *Technical report, Space Studies Board*, 2007.

¹⁴ J. Osborn. “The Role of the Science Officer Flight Controller in the Upcoming Era of Lunar Exploration.” *Technical report, NASA Johnson Space Center*, 2006.

¹⁵ G.R. Osinski, T. Barfoot, N. Ghafoor, P. Jsiobedski, J. Tripp, R. Richards, T. Haltigin, N. Banerjee, M. Izawa, and S. Auclair. “Optimizing Lunar Surface Activities: Lidar and MsM as Scientific Tools” *Joint Annual Meeting of LEAG-ICEUM-SRR*, Number 4061, 2008.

¹⁶ Schaber G. “The U.S. Geological Survey, Branch of Astrogeology – A Chronology of Activities from Conception Through the End of Project Apollo (1960-1973).” *Technical Report 2005-1190*, USGS, 2005.

¹⁷ D. Schreckenghost, T. Fong, T. Milam, E. Pacis, and H. Utz. “Real-time assessment of robot performance during remote exploration operations.” *Proceedings of the IEEE Aerospace Conference*. IEEE, March 2009.

¹⁸ D. Schreckenghost, D., T. Fong, T. Milam, and H. Utz. “Measuring Robot Performance in Real-time for NASA Robotic Reconnaissance Operations.” *NIST PerMIS Workshop*. Sep 2009.