

Productivity Challenges for Mars Rover Operations: A Case Study of Mars Science Laboratory Operations

**Daniel Gaines, Gary Doran, Heather Justice, Gregg Rabideau,
Steve Schaffer, Vandana Verma, Kiri Wagstaff, Ashwin Vasavada,
William Huffman, Robert Anderson, Ryan Mackey, Tara Estlin**

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109
`{firstname.lastname}@jpl.nasa.gov`

Abstract

Achieving consistently high levels of productivity for surface exploration missions has been a challenge for Mars missions. While the rovers have made major discoveries and accomplished a large number of objectives, they often require a great deal of effort from the operations teams and achieving objectives can take longer than anticipated. In order to better understand the productivity challenges faced by surface missions, we conducted a case study of selected MSL science campaigns. For each campaign, we examined the activity performed by the operations team, how the team decided what to accomplish each day, how well these objectives were achieved and how results from one day fed into and informed objectives for the next and how the team allocated vehicle resources during the campaigns. We also conducted a series of interviews with operations personnel to get their broader perspective on surface mission productivity factors. This document presents the results of the case study. We begin with an overview of MSL operations to provide context for the study followed by a description of the design and methodology for the case study. We then present results beginning with our interviews with operations personnel. Next, we present each of the selected science campaigns, describing the science objectives, descriptions of the science and engineering activities performed during the campaign and an analysis of the resource allocations, significant productivity factors and significant ground-in-the-loop decisions. We then discuss comparisons among the selected case study. As a result of the case study, we have identified several factors that pose challenges to surface mission productivity. Some of these challenges may become an even larger liability on future missions due to the potential need to rely on non-synchronous relay orbiters. Our study also showed that there is available vehicle resources to provide an opportunity to increase productivity if these challenges can be overcome.

Executive Summary

Achieving consistently high levels of productivity for surface exploration missions has been a challenge for Mars missions. While the rovers have made major discoveries and accomplished a large number of objectives, they often require a great deal of effort from the operations teams and achieving objectives can take longer than anticipated. We are in the early stages of a multi-year project to investigate solutions for enhancing surface mission productivity. A primary focus of this early stage is to conduct in-depth studies of Mars Science Laboratory science campaigns to gain a deeper understanding of the factors that impact productivity, and to use this understanding to identify potential changes to flight software and ground operations practices to increase productivity.

This report documents the MSL productivity case study. The report begins with background on the MSL mission and an overview of how MSL operations are conducted. Next, we describe the design of the case study which included interviews with MSL operations personnel and detailed studies of specific mission campaigns. This section includes a conceptual model of how surface missions achieve objectives that we developed to guide our study. We then present the results of the case study beginning with a description of the interviews, followed by a detailed discussion of each of the selected science campaigns. We then discuss comparisons among the campaigns we selected for study. We conclude the report with a summary of the significant productivity factors identified in the study and thoughts for future directions of inquiry.

Introduction

The Curiosity rover is a roughly car-sized, six-wheeled rover with a mass of 899 kilograms (1,982 pounds). It is powered by a multi-mission radioisotope thermoelectric generator (MMRTG) that converts heat from radioactive decay into electrical energy. The rover has radios that allow the rover to communicate directly with Earth or via relay through Mars orbiting spacecraft. The rover contains a rich and varied suite of science instruments including several instruments mounted on a Remote Sensing Mast (RSM) and on a five-degree-of-freedom arm. A percussive drill mounted on the rover's arm acquires rock samples which can be delivered to a pair of sample analysis instruments within the rover.

MSL has developed a tiered approach for operating this complex rover: Strategic, Supratactical and Tactical operations. Strategic planning focuses on developing long-term plans, typically spanning weeks or months, to achieve high-level objectives. These strategic plans provide vital guidance in achieving mission objectives, but they are not directly executable. They must be adapted to take into account the current conditions and adjusted to respond to unanticipated conditions as the rover explores the environment.

The supratactical stage provides a bridge between the long-term strategic plan and the day-to-day, highly reactive tactical process. The process is designed to coordinate the complex science instruments and manage the constraints and resources required to conduct campaigns. The supratactical process produces "look-ahead plans" which span several sols, typically a week, of activity. These plans help coordinate the use of vehicle resources across multi-sol activities and they also help to coordinate the large science team spread across the globe. The process feeds into the tactical process by delivering a "skeleton" plan for each sol of tactical planning. The skeleton provides the tactical team with the major objectives for the plan. It includes a rough structure of the activities, including coordination of science activities around communications windows and other engineering activities and guidelines on how much resources, such as energy, time and available data volume, can be expended during the execution of the plan.

The tactical planning process forms the highly reactive phase of surface operations. It includes an assessment of the state of the vehicle and the performance of the previous plan's activities. During activity development stage, specific science and engineering objectives are identified based on the high-level objectives of the current campaign and guidelines provided in the skeleton plan. The plan may cover a single sol of the rover's activity, or multiple sols, referred to as multi-sol plans, in cases such as weekends, when the team wishes to generate command products that cover multiple sols in a single tactical shift. The developed activity plan is translated into sequencing command products to be executed on the vehicle. These command products are verified, reviewed and delivered for uplink to the rover. The tactical timeline is organized around several meetings which are used for team coordination and to review the evolving command products.

A typical sol in the life of the rover consists of a morning communication window in which the ground operations team uplinks the command products that will control the rover's activity for the sol, or multiple sols for multi-sol plans. The rover's sol will typically include 3 to 5 additional communication windows with relay orbiters in which it will downlink collected data. Given the relative timing between Mars and Earth, along with orbiter relay latencies, not

all of the data from these communication passes will be available to the team by the start of the next tactical shift. The last pass for which data will be available is referred to as the **decisional pass** because this is the last pass in the plan for which data can be used to make decisions during the next tactical shift. When developing the tactical plan, the team must carefully consider the timing of activities and the amount of data that will be collected to ensure that activities that will produce decisional data are performed prior to the decisional pass and that the volume of decisional data acquired will fit within the predicted downlink capacity of the passes leading up to and including the decisional pass. Further, the team must avoid conducting activities following the decisional pass that would compromise their ability to make use of the decisional data, for example, by driving away from the location the rover was in at the decisional pass.

The team typically organizes groups of related activities into **blocks**. For example, a typical drive sol will include a pre-drive science block, a drive block and a post-drive imaging block. Each block of activity includes margin, in case activities take longer than predicted, and cleanup, in case activities exceed allocated margin. The rover often performs multiple activities in parallel. It is common that heating in preparation for future activity will be conducted in parallel with other activity. Similarly, the rover will typically perform background environmental observations through much of the rover's diurnal cycle.

To conserve energy, and allow opportunity for the rover's batteries to be recharged, each sol typically includes several **sleep** periods in which the rover's computer is powered off. Some of the rover's instruments are capable of remaining active, collecting data, while the rover is asleep. The operations team decides which instruments to leave on while the rover is sleeping and schedules activities to collect the accumulated data when the rover is awake.

Case Study Design

We designed and conducted an exploratory case study to understand the factors that contribute to and detract from surface mission productivity. The diverse and dynamic nature of the mission posed a challenge for the design of our case study. We wanted to obtain an in-depth understanding of the productivity factors that affect surface operations over a broad range of mission and environmental conditions. In order to manage the scope of the study, we chose to use select science campaigns and perform a detailed study of the team's day-to-day activities during this time. We conducted interviews with operational personnel to contribute to the broad perspective of productivity.

To manage the scope of our research project, we decided to primarily focus on remote sensing and navigation operations. As such, our selection of case study campaigns was biased toward those that emphasized these types of activities. There are two major types of campaigns that MSL conducts that emphasize remote sensing and navigation: characterizing a geographical area and driving to a distant, strategically identified location. We selected three campaigns, two of which had the objective of exploring a geographical area and a third with the objective of performing a strategic drive. While the primary focus was on remote sensing and driving campaigns, we are also interested in addressing challenges of other types of activity. As such, the selected campaigns also included activities with the rover's arm-mounted instruments.

The following table summarizes the cases we selected for the study:

Campaign	Science Objective	Martian Season	Sols	Number of Sols	Drive Distance
Pahrump Hills Walkabout Pass 1	Characterize the Pahrump formation, making up the Mount Sharp basal layer	Spring	780 - 798	19	152m
Artist's Drive	Drive to higher levels of Mount Sharp, characterizing stratigraphy along the way	Late Summer	949 - 972	24	567m
Marias Pass	Characterize the contact between Murray and Stimson formations	Late Summer, Early Fall	991 - 997; 1027 - 1043	24	130m

Summary of selected science campaigns.

To guide the case study, we developed a conceptual model for how the operations team accomplishes objectives on a surface mission.

The general flow of the model begins with the team identifying candidate activities that can be used to accomplish their intent. These activities are developed and refined during operations planning until a set of command products is ready to be uplinked to the vehicle. The vehicle executes these activities and produces results which are conveyed back to Earth through telemetry and data products. This information, in turn, is used to support the development of subsequent activities and, potentially, new intent. Throughout this process, various factors can result in limitations or loss of activities. During operations planning, this can include restricting the scope of an activity, deferring an activity to a later planning day or even descoping an activity entirely. Reasons for these changes include factors such as timeline complexity, prediction of available vehicle resources and the availability of data from the rover needed to make decisions. During execution, loss of activity can include partial or complete failure of an activity.

An important concept related to availability of data is the notion of **restricted sols**. Restricted sols occur when, due to the relative timing between Mars and Earth, data from the rover's end-of-day communication pass is not available to the operations team until late in the day on Earth. Unless the team chooses to work during the night on Earth, they will not have sufficient time to develop a new tactical plan in time for the rover's next uplink. Instead, the team will create a plan to send to the rover on the following day. These sols are referred to as restricted sols because the team is limited in what they allow the rover to do on the sols between the time decisional data is sent and the next plan is received, in order to avoid the rover taking action that would compromise the team's use of the decisional data.

We used the conceptual model to develop a schema to guide the collection of data for the study. We collected data from the mission's activity plans, operations reports and vehicle telemetry. Data was organized into spreadsheets according to the schema. We also included a free-form section to enable the collection of data that did not follow the schema.

Our study includes some important caveats. MSL is a vast and varied mission. While we believe we have selected campaigns that reflect issues common throughout the mission and are relevant to future surface missions, the campaigns certainly do not capture all aspects of the mission. In particular, we did not include sampling campaigns in our study. This would be a valuable area to consider in future extensions of this work.

One of the biggest challenges we faced in the data collection process was determining the intent for a given activity. Operations personnel were generally very good about documenting in MSL Reports the specific reason an activity was included in the plan. However, it was not always clear how this rationale related to specific campaign objectives. For our objective of understanding productivity factors affecting achieving objectives, we concluded that it was sufficient to perform a "high level" categorization of activity intent. As such, we defined three high level intent categories for activities: Campaign Science, activities that directly contributed toward the campaign's objectives; Other Science, activities that contributed to science objectives unrelated to the campaign under study; and Engineering, activities carried out primarily for monitoring and maintaining the health of the vehicle.

Another important limitation of the study is that our information about what activities were considered during a given planning session is limited to what was documented in MSL Reports. While the reports provide information about activities that were part of an earlier plan but removed before uplink, and they often document activities that were considered but not included in a plan, they do not document all activities that the team considered. By this point in the mission the team had developed significant experience on the activity that could be included in a plan due to complexity limits and activity constraints. As such, there is a significant amount of activity pruning that is never documented.

Interviews with Operations Personnel

The objectives of the interviews was to gain a broad perspective on mission productivity that would complement the more detailed investigation of productivity in our campaign studies. Toward that end, we selected members of the MSL science and engineering operations team that performed roles that had a strong connection to productivity: Scientists, Rover Planners, Tactical Uplink Leads and Supratactical Uplink Leads. The interviews were structured as guided conversations. We conducted a total of thirteen separate interviews, each interview running between 45 and 90 minutes. We selected 5 members of the science team (three with a geo-chemical emphasis and two with an environmental emphasis), 3 Rover Planners, 1 Tactical Uplink Lead, 1 Supratactical Uplink Lead and 3 team members who performed both Supratactical Uplink Lead and Tactical Uplink Lead roles.

We began by asking participants to describe what they consider mission productivity to be. The majority of participants described mission productivity as a measure of how much science is accomplished each sol, or how many sols it takes to achieve science objectives. This description is consistent with the definition of productivity as a measure

of output per unit of input with sols taken as an aggregate measure of input. Two participants described productivity in terms of accomplishing objectives better, faster, and with less “churn”, which is related to efficiency, or a measure of the amount of input required to produce output. Productivity and efficiency are closely related terms, with productivity defined as an average measure of the efficiency of production. One participant noted the importance of taking into account team longevity in determining sustainable levels of mission output.

We then asked participants to describe the factors that contribute toward and detract from productivity. Following is a summary of some of their responses, in no particular order:

Ability to Exploit Supratactical Work: If the tactical team follows the supratactical plan, they are able to leverage work performed on the supratactical timeline and productivity benefits. If decisional data prompts the tactical team to change course, they may need to start from scratch and are often not able to accomplish as much.

Overhead in Command Product Generation: The effort involved in developing and validating command products can limit the amount of activity that can be accomplished.

Managing Tactical Timeline Complexity: Both the supratactical and tactical teams must judge how much activity can be managed within the tactical timeline. The ability to accurately make this judgment can impact productivity.

Predicting Available Vehicle Resources: Inaccuracies in resource modeling, including activity power and duration requirements, can result in unnecessarily restricting planned activity.

Interpersonal Communication: Effective communication was identified as a significant productivity factor by many of the participants in each of the operations roles we interviewed. This includes communication between the science and engineering teams and among roles within science and engineering teams.

Science Team Engagement: It is important for science team members to be aware of the current and past context of the mission to make informed science decisions. A significant challenge to engagement is that many team members work on the mission part time.

Restricted Sols: Restricted sols pose a significant productivity challenge when the rover is driving. The team requires post-drive data in order to perform targeted observations and to effectively plan the next drive.

Time to Analyze Data, Make Decisions: The time allocated to the scientists in the tactical timeline to analyze data, make decisions and develop plan fragments for their activities is a significant factor in productivity.

We also asked participants for their thoughts on how productivity could be improved. Following are some of their responses:

- Improved resource modeling
- Better integration between Rover Planner tool and activity plan development tool.
- Simplified tools to make it easier to access science data.
- More intuitive methods for developing science plan fragments.
- Tools to assist developing Rover Planner activities.
- Support for documenting intent, to facilitate communication between supratactical and tactical operations.
- Allow the vehicle to insert activities onboard if extra resources are available.
- Reduce the duration of the tactical timeline to reduce the number of restricted sols.
- Develop a capability to enable a rover trajectory following mode.
- Create a tool to help scientists find and explore collected data.
- Increase orbiter relay bandwidth and reduce relay latencies.
- Staff the mission with a co-located, fully engaged team that is on the mission full time.

Pahrump Hills Walkabout Pass 1 Campaign (Sols 780 - 798)

Pahrump Hills was a significant destination for Curiosity because HiRISE imagery indicates that the outcrop represented the rover's first encounter with the Murray formation, the strata recognized as lower Mount Sharp. The team chose to explore the area employing a "walkabout" approach used by geologists on Earth, in which the geologist first walks the area to select locations for further, more detailed study. The objective of this initial pass of the Pahrump Hills Walkabout was to perform a reconnaissance imaging survey using the rover's remote sensing suite of instruments to identify and characterize the lithologies present in the formation.

The geography of Pahrump Hills was well suited for developing a strategic plan for the initial walkabout. The formation was laid out in front of the rover with an upward slope affording an excellent view of the exposed outcrop and potential paths for the rover to traverse. From this vantage, the scientists identified specific outcrops of interest. Working with Rover Planners, a strategic traverse plan was developed with anticipated mid-drive and end-of-drive imaging locations. These locations were refined as the campaign was carried out and higher-quality imagery was made available as the rover approached these areas. Scientists also used imagery acquired at each end-of-drive location to select specific targets for targeted remote sensing observations with the ChemCam and Mastcam instruments. The team also collected a "sidewalk" video documenting the traverse through the area using the MARDI instrument.

The Pahrump Hills walkabout was a highly successful science campaign. Data collected from the exploration enabled the team to develop a chemostratigraphic characterization of the formation. Video enabled tracking the distributions and relationships of sedimentary structures observed in the formation. The collected reconnaissance survey informed the follow-up observations with close contact instruments.

We performed a series of analyses to understand how mission resources were used during the campaign. This helps shed light on how effectively the team was able to make use of resources during the execution of the campaign. The analysis showed a general decrease in overall activity across multi-sol plans. This is likely due to limitations in how much activity can be developed during the tactical timeline. In addition, the team is limited in the types of activity that can be performed after a drive without ground-in-the-loop. The analysis also showed a significant decrease in the allocation of resources to campaign objectives following drives during multi-sol plans. This is because ground-in-the-loop is required to perform the majority of the activities needed to accomplish the campaign objectives such as picking science targets and planning drives.

The resource analysis indicated that the team was not constrained by energy during this campaign. **There was sufficient unused energy and sufficient non-productive vehicle awake time to support an estimated additional 72 hours of campaign-related activity over the span of the 19 sols.**

We observed a few significant factors that limited productivity during this campaign. As mentioned, restricted sols impacted the amount of campaign activity that could be included in several plans. There were two cases in which a planned drive had to be restricted in length due to a lack of terrain visibility from the rover's current location. There was a multi-sol plan in which significant campaign objectives were deferred to a later plan due to unanticipated interactions with separate, high-priority science objective.

Artist's Drive Campaign (Sols 949 - 972)

After completing an extensive exploration of the basal layer of Mount Sharp at Pahrump Hills, the team turned their attention to reaching higher layers of the mountain. In particular, they were interested in driving to a location, called Logan Pass, where orbital imagery showed a second formation, known as the Stimson Formation, made contact with the Murray formation of Pahrump Hills.

With the Logan Pass objective in mind, the scientists and engineers developed a strategic traverse plan to reach this location along a route named the Artist's Drive. The science team also identified strategic imaging locations from which the rover would have a good viewshed for imaging the surrounding topography. The route and specific imaging locations would be refined during the course of the campaign as higher-resolution data was obtained from the rover. The traverse through Artist's Drive valley originally aimed for Logan Pass, but during the campaign, the team discovered rock lenses sticking out of a hillside (later dubbed Logan's Run) and diverted to inspect them up close.

The team did not actually achieve the original goal of getting to Logan Pass due to excess slip encountered along the way. However, the Artist's Drive campaign successfully achieved the objectives of reaching a Stimson / Murray formation contact at the alternative location at Marias Pass. The campaign also provided the team with opportunities for imaging the topography along the route. In addition, a brief detour provided the opportunity to explore the intriguing outcrop at Logan's Run. The Logan's Run excursion is a good example of how the team identifies new science

opportunities in data collected by the rover, and how the team balances these new opportunities with their other objectives. The science team must weigh the potential science gains with the cost (e.g. sols, wheel wear) in accomplishing these new goals.

In terms of productivity, the campaign benefited from beginning near the start of a non-restricted period. This resulted in a relatively large number of unrestricted single-sol plans and, as a result, a high level of campaign activity during the early and middle sols. The mission re-entered restricted sols toward the end of the campaign and a decrease in productivity can be observed. The challenging terrain along Artist's Drive was another significant factor in the campaign's productivity, with several drives limited by available viewshed and a drive that ended early due to an obstacle in the rover's path. The need to re-do the APXS integration on a sample dump pile also impacted productivity as the team had to defer campaign objectives for a sol while the APXS observation was re-acquired. As with Pahrump Hills, our analysis indicated that there is opportunity for increasing productivity during the campaign by making use of shunt energy and otherwise unproductive RCE duration, with an estimated 62 hours of potential additional duration for campaign activities.

Marias Pass Campaign (Sols 991 - 997; 1027 - 1043)

The objective of the Marias Pass campaign was to study the contact between the Murray and Stimson formations at a location observed from HiRISE imagery. In contrast to Pahrump Hills, the team chose to explore the area using a linear approach, traditionally the more conventional strategy for rover exploration. The Marias Pass area presented more challenging terrain than Pahrump Hills including a hill to reach the contact, a sandy bowl near the contact and various terrain features that posed visibility occlusions. Toward the middle of the campaign, the team discovered unexpectedly high silica and hydrogen readings from prior locations. This motivated the team to backtrack to these previous areas for further study.

Data gathered during the campaign enabled the team to characterize the sedimentary facies and architectural elements preserved within the Stimson formation, and to determine the depositional history and regional relationships with the underlying Murray formation. These results are important in understanding the later history of the infilling of Gale crater and the potential habitability of the environment. The study of high-silica targets led to the conclusion that the introduction of silica represented one of the most recent water-rock interactions observed in Gale crater and has contributed to the study of areas the team visited later in the mission.

As with the other campaigns, we performed a series of analyses to understand how mission resources were used during the campaign. We observed similar patterns in that multi-sol plans resulted in an overall reduction in per-sol activity. And, as with other cases, the allocation of resources to campaign objectives was significantly reduced following drives in multi-sol plans. The resource analysis indicated that the team was not constrained by energy during this campaign. There was sufficient unused energy and sufficient non-productive vehicle awake time to support an estimated additional 70 hours of campaign-related activity over the span of the 24 sols.

In terms of significant productivity factors for this campaign, as with other campaigns, restricted sols resulted in a reduction in campaign-related activity. In addition, the challenging terrain in this area resulted in the need for additional drives to complete objectives. Two drives were constrained in distance by limited viewshed for route planning. Two drives ended early due to insufficient terrain features to support onboard position estimation. One drive was limited due to complexity in planning a route to circumvent a sandy area and another drive resulted in a placement of the rover with insufficient stability to perform the intended science objectives. This campaign also had two separate sols that were highly limited due to being part of extended, multi-sol plans to cover holiday weekends. The campaign included a few examples of received data prompting the team to repeat activity. There were cases where ChemCam observations did not hit the anticipated target, due to challenges entailed by the precise pointing requirements of the instrument, as well as a case in which imagery was re-acquired due to lighting conditions. The team's decision to re-trace their steps to investigate the high silica and hydrogen readings is another example, in this case requiring several sols of activity to acquire additional data.

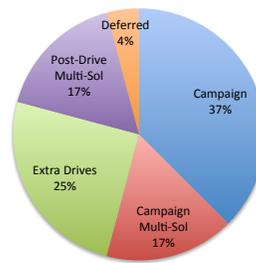
Discussion

The following table presents a rough breakdown of the sol-by-sol activity conducted in each campaign in terms of how activity on each sol contributed toward campaign objectives:

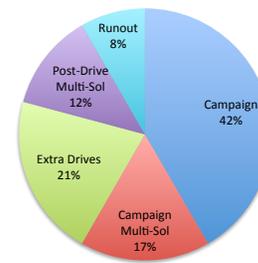
Sol Type	Pahrump Hills	Artist's Drive	Marias Pass
Campaign	6	9	10
Campaign Multi-Sol	2	4	4
Extra Drives	2	6	5
Post-Drive Multi-Sol	5	4	3
Deferred	3	1	0
Runout	1	0	2
Total Sols	19	24	24



Pahrump Hills



Artist's Drive



Marias Pass

Breakdown of sols for all campaigns.

Sols labeled “Campaign” were those that directly contributed to the campaign objectives with remote sensing and/or drives. “Campaign Multi-Sol” sols are those in which significant activity was performed toward the campaign objectives as part of a multi-sol plan, either due to a weekend or restricted planning. The reason for calling these sols out separately is that the presence of the multi-sol plan limited the team’s options for these sols. For example, had there not been a multi-sol plan, the team may have opted to move up activity that was performed in a subsequent plan (e.g. a drive activity) which would have reduced the overall number of sols required to achieve the campaign objectives. The “Extra Drives” label denotes sols in which unexpected drives were required. Sols labeled “Post-Drive Multi-Sol” were those sols in which the team was not able to achieve substantial campaign objectives due to lack of data following a drive during a multi-sol plan. “Runout” are sols of very low activity used in cases when the team had to create multi-sol plans but the tactical timeline capacity did not allow for sufficient time to develop activities for all sols of the plan. Finally, the sols labeled “Deferred” were sols in which campaign objectives were unexpectedly deferred due to the need to respond to an issue identified during tactical plan development or in response to an event from received downlink data. For the Pahrump Hills campaign, the deferred sols were due to an unexpected interaction, identified during tactical planning, between Pahrump Hills objectives and high-priority observations of the comet Siding Springs making its closest approach to Mars. For Artist’s Drive, the deferred sol was due to the need to repeat an activity from the previous sol, un-related to the Artist’s Drive objectives, as received data showed the activity did not have the intended result.

Comparing the campaigns we note that despite having different high-level objectives, the sol breakdown for Artist’s Drive and Marias Pass appear to be the most similar. This is due to these campaigns having a similar number of restricted plans and both being conducted in similar, challenging terrain conditions.

Comparing the sol breakdown for Pahrump Hills with the other two campaigns shows that restricted sols have a major productivity impact for these types of campaigns. The following table shows the number of nominal vs. restricted shifts for each campaign:

Sol Type	Pahrump Hills	Artist's Drive	Marias Pass
Nominal Shifts	2	16	12
Restricted Shifts	7	4	4
Total Shifts	9	20	16

Summary of shift types for all campaigns.

Pahrump Hills had a total of 9 tactical shifts of which 7 were during restricted periods of the mission. In contrast, Artist’s Drive and Marias Pass had more total shifts and fewer restricted shifts than Pahrump Hills.

The reason for the differences in number of restricted sols between Pahrump Hills and the other campaigns is largely luck of campaign timing. The Pahrump Hills Walkabout campaign happened to begin just as a restricted period was about to start. On the other hand, the Artist’s Drive campaign began just after a restricted period had ended. Marias Pass began toward the end of a nominal period but solar conjunction began before the restricted period began. By the time conjunction was over and the team returned to operations, the restricted period had completed. Thus, it was only toward the end of the Marias Pass campaign that another restricted period impacted operations.

The number of restricted sols is anticipated to increase for future missions as the current fleet of sun-synchronous orbiters are replaced with non-sun-synchronous orbiters. There are important science motivations for non-sun-synchronous orbiters, such as studying the Recurring Slope Lineae (RSL). However, such an orbit does not provide the consistent pattern of passes at the end of the rover’s day. This will result in many sols in which the operations team does not have sufficient time to develop new plans in response to the latest rover data, thus increasing the number of restricted sols.

The following table summarizes the traverses performed in each campaign:

Sol Type	Pahrump Hills	Artist’s Drive	Marias Pass
Nominal	5	6	3
Viewshed Limited	2	5	2
Drive Fault	0	1	2
Insufficient Stability	0	0	1
Total Traverses	7	12	8

Summary of traverses for all campaigns.

Note that the Artist’s Drive traverse marked as “Drive Fault” was also limited by viewshed. Rather than double-count it, we counted it as “Drive Fault” and not “Viewshed Limited”. The increased terrain occlusions encountered during the Artist’s Drive campaign lead to a larger number of traverses being limited by viewshed than encountered during the Pahrump Hills Campaign. Although the terrain at Marias Pass was more challenging than at Pahrump Hills, it had the same number of traverses limited by viewshed. This is likely because the Marias Pass campaign included returning to previously explored areas, allowing the team to make use of terrain imagery collected on previous sols. The more challenging terrain of Artist’s Drive and Marias Pass resulted in drive faults and rover stability issues in the associated campaigns.

It is interesting to compare the Pahrump Hills and Marias Pass campaigns as they had similar high-level objectives but were conducted with different exploration strategies. Unlike the Pahrump Hills Walkabout, the Marias Pass campaign did not have an extensive strategic plan to direct activity. This was largely due to the geography of the Marias Pass valley. HiRISE imagery provided a high level overview of the region, with sufficient detail to indicate that the area included a promising contact between the Stimson formation and Murray formation, but contained insufficient detail to form a strategic plan for exploring the location. Because the valley was elevated above the Artist’s Drive route the rover had been following, it was not possible to obtain the same type of Mastcam panorama that was available for planning the Pahrump Hills Walkabout.

Despite the absence of a detailed strategic plan for the Marias Pass campaign, the team was able to make quick tactical decisions and respond to new data as it arrived such as identifying drive routes and selecting key science targets. This can explain why there was a similar number of campaign-oriented sols as Artist’s Drive. **In other words, it seems that the number of restricted shifts and terrain challenges was a bigger factor than the availability of a guiding strategic plan, given the team’s ability to react.**

It is also interesting to compare the walkabout approach employed at Pahrump Hills vs. the linear approach used at Marias Pass. Although the team intended to use a linear strategy at Marias Pass, they ended up backtracking to explore data collected near the Sol 992 location. There was additional backtracking in the sols that followed the end of our case study sol range. It was suggested by one of the scientists in our interviews that perhaps the Marias Pass campaign would have been overall more efficient had it employed a walkabout approach. One of the objectives we have with this case study is to leverage what we have learned from these productivity challenges to identify flight and ground approaches that can reduce the overhead of employing a walkabout approach.

We performed a series of analyses on how the team allocated vehicle resources throughout each campaign. This included tracking of predicted and actual allocations of flight computer duration, energy and data volume. The analyses of resource allocations followed a similar pattern for each of the campaigns. The impact of multi-sol planning due to weekends and restricted sols had the largest impact on how effectively the team was able to allocate resources toward

campaign activity. The analysis showed a general decrease in overall activity across multi-sol plans. For 2-sol plans, there was a 12% reduction of RCE duration and 7% reduction in energy use. For 3-sol plans, there was a 15% reduction of RCE duration and 11% reduction in energy use. This is likely due to limitations in how much activity can be developed during the tactical timeline. In addition, the team is limited in the types of activity that can be performed after a drive without ground-in-the-loop. The analysis also showed a significant decrease in the allocation of resources to campaign objectives following drives during multi-sol plans. This is because ground-in-the-loop is required to perform the majority of the activities needed to accomplish the campaign objectives.

The resource analysis indicated there were available resources that could provide an opportunity for increasing productivity. There was sufficient unused energy and sufficient non-productive vehicle awake time to support an estimated additional 72 hours, 62 hours and 70 hours of campaign activity could have been performed during the Pahrump Hills, Artist's Drive and Marias Pass campaigns, respectively.

In summary, following is a list of the significant productivity factors identified in the case study:

Ground-in-the-loop requirements for target selection and effective drive planning: This results in a significant drop in productivity on sols that follow drives during restricted periods of the mission. Even during non-restricted sols, it constrains the timing of activity that can change the state of the vehicle and activity that acquires decisional data to occur prior to the decisional pass.

Capacity of tactical timeline to fill multi-sol plans: Due to the time required to develop and validate command products, the amount of overall activity across a multi-sol plan is generally lower than the amount of activity across a similar number of single sol plans.

Ground-in-the-loop requirements to respond to outcome of activity: We observed several instances where the team decided to re-do an activity, or return to a previous location, after observing the data received from the vehicle. This included re-doing an APXS integration during Artist's Drive, re-acquiring Mastcam and ChemCam observations during Marias Pass, and returning to a previous location in response to interesting instrument readings during Marias Pass. There were also several cases where drives ended early due to unexpected terrain conditions and had to be re-planned.

Use of margin and cleanup duration allocations: Because there is uncertainty in the actual run time of activities, ground operations allocates margin to allow activities to run long without impacting future activities. And to protect against activities exceeding margin, additional time is reserved for cleaning up activities. This results in non-productive time when activities run within predicted durations.

Ability to Exploit Supratactical Work: If the tactical team follows the supratactical plan, they are able to leverage work performed on the supratactical timeline and productivity benefits. If decisional data prompts the tactical team to change course, they may need to start from scratch and are often not able to accomplish as much.

Overhead in Command Product Generation: The effort involved in developing and validating command products can limit the amount of activity that can be accomplished.

Managing Tactical Timeline Complexity: Both the supratactical and tactical teams must judge how much activity can be managed within the tactical timeline. The ability to accurately make this judgment can impact productivity.

Predicting Available Vehicle Resources: Inaccuracies in resource modeling, including activity power and duration requirements, can result in unnecessarily restricting planned activity.

Interpersonal Communication: Effective communication was identified as a significant productivity factor by many of the participants in each of the operations roles we interviewed. This includes communication between the science and engineering teams and among roles within science and engineering teams.

Science Team Engagement: It is important for science team members to be aware of the current and past context of the mission to make informed science decisions. A significant challenge to engagement is that many team members work on the mission part time.

Time to Analyze Data, Make Decisions: The time allocated to the scientists in the tactical timeline to analyze data, make decisions and develop plan fragments for their activities is a significant factor in productivity.

Following is a summary of common types of ground-in-the-loop decisions across the campaigns:

Selecting targets for ChemCam, Mastcam and contact science: While distant imagery of terrain provided sufficient information to indicate the value of traversing to an area, the scientists required the higher quality imaging of the area, obtained when the rover arrives at the site, to select specific targets.

Drive planning: Post-drive imagery is also used to provide the data necessary to plan the next traverse, including allowing the scientists to refine their selection of the end-of-drive location and the engineering team to design a route for the rover to follow.

Stability assessment for contact science: Prior to deploying the arm and performing contact science, the team must use data from the rover's current position to assess the vehicle's stability.

Responding to problems in activity executions: Each campaign had several examples of the need to respond to partially successful or failed activities including drives that ended earlier than expected due to unexpected terrain and remote sensing observations that had less than satisfactory lighting or that missed their target.

Science Discovery: Both the Artist's Drive and the Marias Pass campaigns contained examples of the importance of analyzing collected data for unexpected science discoveries.

Conclusions

A successfully deployed Martian rover represents an immensely valuable asset for science exploration and discovery. As such, surface missions have a strong interest in getting the most out of the vehicle to increase the return on investment. This is further motivated by the fact that the rover's capabilities will inevitably degrade over time. As such there is a strong interest in enhancing the productivity of future surface rover missions.

We conducted the case study of MSL campaigns with the objective to better understand the productivity challenges facing surface missions. The study included interviews with mission scientists and engineers along with detailed study of three science campaigns. The responses from the interview participants and our analysis of the campaigns showed that there are opportunities for increasing surface mission productivity. In particular, we observed that it is often the case that the vehicle has more available resources than the operations team is able to use. The case study identified a variety of issues that are limiting the productive use of these resources.

Perhaps the largest factor observed in our study of MSL campaigns was due to restricted sols and, more generally, the reliance on ground-in-the-loop to inform a large portion of the rover's activities. The ground-in-the-loop reliance is expected to become an even more significant liability to surface missions as the fleet of aging sun-synchronous orbiters are replaced with non-sun-synchronous orbiters that do not provide a consistent "end of day" downlink of the rover's state. The next biggest productivity factor across all campaigns was the need for additional drives to reach objectives due to terrain interactions. There were several cases in which drives were limited due to available viewshed for route planning. In other cases, drives faulted out early due to issues such as visual odometry failure.

The interview participants identified additional significant productivity factors. There was large consensus that effective communication was key to productive operations. Several participants agreed that the complexity and overhead of producing command products was a primary factor in explaining why the team often does not use all available vehicle resources. The scientist participants also emphasized the importance and challenge of engagement and situation awareness along with the need for time to analyze downlink results and consider options for activities.

As discussed earlier, we focused our study on campaigns that emphasized driving and remote sensing. Though two of the campaigns we chose included contact science activity, there would be value in studying campaigns that made an emphasis on contact science as well as sampling campaigns. We believe many of the productivity factors identified in the current study are relevant for contact science and sampling campaigns, but there are likely to be additional important productivity factors identified in these other campaign types.

Our next objective is to identify changes to flight systems and ground operations practices to overcome these challenges and enable high levels of productivity for future surface missions. The findings from this study will guide the design and development process by helping to define the capabilities required to meet these productivity challenges. We will also leverage examples from the campaigns studied to define scenarios that will be used to focus the development and to evaluate the performance of our work.

Contents

1	Introduction	16
1.1	Objective of the Case Study	16
1.2	Background on the Curiosity Rover	16
1.3	Background on MSL Operations	19
1.4	An Example Sol in the Life of the Rover	21
2	Case Study Design	24
2.1	Selection of Cases	24
2.2	Conceptual Model of Achieving Objectives	26
2.2.1	Step (A): Activity Development	27
2.2.2	Step (B): Plan Complexity Analysis	29
2.2.3	Step (C): Activity Refinement	29
2.2.4	Step (D): Uplink	29
2.2.5	Steps (E) and (F): Execution	30
2.2.6	Steps (G): Downlink	30
2.3	Data Collection Methodology	30
2.4	Caveats	33
3	Interviews with Operations Personnel	34
3.1	Interview Design	34
3.2	Defining Mission Productivity	36
3.3	Identifying Productivity Factors	36
3.4	Thoughts on Improving Productivity	40
4	Pahrump Hills Walkabout Pass 1 Campaign	41
4.1	Overview of Campaign	41
4.1.1	Science Objectives	45
4.1.2	Summary of Activity	45
4.1.3	Outcome of Campaign	54
4.2	Analysis of Resource Usage	54
4.3	Significant Productivity Factors	64
4.4	Significant Ground-in-the-Loop Cycles	68
4.5	Conclusions	69
5	Artist's Drive Campaign	70
5.1	Overview of Campaign	70
5.1.1	Science Objectives	72
5.1.2	Summary of Activity	72
5.1.3	Outcome of Campaign	82
5.2	Analysis of Resource Usage	84
5.3	Significant Productivity Factors	90
5.4	Significant Ground-in-the-Loop Cycles	91
5.5	Conclusions	93
6	Marias Pass Campaign	93
6.1	Overview of Campaign	93
6.1.1	Science Objectives	95
6.1.2	Summary of Activity	96
6.1.3	Outcome of Campaign	107
6.2	Analysis of Resource Usage	108
6.3	Significant Productivity Factors	114
6.4	Significant Ground-in-the-Loop Cycles	118

6.5 Conclusions	119
7 Discussion	119
8 Conclusions	123

List of Figures

1	Case study questions.	17
2	Curiosity rover.	17
3	Overview of MSL operations.	19
4	Example sol in the life of the rover.	22
5	Selected campaigns for case study (NASA/JPL-Caltech/MSSS/Univ. of Arizona).	25
6	Seasonal temperatures at Gale Crater (NASA/JPL-Caltech/CAB(CSIC-INTA)).	26
7	Factors impacting surface mission productivity.	27
8	Mars activity vs. Earth planning.	28
9	Data collection schema.	31
10	Interview questions.	35
11	Curiosity's route from landing to Pahrump Hills (NASA/JPL-Caltech/Univ. of Arizona).	42
12	Contact (dashed line) between Gale Crater plains (Aeolis Palus) and Mount Sharp (Aeolis Mons) showing Pahrump Hills at the start of Murray formation. (NASA/JPL-Caltech/Univ. of Arizona).	43
13	Geologic Layers of Mount Sharp (NASA/JPL-Caltech/Univ. of Arizona).	44
14	Curiosity's at Pahrump Hills (NASA/JPL-Caltech/Univ. of Arizona).	45
15	Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of campaign.	46
16	Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of Sol 785 planning.	49
17	Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of Sol 789 planning.	51
18	Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of Sol 792 planning.	52
19	Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of Sol 794 planning.	52
20	Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of Sol 796 planning.	53
21	Traverse map of Curiosity during Pahrump Hills Walkabout Pass 1 (NASA/JPL-Caltech/Univ. of Arizona).	55
22	Allocation of RCE duration for Pahrump Hills Walkabout Pass 1 campaign.	57
23	Allocation of data volume for Pahrump Hills Walkabout Pass 1 campaign.	58
24	Downlink performance for Pahrump Hills Walkabout Pass 1 campaign.	60
25	Summary of device warmup activity for Pahrump Hills Walkabout Pass 1 campaign.	60
26	Allocation of energy for Pahrump Hills Walkabout Pass 1 campaign.	61
27	Summary of battery state of charge for Pahrump Hills Walkabout Pass 1 campaign.	62
28	Summary of battery and MMRTG energy for Pahrump Hills Walkabout Pass 1 campaign.	63
29	Estimate of extra duration availability for Pahrump Hills Walkabout Pass 1 campaign.	63
30	Allocation of RCE duration during daylight hours for Pahrump Hills Walkabout Pass 1 campaign.	64
31	Tactical timeline durations for Pahrump Hills Walkabout Pass 1 campaign.	65
32	Visualization of the Pahrump Hills Walkabout 1 Campaign.	67
33	View of Logan Pass from HiRISE (NASA/JPL-Caltech/Univ. of Arizona).	71
34	Strategic traverse route for Artist's drive including anticipated imaging locations.	71
35	Curiosity's planned drive direction along Artist's Drive at the start of planning for Sol 950.	73
36	Ripples as seen by Curiosity on Sol 951 for planning Sol 952.	74
37	Southwest traverse route chosen at decision point on Sol 956.	77
38	Drive direction ahead for Curiosity for planning Sol 957.	78
39	Long drive planned on Sol 960, deviating north of the Strategic Traverse Route.	79
40	The Mt. Shields area, which appears to a valley carved in bedrock and later filled with sedimentary material. This area is the target of the diversion away from Logan's Pass called "Logan's Run." (NASA/JPL-Caltech/MSSS).	80
41	This NavCam image taken after the Sol 963 drive shows the nearby large rock that autonomous navigation deemed a hazard and stopped the drive early (NASA/JPL-Caltech).	81
42	A sketch of the sol 967 drive, approaching an outcrop at the lower part of Mt. Shields (NASA/JPL-Caltech).	82
43	Traverse path during the Artist's Drive campaign (NASA/JPL-Caltech/Univ. of Arizona).	83
44	Allocation of RCE duration for Artist's Drive campaign.	84
45	Allocation of data volume for Artist's Drive campaign.	85

46	Downlink performance for Artist's Drive campaign.	86
47	Summary of device warmup activity for Artist's Drive campaign.	87
48	Allocation of energy for Artist's Drive campaign.	87
49	Summary of battery state of charge for Artist's Drive campaign.	88
50	Summary of battery and MMRTG energy for Artist's Drive campaign.	88
51	Estimate of extra duration availability for Artist's Drive campaign.	89
52	Allocation of RCE duration during daylight hours for Artist's Drive campaign.	89
53	Tactical timeline durations for Artist's Drive campaign.	90
54	Visualization of the Artist's Drive Campaign.	92
55	Logan Pass and Marias Pass (NASA/JPL-Caltech/MSSS).	94
56	Traverse map through Sol 991 showing contact between Stimson and Murray formation at Marias Pass (NASA/JPL-Caltech/Univ. of Arizona).	94
57	Marias Pass, showing the contact between Stimson Formation and Murray formation (NASA/JPL-Caltech/ MSSS).	95
58	The hill to climb on Sol 991.	96
59	View of the Marias Pass contact following the Sol 991 drive showing the locations selected for closer study.	98
60	Navcam image taken after the Sol 991 drive, showing the Elk and Lamoose targets.	99
61	View of the Marias Pass contact following the Sol 992 drive.	100
62	Navcams following the Sol 995 drive, showing the Missoula and Ronan targets.	101
63	MAHLI mosaic (40cm across) of the Missoula target at the contact between Stimson and Murray formations at Marias Pass (NASA/JPL-Caltech/MSSS).	104
64	Hazcam images following the drives back to the Sol 991 location.	106
65	Allocation of RCE duration for Marias Pass campaign.	109
66	Allocation of data volume for Marias Pass campaign.	110
67	Downlink performance for Marias Pass campaign.	111
68	Summary of device warmup activity for Marias Pass campaign.	111
69	Allocation of energy for Marias Pass campaign.	112
70	Summary of battery state of charge for Marias Pass campaign.	112
71	Summary of battery and MMRTG energy for Marias Pass campaign.	113
72	Estimate of extra duration availability for Marias Pass campaign.	113
73	Allocation of RCE duration during daylight hours for Marias Pass campaign.	114
74	Tactical timeline durations for Marias Pass campaign.	115
75	Visualization of the Marias Pass Campaign.	116
76	Per-sol allocation of resource allocations for 1-sol, 2-sol and 3-sol plans.	122

List of Tables

1	Breakdown of sols for the Pahrump Hills Walkabout Pass 1 campaign.	68
2	Predicted vs. actual critical data volumes for decisional limited sos during Artist's Drive campaign.	86
3	Breakdown of sols for the Artist's Drive campaign.	91
4	Breakdown of sols for the Marias Pass campaign.	115
5	Breakdown of sols for all campaigns.	120
6	Summary of shift types for all campaigns.	120
7	Summary of traverses for all campaigns.	121

1 Introduction

The Curiosity rover has been exploring Gale Crater and Mount Sharp since its landing in August 2012. During this time, the Mars Science Laboratory (MSL) mission has accomplished many significant objectives. It has achieved the success criteria for the prime mission, collected evidence that indicates Mars was once habitable, collected over a dozen samples and driven more than 12 kilometers [Grotzinger et al.(2015), Vasavada et al.(2014)]. Curiosity is currently in its extended mission and continues to make new discoveries as it explores Mount Sharp.

While the Mars rovers, including Spirit, Opportunity and Curiosity, have demonstrated an incredible ability to survive far beyond their designed lifetimes, they still represent limited opportunities to explore the planet. As such, there is great interest in getting the most out of these landed assets over the course of the missions.

Maintaining high levels of productivity for the Curiosity rover is challenging. While the operations team has made significant accomplishments with the rover, doing so often requires a large amount of human effort in planning, coordinating, sequencing and validating the development of command products for the rover. Further, limitations in communication opportunities and anomalies on the vehicle can sometimes cause delays in accomplishing the team's objectives. These productivity challenges can result in the under-utilization of the vehicle's resources.

In order to better understand the productivity challenges faced by surface missions, we conducted a case study of selected MSL science campaigns. For each campaign, we examined the activity performed by the operations team, how the team decided what to accomplish each day, how well these objectives were achieved and how results from one day fed into and informed objectives for the next and how the team allocated vehicle resources during the campaigns. We also conducted a series of interviews with operations personnel to get their broader perspective on surface mission productivity factors.

This document provides a report on this case study. We begin with a summary of the specific objectives of the case study. Next we present an overview of the MSL mission, including a description of the Curiosity rover's capabilities and a high-level summary of MSL operations. Section 2 presents the design of the case study, describing the selection of cases and the methodology for collecting and analyzing data. Section 3 presents the results from the interviews we conducted with operations personnel. We provide the questions used to guide the interviews and a summary of the resulting discussions.

The following sections present the results of each of the three science campaigns we selected for the study. For each campaign we provide background on the science objectives for the campaign, a description of the team's activity during the campaign, and an analysis of how resources were allocated throughout the campaign. We use these results to provide an assessment of the major productivity factors that influenced the campaign and a summary of the significant ground-in-the-loop decisions that were made. In Section 7 we compare and contrast the three different science campaigns. We conclude with a summary of what the case study has revealed regarding surface mission productivity.

1.1 Objective of the Case Study

The specific objective of the case study was to develop an understanding of the factors that contribute toward and detract from surface mission productivity. Figure 1 summarizes the questions we wished to explore through the case study.

The case study is in support of a multi-year research project with the objective of increasing productivity for future surface missions. By answering these questions we hope gain a deeper understanding of surface mission productivity that will enable us to develop suggestions for changes to rover flight systems and ground operation practices to enhance productivity for future missions. In Section 2 we will describe how the case study was designed to address these questions.

1.2 Background on the Curiosity Rover

Figure 2 shows major systems of the Curiosity rover. The rover is a roughly car-sized, six-wheeled rover 2.2 meters (7 feet) tall, 2.6 meters (9 feet) wide, and 3 meters (10 feet) long with a mass of 899 kilograms (1,982 pounds).

Power for the rover is powered by a multi-mission radioisotope thermoelectric generator (MMRTG). The MMRTG converts heat from the decay of plutonium-238 dioxide into electricity. At the start of the surface mission, the MMRTG's output was about 110 watts. This output decreases over time as the heat source degrades. The energy is used to charge two lithium-ion rechargeable batteries, each with a capacity of 42 amp-hours. The batteries enable the

Primary Case Study Questions:

1. How does the operations team define productivity?
2. What factors tended to increase productivity?
3. What factors tended to decrease productivity?

Secondary Case Study Questions:

1. How did the team select activities to accomplish intent?
2. What types of ground-in-the-loop decisions were made?

Figure 1: Case study questions.

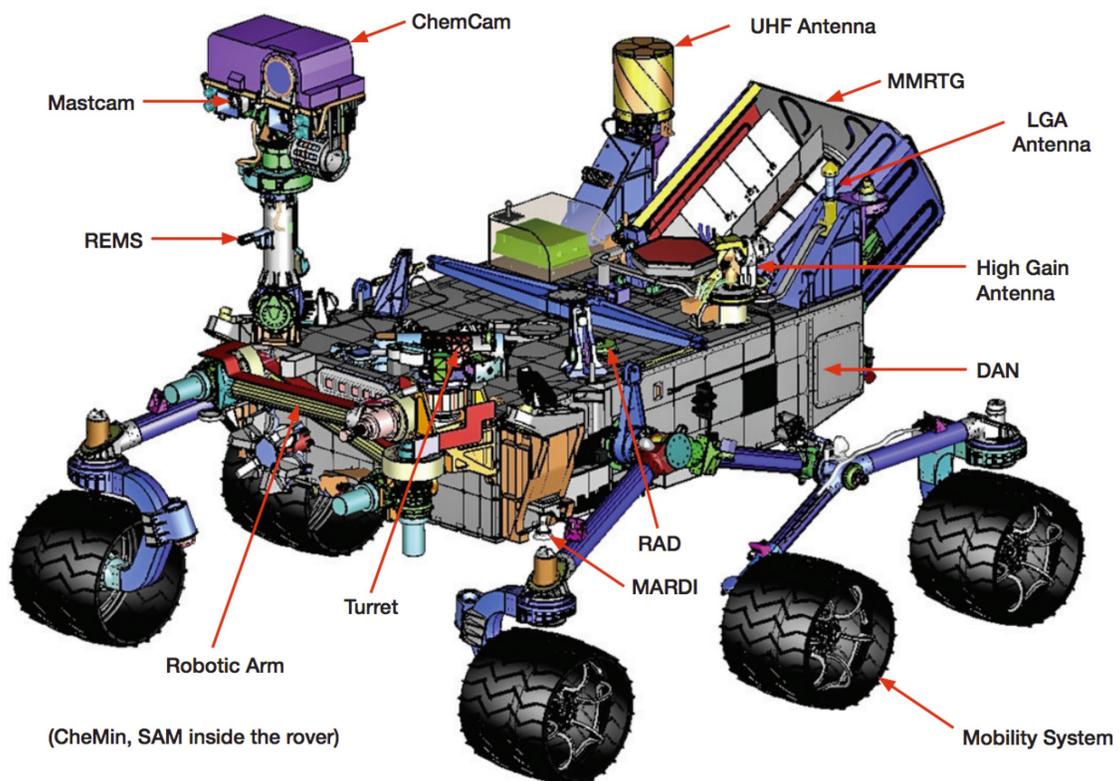


Figure 2: Curiosity rover.

rover to store energy to support activities that temporarily exceed the energy output of the MMRTG. Excess energy that is not put into the batteries is shunted.

The rover has three different antennas used for communication. A High Gain Antenna (HGA) and a Low Gain Antenna (LGA) are used for X-Band (7-8 gigahertz) communication directly to and from Earth. These are relatively low bandwidth communication channels for the rover. As such, the mission uses them primarily for uplinking command products to the rover, which tend to be small in size. The vast majority of the data transmitted by the rover is relayed through Mars orbiting spacecraft through the UHF antenna (about 400 megahertz).

The Rover Compute Element (RCE) provides the rover's computational power and storage for data. The RCE contains the CPU, a radiation-hardened RAD750 CPU operating at 133 MHz with 128 MB of onboard RAM, 256 MB of off-board DRAM and 4 GB of flash memory for data storage. In fact, there are two copies of the RCE, but only one is actively operating the vehicle. The other is a backup in case problems occur with the primary computer.

The rover has a 1.9 meter (6.2 feet) long, five-degree-of-freedom arm with instruments and devices to support contact science. In addition to the MAHLI and APXS instruments, described below, it holds a Dust Removal Tool (DRT) for brushing away the dust layer on rock surfaces and a percussive drill to acquire rock samples.

The mobility system consists of a rocker-bogie suspension system with six wheels. The four corner wheels are independently steerable, enabling the rover to turn in place. Late in the prime mission, the operations team began noticing an increase in the amount of wear the wheels were experiencing due to traveling across rough terrain [Jet Propulsion Laboratory Press Release(2013)]. As a result, the team became more cautious in the drive paths that were selected and began periodically monitoring the wheel health.

To support navigation, the rover has a set of stereo hazard avoidance cameras (aka Hazcams) mounted lower on the front and rear of the vehicle. These are wide-angle (about 120° field of view) cameras providing a view of terrain near the rover. A pair of stereo navigation cameras (aka Navcams) are mounted on the rover's mast which provide about 45° field of view used to view the terrain further from the rover.

The rover has a wide range of science instruments to assist in its exploration:

Mastcam: The Mastcam consists of a pair of two-megapixel color cameras. The left camera has a 34-millimeter lens providing a medium field of view of about 18°. The right camera has a 100-millimeter lens with a narrower field of view of about 6°. The cameras are also capable of capturing video at 720p at four to seven frames per second.

ChemCam: The ChemCam instrument is able to fire a laser allowing it to analyze the elemental composition of vaporized material. The laser is also able to clear away dust from rocks. ChemCam also includes a telescopic camera called the Remote Microscopic Imager (RMI) with a very small field of view of about 21 milliradians (or about 1.2°). The RMI is used to provide context for the laser firings.

MARDI: The Mars Descent Imager (MARDI) instrument was primarily designed to provide color video of the rover's descent onto the surface. Since landing, the downward facing camera has been used to collect images and video of the terrain under the vehicle. The camera has a field of view of 70° by 55°.

DAN: The Dynamic Albedo of Neutrons (DAN) instrument is used to detect hydrogen at or near the surface by monitoring the speed of neutrons. The instrument can be operated in two modes. An active mode in which the instrument actively emits neutrons or a passive mode in which it monitors neutrons dislodged by cosmic rays.

RAD: The Radiation Assessment Detector (RAD) instrument is used to characterize radiation on the Martian surface in preparation for future human exploration.

REMS: Rover Environmental Monitoring Station (REMS) is a weather station capable of measuring humidity, pressure, temperatures, wind speeds, and ultraviolet radiation.

APXS: The Alpha Particle X-Ray Spectrometer (APXS) is an arm-mounted instrument that irradiates targets with alpha particles and maps the spectra of the X-rays that are emitted to determine the material's chemical composition.

MAHLI: Another arm-mounted instrument, the Mars Hand Lens Imager (MAHLI) is a color imager capable of imaging features as small as 12.5 micrometers.

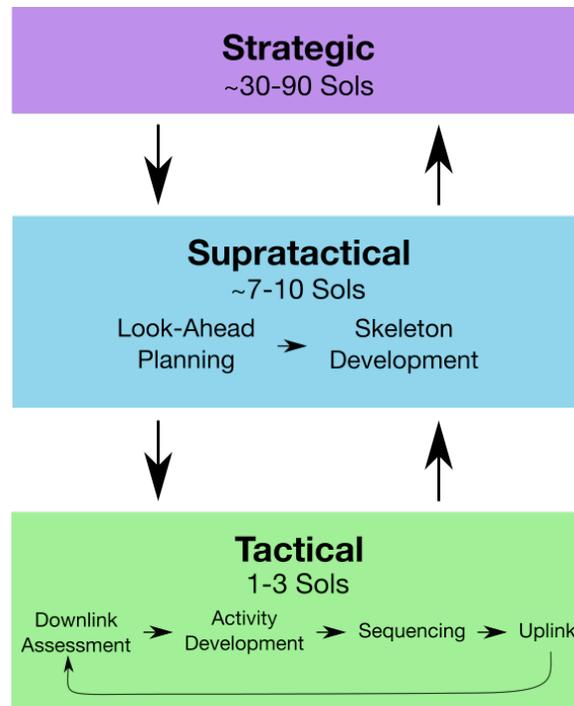


Figure 3: Overview of MSL operations.

CheMin: The Chemistry and Mineralogy (CheMin) instrument is one of two sample processing instruments housed inside the rover. It uses X-ray diffraction and X-ray fluorescence to identify and quantify minerals present in samples.

SAM: Sample Analysis on Mars (SAM) is the second sample processing instrument inside the rover able to analyze organics and gasses in atmospheric and solid samples.

1.3 Background on MSL Operations

One of the challenges a surface mission has compared to an orbital mission is that a surface mission is impacted more significantly by a priori unknown and changing environmental conditions. While orbital imagery provides valuable information to guide activity, it does not capture all the conditions that affect the rover. For example, while orbital imagery may indicate that exploring a particular region is promising to achieve a science objective, the specific science targets are not known until additional data is collected from the rover itself, such as images from its mast-mounted cameras. Further, orbital data is insufficient for fully predicting specific terrain conditions that will impact the rover's traversability and its ability to perform close-contact operations on targets of interest.

As such, surface operations must be reactive and respond to the results of activity carried out during the previous sol (Martian day). This daily planning activity is referred to as **tactical** operations and is patterned after the tactical operations developed for the Mars Exploration Rovers [Mishkin et al.(2006)]. In order to assist in coordinating the complex operations of the rover as it conducts science campaigns, MSL operations also includes supratactical and strategic phases.

Figure 3 shows the relationship among these three major phases of operations — strategic, supratactical and tactical [Chattopadhyay et al.(2014)]. These processes are structured to enable the team to achieve long-term science objectives, managing the rover's limited resources, while still responding to the dynamic nature of surface exploration.

Strategic planning focuses on developing long-term plans, typically spanning weeks or months, to achieve high-level objectives. For example, in the Pahrump Hills campaign, discussed in Section 4, the strategic plan spanned several months and specified a multi-pass approach to exploring the region in order to achieve the high-level objectives of performing a comprehensive study of the formation. Strategic planning for the Artist's Drive campaign, presented in

Section 5, included the development of a Strategic Traverse Route (STR) to provide guidance for selecting paths for the rover on its longer-term objective to reach higher levels of Mount Sharp. These strategic plans provide vital guidance in achieving mission objectives, but they are not directly executable. They must be adapted to take into account the current conditions and adjusted to respond to unanticipated conditions as the rover explores the environment. For example, while the STR provides guidance on the direction the rover should travel, the actual tactical routes may deviate from the route in order to respond to local terrain conditions. And in some cases, significant alterations to the STR were required when a particular path was discovered to be non-traversable or in cases where unexpected science objectives were identified.

The supratactical stage provides a bridge between the long-term strategic plan and the day-to-day, highly reactive tactical process. The process is designed to coordinate the complex science instruments and manage the constraints and resources required to conduct campaigns. The supratactical process produces **look-ahead plans** which span several sols, typically a week, of activity. These plans help maximize the use of vehicle resources. For example, if an energy-intensive, multi-sol sampling experiment is coming up, the look-ahead plan provides guidelines on how much energy can be used each sol of operations. The process also helps with coordination among the large science team spread across the globe.

The supratactical process feeds into the tactical process by delivering a **skeleton** plan for each sol of tactical planning. The skeleton provides the tactical team with the major objectives for the plan, e.g. drive toward a particular location, or perform close-contact operations. It includes a rough structure of the activities, including coordination of science activities around communications windows and other engineering activities and guidelines on how much resources, such as energy, time and available data volume, can be expended during the execution of the plan.

The science team begins working with the skeleton plan during the planning cycle that precedes the tactical planning cycle for that skeleton. For example, on Monday, the science team will begin working with the skeleton for Tuesday's planning. This gives the team extra time to think through options for the next plan. Depending on the type of campaign the team is in, this work may be relatively abstract. For example, if the rover is in a driving campaign, the team may not know specific targets they will want to observe as they will not yet have data down from the location the rover will be in. However, given the estimated available resources from the skeleton, they can decide how to allocate resources for remote sensing based on the types of observations they are likely to perform (e.g. number and type of ChemCam and Mastcam observations). If the rover will be in the same area for a while, e.g. while conducting an extended contact science campaign, the team will be able to develop detailed plans ahead of the tactical timeline.

This use of the supratactical process "virtually" extends the tactical timeline, providing additional time to develop activities beyond the limits of the tactical timeline. This does, of course, carry the risk that an unexpected event may arise, causing the need to change the direction of the tactical plan. In this case, some or all of the supratactical work may no longer apply and this work must be re-done on the tactical timeline.

The tactical planning process forms the highly reactive phase of surface operations. It includes an assessment of the state of the vehicle and the performance of the previous plan's activities. During activity development stage, specific science and engineering objectives are identified based on the high-level objectives of the current campaign and guidelines provided in the skeleton plan. The developed activity plan is translated into sequencing command products to be executed on the vehicle. These command products are verified, reviewed and delivered for uplink to the rover.

The tactical timeline is organized around several meetings which are used for team coordination and to review the evolving command products. The meetings are as follows:

TACT: The Tactical Activity Coordination Tagup (TACT) meeting is held early in the tactical timeline. The engineering team provides a brief, preliminary assessment of the health of the vehicle based on the latest downlink and summarizes any significant constraints or issues that may impact that day's planning. The science team provides a preview of the significant science activities that are anticipated for the plan. A primary objective of this meeting is to ensure good coordination between the science requests for the Rover Planners (e.g. contact science targets and activities of interest, or drive destinations).

SOWG: After TACT, the science and engineering teams work separately on their portions of the day's plan. These pieces of the plan, aka plan fragments, are delivered to the Science Planner, who combines them with the skeleton plan to create a combined activity plan. The team then convenes for the Science Operations Working Group (SOWG) meeting. The team reviews the plan and verifies the planned activities fit within predicted

resource constraints such as time, energy and data volume. If constraints are violated, the team works to resolve those constraints, which may include limiting activities or descopeing them from the plan.

APAM: There are typically refinements, such as adjusting activity parameters, that need to be made to the plan following SOWG. The teams perform these adjustments on their portions of the plan and the Science Planner integrates their updates back into the activity plan. The team meets again for the Activity Plan Approval Meeting (APAM) in which any significant changes made to the plan since SOWG are reviewed and a final check of resources is performed. If everything looks good, the team officially approves the plan and moves on to developing sequences to implement the plan activities.

Master / Submaster Walkthrough: After APAM, the science and engineering teams develop sequences to implement the activities in the plan. At the Master / Submaster Walkthrough, the team reviews a subset of these sequences. The Master sequence is the sequence that controls the overall execution of the plan. It commands the wakeups and shutdowns of the vehicle, initiates blocks of activity and ensures that activity completes when expected and cleans up after their completion. Submaster sequences are invoked by the Master sequence and are responsible for the execution of a block of activity, e.g. a pre-drive imaging block of remote sensing activity. The team reviews the Master and Submaster sequences along with the Rover Planner sequences to ensure appropriate coordination among the science and engineering activities in the plan.

Sequence Report Walkthrough: Following the Master / Submaster Walkthrough, the science and engineering teams perform independent validation of their sequences and deliver them to the Sequence Integration Engineer who integrates the sequences and performs combined validation checks. Validation checks include simulation of the execution of the sequences along with manual and automated flight rules checks. At the Sequence Report Walkthrough, the team reviews the integrated command products and dispositions any flight rule issues.

CAM: The final meeting of the day is the Command Approval Meeting (CAM). At this point, all command products have been prepared for radiation to the spacecraft. The team reviews final checks on the prepared products and reviews the product that will be used to coordinate the radiation of the command products to the rover. At this meeting, the team makes the final approval for sending the products to the vehicle.

The project has reduced the duration of the tactical timeline over the course of the mission. At the start of the surface mission, the tactical timeline was 16 hours, consisting of two separate shifts of approximately 8 hours each. The acquisition of team experience and the development of ground tools enabled tactical operations to be reduced in duration and combined into a single shift. At the time of the Pahrump Hills campaign, the first campaign of our case study, the nominal tactical timeline duration was 10 hours. By the time of the Artist's Drive and Marias Pass campaigns, the second and third campaigns of our study, the nominal tactical timeline duration was 9.5 hours.

1.4 An Example Sol in the Life of the Rover

Figure 4 illustrates a typical drive sol for the rover, based on a simplified version of the Sol 780 activity plan. The “Uplink” window near the beginning of activity in Figure 4 represents the X-band communication window used by the team to send the Sol 780 command products to the rover. Nearly every sol has a similar uplink window in the rover's morning to enable the team to command the vehicle. In some cases, such as in preparation for weekends, the team will create multi-sol plans which cover multiple, consecutive sols. In these cases, the team will uplink the command products for all of the sols in the multi-sol plan to the rover on the first sol of the plan. The rover will still perform the morning uplink windows on the subsequent sols of the plan, in case the team needs to respond to some unanticipated situation, but in most cases, no new command products will be received on these sols.

The operations team creates a **Master** sequence to coordinate the activity in a sol's activity plan and includes the sequence with the other command products in the uplink to the rover. The Master for a given sol begins just after the uplink window for the sol through the uplink window for the next sol. In this case, the Sol 780 Master starts just after the “Uplink” window in Figure 4, which ends at 10:00 on Sol 780, and then continues into the morning of Sol 781 until the end of the Sol 781 “Uplink” window, which ends at 10:00 on Sol 781. The mission refers to the transition from one Master to the next as **handover**. As such, when the mission refers to the plan for a given sol or activity occurring within a give sol, they are usually referring to the period of time from handover on that sol through handover on the next sol. We adopt this convention in this report. At first, this can be confusing as, in this terminology, the

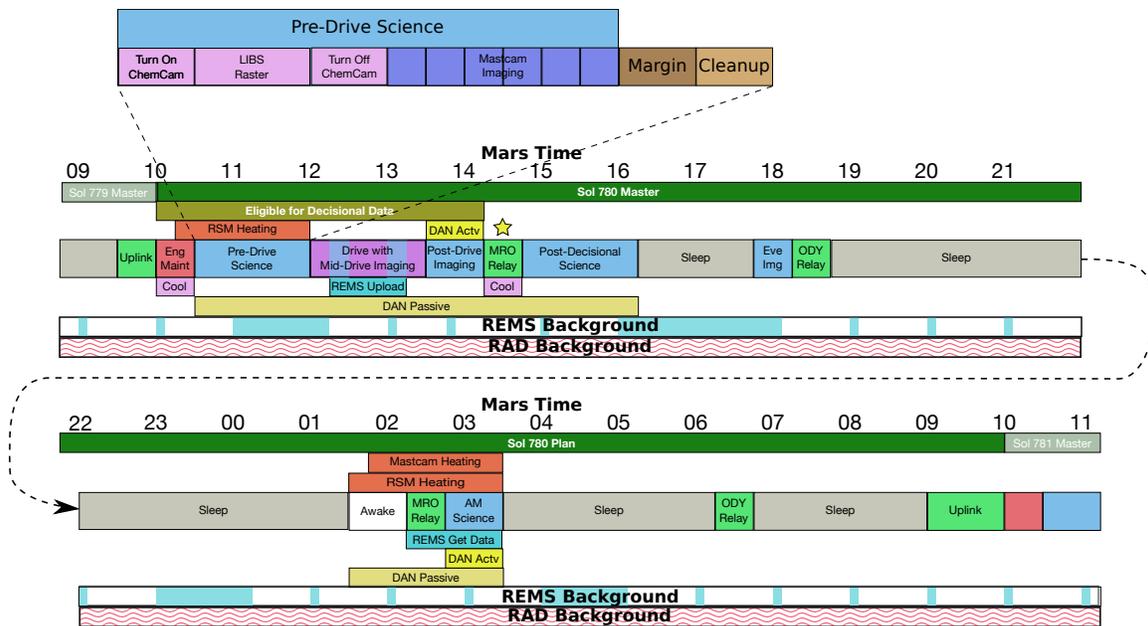


Figure 4: Example sol in the life of the rover.

“AM Science” activity in Figure 4 occurs at 02:45 in the morning of Sol 781, but it is part of the Sol 780 plan and coordinated by the Sol 780 Master.

The rover’s sols are organized into blocks of activities. The Master sequence coordinates these blocks, ensuring associated activity begins at the appropriate time, and forcing activity to stop if it runs too long. As with most sols, Sol 780 begins with a block of “Engineering Maintenance” activities. This typically includes managing onboard data products (e.g. deleting products that have been received on ground, re-transmitting products that were only partially received) and creating data products to document various aspects of the vehicle’s state.

On this particular sol, additional activity is starting in parallel with the engineering activity in preparation for an upcoming “Pre-Drive Science” block. The science block will use the RSM (Remote Sensing Mast) along with the ChemCam and Mastcam instruments. The actuators in the RSM and the electronics in the Mastcam must be warm enough for safe operation. Given the current season and time of day, the Mastcam electronics were warm enough and did not require additional heating. However, the RSM did require heating, so the Sol 780 Master initiated warmup heating during the engineering activity so that the RSM would be warm enough for safe use at the start of the science block. After warmup heating has completed, the rover’s thermal system switches to maintenance heating to maintain safe operating temperatures through the device’s use. In contrast, the ChemCam instrument requires cooling to improve science quality. Therefore, a cooling activity is started to cool down the instrument in time for its scheduled activity.

Following the engineering maintenance block, the Master sequence initiates the “Pre-Drive Science” block. Figure 4 shows a blow-up of this portion of the plan to provide more detail. Most drive sols include a block of pre-drive science. This gives the team an opportunity to collect observations on targets of interest in the current area before the rover drives away. The targets in this block are typically selected from the imagery that was acquired after the previous drive, referred to as **post-drive imagery**.

In this particular pre-drive science block, the team chose to acquire a ChemCam raster of one target and a Mastcam image of the same ChemCam target and a variety of additional Mastcam mosaics. The team typically acquires a Mastcam of each ChemCam target to provide additional context for the ChemCam observation.

The team generally has a reasonable prediction of how long each activity is expected to take. But there are some uncertainties, such as time needed for auto-exposure, auto-focus, and image compression, along with other timing that is not precisely modeled, such as mast slew durations. These uncertainties can add variance in the actual execution time of activities. To guard against this, the team allocates “Margin”, as shown in the blow-up area of Figure 4, so that activity can run a little long without disrupting subsequent activity.

Although unlikely, the team must also guard against the case that an activity runs so long that it exceeds the allocated margin. Therefore, the Master sequence includes a series of “Cleanup” actions to ensure that all activity from the associated block has finished and instruments are returned to known, consistent states in the event they were interrupted. If activity finished on time, or early, the Master sequence still waits for the cleanup time and performs these actions before moving on to the next block of activity. Note that for readability Figure 4 only shows one example of margin and cleanup, but almost all blocks of activity include them.

Figure 4 shows a long DAN Passive observation beginning at the start of the pre-drive science block and extending until the rover goes to sleep. It is common practice for the team to schedule DAN Passive observations whenever the vehicle is awake for at least an hour.

Following the pre-drive science block, the rover started the “Drive with Mid-Drive Imaging” block. Drives do not always include mid-drive imaging, but this was requested by the science team to better characterize the terrain in this area. Following the drive, “Post-Drive Imaging” was performed. This is important as it collects imagery of the terrain surrounding the rover at its new location. The ground team uses this imagery to select targets for the next plan’s pre-drive science (if any) and to plan the next drive. If the team intends to perform contact science, i.e. observations with arm-mounted devices and instruments, at this location, a different set of imagery will be acquired to support the necessary analyses for contact science.

In the middle of the drive, the team has scheduled a “REMS Upload” activity. This activity uploads a **schedule table** to the REMS instrument which is used to control the REMS instrument, scheduling different meteorological observations throughout the sol.

Following the drive is an “MRO Relay” block that represents the rover’s UHF communication window with the MRO relay orbiter. There are usually three to five UHF windows, aka passes, throughout each sol. While the mission has the capability of sending command products to the rover through the relay windows, they are predominantly used for sending data from the rover to Earth. The amount of data that can be transferred in each UHF pass varies based on the type of orbiter (MRO typically has higher bandwidth than Odyssey) and its elevation relative to the rover.

Figure 4 uses a star to mark the post-drive MRO relay window. This is because, due to relay latencies and relative timing between Mars and Earth, this is the last communication pass on this sol that would provide data to Earth in time for the team’s next tactical shift, which would be to plan sol 781. This pass is referred to as the **decisional** pass because any data needed to make decisions for the next tactical shift must be acquired prior to this pass. Note that in this case, the decisional pass was the first UHF pass in the plan. This is not always the case. As noted, it depends on the relay latencies and the current relative difference between Mars and Earth time.

Figure 4 includes the label “Eligible for Decisional Data” to indicate the period of time in which data acquired by the rover may be available by the start of the team’s next tactical shift. How much of this data will actually be available to the team depends on the bandwidth of the relay pass and the size and priority of the data collected. The team carefully considers the predicted size of collected data and assigns priorities to the data products in an effort to ensure the necessary decisional data fits within the predicted data volume of the passes leading up to and including the decisional pass.

Any data collected after the decisional period will not be available for the team to make decisions in the next tactical shift. It may be that some of the data will arrive during the shift. But, except for unusual situations, the team will not base decisions on this data as there is usually insufficient time in the tactical shift to make significant changes late in the shift.

During the MRO pass, the team scheduled another ChemCam cooling activity in preparation for an upcoming science block that will use the instrument. Following this “Post-Decisional Science” block, the rover shuts down to sleep. Shutting down means that the RCE is powered off. The team periodically powers off the RCE to conserve energy and allow the batteries to recharge. Part of the team’s planning process is to determine how much RCE awake time can be supported in the plan given the overall activity for the plan and energy needs for upcoming plans, as tracked by the Supratactical team.

While the RCE is powered on and off throughout the sol, other subsystems on the rover are always powered. Some of the devices that remain powered are required for survival, such as survival heaters, pumps and power management devices. In addition, some instruments are able to operate, and collect data, with the RCE powered off — these include RAD, REMS, APXS, CheMin and SAM. The team can choose to leave these powered on across shutdowns.

The only instruments left on across shutdowns in the Sol 780 plan are RAD and REMS. These instruments are typically conducting **background** activity to collect data throughout each Martian diurnal cycle.

The remainder of the Sol 780 plan follows a similar pattern. The rover wakes up to perform communication passes

and an additional early morning science block. Given the colder temperatures at this time of the sol, longer duration warmup heating is required.

Toward the end of the Sol 780 plan, in the morning of Sol 781, another uplink pass occurs. The team uses this uplink window to load the command products, including the Sol 781 Master sequence, that will be used to control the rover for the Sol 781 plan. The Sol 780 Master hands over to the Sol 781 Master following this uplink window, thus concluding the Sol 780 plan.

2 Case Study Design

The objective of the case study is to understand the factors that contribute to and detract from surface mission productivity. We designed an exploratory case study to investigate this topic [Yin(2014), Thomas(2016)].

MSL is a complex and ever-changing mission. At the time of our case study the surface mission had been operating for almost 4 Earth years, spanning over 1,300 days on Mars. The nature of operations changes as the team pursues different types of science campaigns and as different challenges arise due to the condition of the vehicle and the environment in which it is operating.

The dynamic nature of the mission posed a challenge for the design of our case study. We wanted to obtain an in-depth understanding of the productivity factors that affect surface operations over a broad range of mission and environmental conditions. In order to manage the scope of the study, we chose to use select science campaigns and perform a detailed study of the team's day-to-day activities during this time. We conducted interviews with operational personnel to contribute to the broad perspective of productivity. In this section we describe the selection of the case study campaigns and discuss the methodology used to collect data.

2.1 Selection of Cases

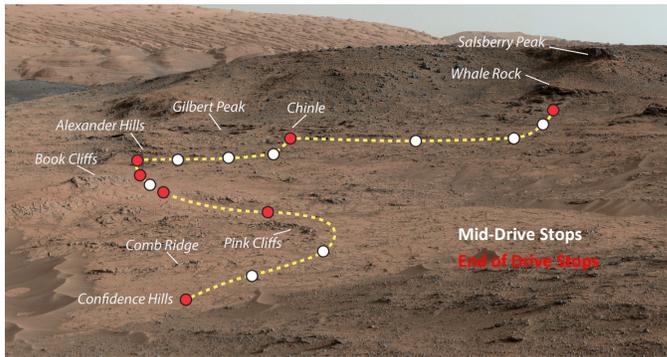
The primary purpose for conducting our case study was to inform our research project in which we will develop technologies and flight operational practices for enhancing surface mission productivity. To manage the scope of our research project, we decided to primarily focus on remote sensing and navigation operations. As such, our selection of case study campaigns was biased toward those that emphasized these types of activities. However, since we are interested in addressing challenges of other types of activity, the selected campaigns also included activities with the rover's arm-mounted instruments.

There are two major types of campaigns that MSL conducts that emphasize remote sensing and navigation: characterizing a geographical area and driving to a distant, strategically identified location. Because different conditions can impact operations, we wanted to be able to compare two different campaigns of the same type, but performed in different contexts. We selected three campaigns, two of which had the objective of exploring a geographical area and a third with the objective of performing a strategic drive. Figure 5 shows the campaigns that we have selected for the study.

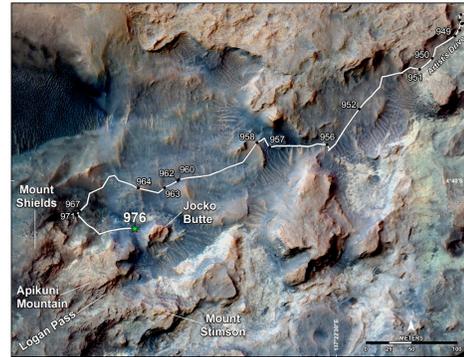
Figure 6 shows seasonal temperatures at Gale Crater as measured by the REMS instrument [Jet Propulsion Laboratory Press Release(2016)]. According to the figure, the Pahrump Hills campaign was conducted in the warmest part of the Martian year. The Artist's Drive and Marias Pass campaigns were conducted at a time when temperatures began to decrease.

Curiosity began exploring Pahrump Hills (Figure 5 (a)) in the fall of 2014 [Stack et al.(2015)]. Pahrump Hills is an interesting case study for multiple reasons. The light-toned outcrop of Pahrump Hills was the first exposure of bedrock making up the base of Mount Sharp that was encountered during the mission. The campaign was also significant in the way in which the exploration of this formation was conducted. The science team decided to conduct a "walkabout", a practice used by field geologists when studying unexplored geological areas on Earth. The team made multiple passes of the area with each pass informing a subsequent, more detailed study. We chose to focus on the first walkabout which explored the region with primarily remote sensing instruments (mast-mounted imagery and spectroscopy) in order to identify locations to return to for more detailed follow-up study with arm-mounted instruments. Another interesting factor was the geography of Pahrump Hills was conducive to developing a strategic plan for the initial walkabout. The sloping hills made it possible to see nearly the complete formation from a single panoramic image allowing scientists and engineers to plan a route to explore the area.

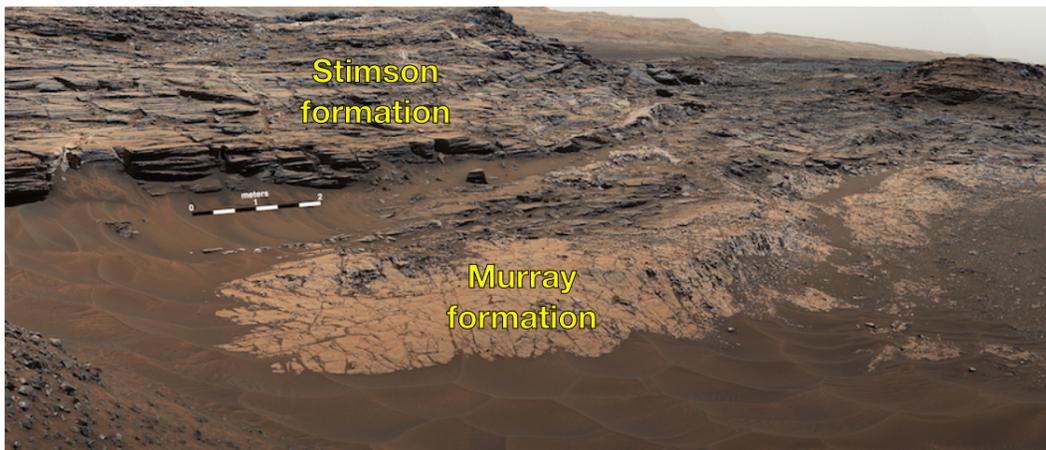
After completing investigations at Pahrump Hills, Curiosity departed the area in the spring of 2015 with the objective to reach higher levels of Mount Sharp for continued exploration. Curiosity followed a route referred to as Artist's



(a) Pahrump Hills
 Martian Season: Spring
 Sols: 780 - 798 (19 sols)
 Drive dist: 152m



(b) Artist's Drive
 Martian Season: Late Summer
 Sols: 949 - 972 (24 sols)
 Drive dist: 567m



(b) Marias Pass
 Martian Season: Late Summer, Early Fall
 Sols: 991 - 997; 1027 - 1043 (24 sols)
 Drive dist: 130m

Figure 5: Selected campaigns for case study (NASA/JPL-Caltech/MSSS/Univ. of Arizona).

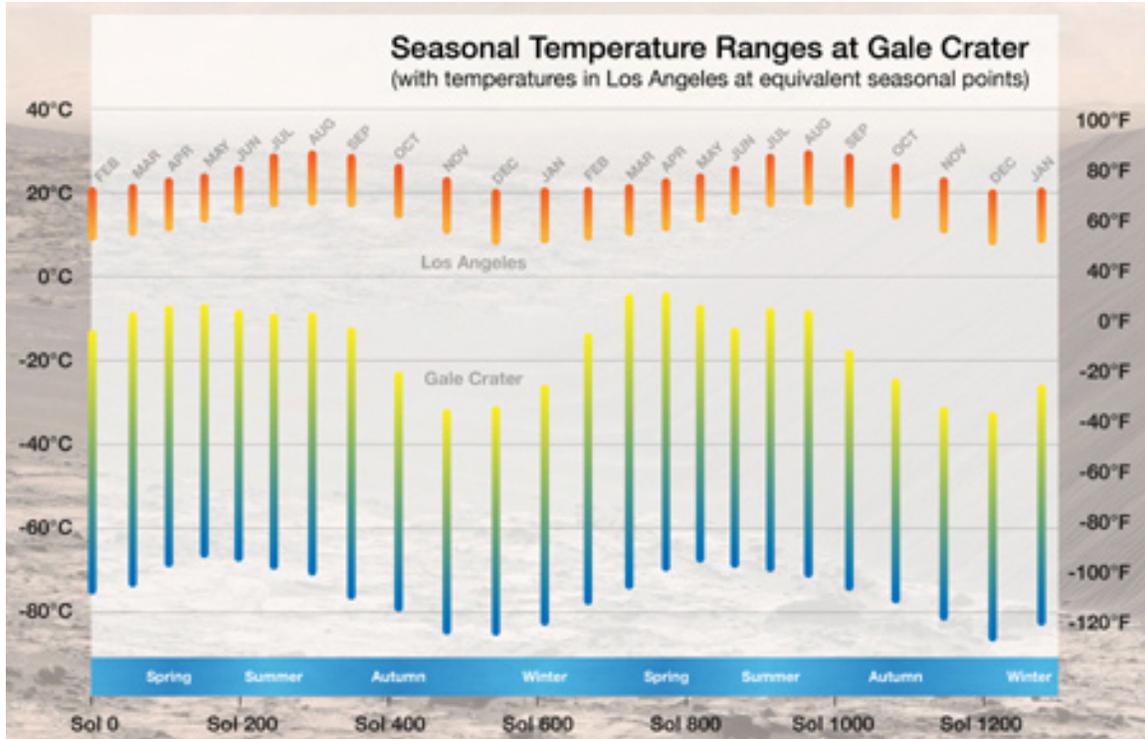


Figure 6: Seasonal temperatures at Gale Crater (NASA/JPL-Caltech/CAB(CSIC-INTA)).

Drive (Figure 5 (b)) [Jet Propulsion Laboratory Press Release(2015j)]. Along the way, the science team conducted a science campaign to capture images of the surrounding topography in order to build a record of the stratigraphy (i.e., layering and structure) of the sedimentary rock layers exposed in the valley walls. The orbital imagery provided by the Mars Reconnaissance Orbiter's HIRISE instrument enabled the team to identify locations where gaps in the surrounding terrain provided the opportunity for imaging the far terrain.

While the orbital imagery enabled the team to develop a strategic route for driving through Artist's drive, the geography of the region made it more challenging for day to day (also referred to as tactical) driving. The orbital data provides good information about the general terrain the rover will encounter, but it is not sufficient resolution for the actual drive path planning. Instead, images acquired from the end of the previous drive are used, in conjunction with orbital data, to plan the next drive. Ridges and valleys in the surrounding terrain often prevented the rover from getting a good view of the terrain in which it would be driving the next day, and wheel wear concerns limited the desirability of using onboard autonomous hazard detection to extend drives into unseen terrain. This often made it difficult for the engineers to plan the next drive path.

Along the route toward higher levels of Mount Sharp, Curiosity took the opportunity to explore an area where the Murray formation (the type of rock from Pahrump Hills) came into contact with an overlying geological formation called the Stimson formation. The contact was explored in an area named Marias Pass (Figure 5 (c)) [Milliken et al.(2016)]. This campaign has interesting similarities and contrasts with the earlier Pahrump Hills campaign. Both campaigns sought to explore and characterize a geological area. However, the more challenging terrain in the area and discoveries made during exploration resulted in a more dynamic campaign than the Pahrump Hills walkabout campaign.

2.2 Conceptual Model of Achieving Objectives

Although we were conducting an exploratory case study, and wanted to minimize unnecessary bias in our study, it is important to have a theory of the topic under investigation to help guide the research [Yin(2014)]. The theory helps identify data to be collected and provides a context for interpreting the results of the study.

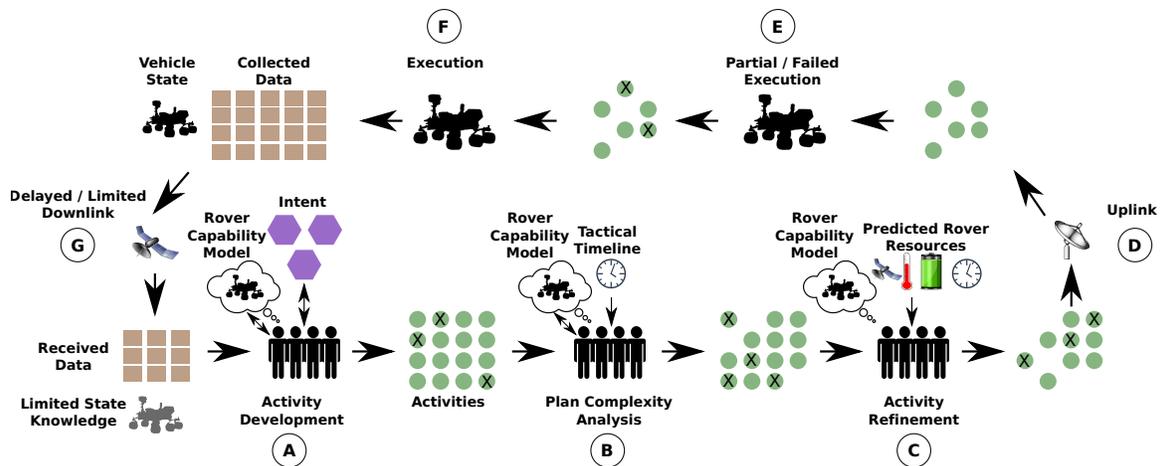


Figure 7: Factors impacting surface mission productivity.

For the purpose of our case study, we began with the general notion that mission productivity had something to do with the team’s ability to accomplish mission objectives. This can include how much effort is required by the team to accomplish a given objective as well as how long it takes, e.g., number of sols, to achieve objectives. We would, of course, be exploring the definition of productivity more objectively as part of the study, but this seemed like a reasonable place to start to form our guiding theory.

We expanded on this premise by developing a conceptual model of how objectives are accomplished in a surface mission. This model is shown in Figure 7. Several of the authors have worked surface operations on the MER and MSL missions and this model is based on the authors’ experience. We included the evaluation of this model in the case study. We presented it to the operations personnel we interviewed to get their feedback and we looked for instances in our data collection where data did not fit well with this model. As will be discussed, we identified ways in which this model should be expanded to encompass additional productivity factors, but we did not identify errors in the productivity factors originally described in the model.

The general flow of the diagram begins with the team identifying candidate activities that can be used to accomplish their intent. These activities are developed and refined during operations planning until a set of command products is ready to be uplinked to the vehicle. The vehicle executes these activities and produces results which are conveyed back to Earth through telemetry and data products. This information, in turn, is used to support the development of subsequent activities and, potentially, new intent. The crossed out activities illustrate typical stages in the conceptual model in which activity is limited in some way. During operations planning, this can include restricting the scope of an activity, deferring an activity to a later planning day or even descopeing an activity entirely. During execution, it can include partial or complete failure of an activity. The following subsections describe each stage and the factors that can limit productivity in more detail.

2.2.1 Step (A): Activity Development

The diagram begins with Activity Development, in Step (A), where the team considers activities that could be performed that would contribute toward achieving their objectives. Objectives may be science objectives, such as characterizing a geological formation, or engineering, such as performing a vehicle maintenance operation of subsystem inspection.

In terms of the operations timeline discussed in the previous section, Activity Development can occur as part of the Supratactical timeline, delivered to Tactical in the Skeleton plan, or during the early stages of the Tactical timeline.

Throughout the planning process, the team makes use of their knowledge of the vehicle’s capabilities to help develop command products that the rover will be able to achieve. In the Activity Development stage, this knowledge is used to help determine if an activity is feasible given the abilities and limitations of the science instruments and other actuators. This includes, for example, understanding the detection sensitivity of science instruments and knowing the range of slopes the mobility system can safely traverse. At times, the team will use their vehicle model to come

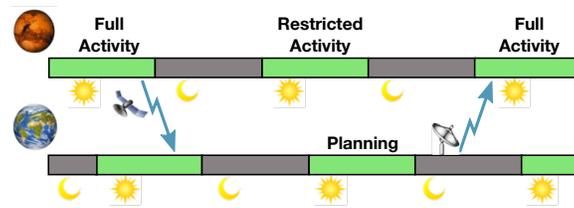


Figure 8: Mars activity vs. Earth planning.

up with creative new ways of using the vehicle’s capabilities in ways not previously considered. For example, after landing the team developed methods for driving and performing arm operations with sample cached in the sampling system. In other examples, the team developed a technique for using the MAHLI instrument as a goniometer [Johnson et al.(2015)], as well as using the rover’s inertial measuring unit to perform a gravimetry survey [Lewis et al.(2016)].

Depending on the type of activity, varying levels of knowledge of the current state of the vehicle may be required. For example, in order to select specific targets for the mast-mounted instruments, knowledge of the position the rover will be in along with navigation images of the surrounding terrain is required. Similarly, activities related to using the arm in close contact with the surface typically require up to date knowledge of the rover and the terrain. In contrast, many activities such as systematic survey imagery and atmospheric measurements do not require as extensive knowledge of the current state of the vehicle.

The level of rover state knowledge required to accomplish an activity may prohibit certain activities from being accomplished on a given sol. The amount of knowledge about the state of the vehicle may depend on the downlinked data from the previous plan as well as activity in the current plan. For example, the downlink which contains the latest information about the state of the rover from the prior plan may be delayed, or the communications window between the rover and relay orbiter may have not transferred sufficient data to support all the desired activity. Or, the current plan may include an event that changes the state of the vehicle, e.g. by driving, such that insufficient state knowledge will be available for performing certain activities, e.g. ground-targeted imaging, after the event.

The rover typically performs activities that result in significant changes to its state during the daytime. There are usually one or two communication windows with relay orbiters during the latter part of the day which allow the rover to relay its latest state and other collected data to Earth. Under what is referred to as “nominal” operations, this data will be received by the operations team in the morning on Earth. The team on Earth will then have all day, during the rover’s night on Mars, to develop command products that will be sent to the rover during the next morning on Mars. Mishkin has referred to this as working the Martian night shift since the operations team works during the Martian night [Mishkin et al.(2006)].

A significant factor in the availability of vehicle state knowledge is the relative duration of a day on Earth and a day (aka sol) on Mars. A Martian sol is approximately 40 minutes longer than an Earth day. As such, if the operations team wishes to continue to work during the Martian night, they must continually shift the times in which they work on Earth. For example, if the team starts their shift at 8:00am one day, they would start their shift at 8:40am the next. Subsequent shift start times would be 9:20am, 10:00am, 10:40am, etc. Over the course of about a month, the team will have transitioned their shift start times around the clock. This mode of operations is referred to as “working Mars time” and is highly taxing to the team. Due to the stress this mode of operations places on the team, the MER and MSL missions limited Mars time operations to the first 3 months of the mission.

The vast majority of the surface mission is conducted with the team restricting operations to the daytime on Earth. The consequence is that the operations team is often out of sync with the activity of the rover on Mars. Figure 8 illustrates the impact this can have on the data available to the team during planning. In the diagram, the end-of-day relay from the rover arrives on the ground late in the Earth day. If the team had still been working Mars time operations, they would arrive to work at this point and begin the tactical process, working through the Earth night to develop command products in time for uplinking to the rover in the morning of its next sol. Instead, the team waits until the next Earth day to begin planning. Meanwhile the rover is waking up for its next Mars day without a new set of command products from Earth. By the time the team has completed the tactical process, they must wait for the subsequent Mars morning to uplink the products to the vehicle.

This often limits what the team can command the vehicle to do during the middle sol of Figure 8. If the vehicle were allowed to make significant changes to its state, in particular driving to a new location, this would significantly

limit the types of activities the team could command on the subsequent sol. These limited activity sols are referred to as “restricted sols” because the latency of data often restricts the type of activity the team can perform.

A similar situation arises when the team takes days off for weekends and holidays. In these cases, the team will create plans that span multiple sols (aka multi-sol plans). Again, activities that result in significant changes to vehicle state are limited since they will impact the activity that can be done in later sols of the plan.

2.2.2 Step (B): Plan Complexity Analysis

Step (A) discussed how the activities considered depend on the general capability of the rover and the team’s knowledge of its state. In Step (B), the team considers the complexity involved in implementing the activities under consideration. As with Step (A), this step may begin with Supratactical planning and continue into the early stages of Tactical operations. The purpose of this stage is to help ensure the team does not take on more activities than can be completely planned and validated during the scheduled shift duration. If the team attempts to perform too much activity it may result in the team not completing the tactical timeline and thus risk missing the next uplink deadline. Or it could mean overloading the team which could lead to mistakes.

As such, the team carefully evaluates the activities it chooses to work on such that they can be completed during the tactical timeline. It is extremely challenging to pick an appropriate set of activities that allows the team to maximize what can be accomplished during the tactical shift without exceeding the capacity of the timeline. It requires a lot of experience and good judgment to make these decisions. Further challenging the decision making is the fact that what can fit in the tactical timeline is continually changing. The first time a new type activity is performed will require more focus and effort from the team. But after that type of activity has been performed several times, it may consume much less of the timeline. New ground tool developments can also result in increasing the capacity of the tactical timeline.

2.2.3 Step (C): Activity Refinement

During Activity Refinement, Step (C), the team takes into account the vehicle resources that are required to perform the proposed activities. The resources the team considers include energy, data volume and time available to perform activities. This stage uses models of the rover and activities to make predictions about the amount of resources that activities will consume and the amount of resources available on the vehicle.

Some of the resource constraints are more or less fixed. For example, the team avoids depleting the battery below a certain level and filling up the data product file system. Other constraints are more transient. For example, the Supratactical team often provides guidelines on the battery state of charge to maintain at the end of the plan. This end-of-plan battery constraint will vary from plan to plan, depending on the activity in the Supratactical look-ahead plan. Similarly, the team may self-impose tighter data collection constraints on itself in preparation for data intensive plans that are known to be upcoming in the near future.

The fidelity of the models used in the stage of operations play an important role. Missions tend to be conservative in their estimates to avoid inadvertently exceeding available resources. The model may overestimate the time and energy consumed by an activity, e.g., by allocating an overly generous margin of time around it. As a consequence, this stage may over-prune activities because the model predicts they would exceed resource constraints when in practice there may have been sufficient resources available.

2.2.4 Step (D): Uplink

Given the complexity of communicating with a tiny robot on a distant, spinning planet millions of kilometers away, the missions have a remarkably reliable channel for sending command products to the rover. However, problems can arise that result in loss of activity during uplink. Noise encountered on the millions of kilometers trip can corrupt the signal beyond the means of error correction codes to correct. Equipment failures on Earth stations can occur with insufficient time to repair before the uplink window. In general, the amount of data required to uplink is very small compared to the amount of data downlinked from the rover. However, there are still rare situations in which the capacity of an uplink window is insufficient to transmit all the desired command products. In each of these cases, some or all of the commanded activity can be lost.

2.2.5 Steps (E) and (F): Execution

After receiving command products from Earth, the rover begins executing the new plan and collecting new data. The commands products are mainly in the form of sequences, files containing lists of commands to execute.

In the large majority of cases, execution proceeds as expected and the rover is able to achieve the desired results. At other times, activities may have partial success or completely fail. There are a variety of causes of unsuccessful execution. Sometimes the command product may have included an uncaught command error which results in a problem during execution. Other times, the current state of the vehicle may have been unexpected. Sometimes sequences are written to take into account uncertainties in the state of the vehicle, but such sequences add complexity to develop, consuming the capacity of the timeline, and the expressivity of the sequencing language can limit what can be sequenced. Different activities have increased levels of autonomy to account for unexpected conditions. For example, the rover is capable of autonomous navigation, which enables the rover to drive to locations without a prior knowledge of the terrain through which it will traverse.

Throughout the plan execution, the rover will produce and collect data products which record the results of its activities. It will also generate telemetry, which includes critical information about the state and health of the vehicle. All of this data is stored onboard awaiting transmission to Earth.

2.2.6 Steps (G): Downlink

The vast majority of the data received from the rover is sent via relay from one of the Mars orbiters. The rover must wait for the orbiter to fly overhead before it can transmit data. The amount of data that can be transferred to the orbiters varies with each window depending largely on the elevation of the orbiter in the sky as it passes over the rover. Data is prioritized by the operations team such that information critical to assessing the health of the vehicle and for planning the next sol's activities is sent earlier in the communication window.

Once the data is onboard the orbiter, it must wait for the orbiter to have a communication opportunity with Earth before the data can finally reach Earth and then get transferred to the operations team for analysis. As with uplink, technical problems may occur during downlink which can result in unexpected delays in data reaching the operations team.

The data becomes input to the next round of planning. It may be used to support the development of further activity, e.g. an interesting target for further study may be identified in a downlinked image, and it may result in new high-level objectives being formed, e.g. unexpected signatures in a spectral analysis may result in a new objective to characterize an area.

2.3 Data Collection Methodology

We used the conceptual model in Figure 7 to develop a schema, shown in Figure 9, to guide data collection for the study. The green boxes in Figure 9 represent the stages of major stages of the activity development and execution process from Figure 7. The blue boxes are the input and output from these stages.

We derived a spreadsheet from this schema that was used to focus the collection of data and organize it after it was gathered. The spreadsheet was organized with a worksheet for each entity in the schema with columns for the attributes of the entity. Each row of a worksheet captures an instance of that entity. Links among entities were captured by references to their identifiers. Following is the schema defined in the spreadsheet. Note that some of the information was manually collected by a member of the team. In other cases we were able to write scripts to automate the collection process. We identify those items below with the tag “(automated)”. In most cases, the automation was made possible thanks to MSL's extensive Elasticsearch database.

Intent: A science or engineering objective.

- Unique identifier
- Description of intent
- Type of intent (e.g. engineering, geological science, environmental science)
- Parent identifier (to reflect hierarchical relationships among intent)
- Link to source documenting the intent

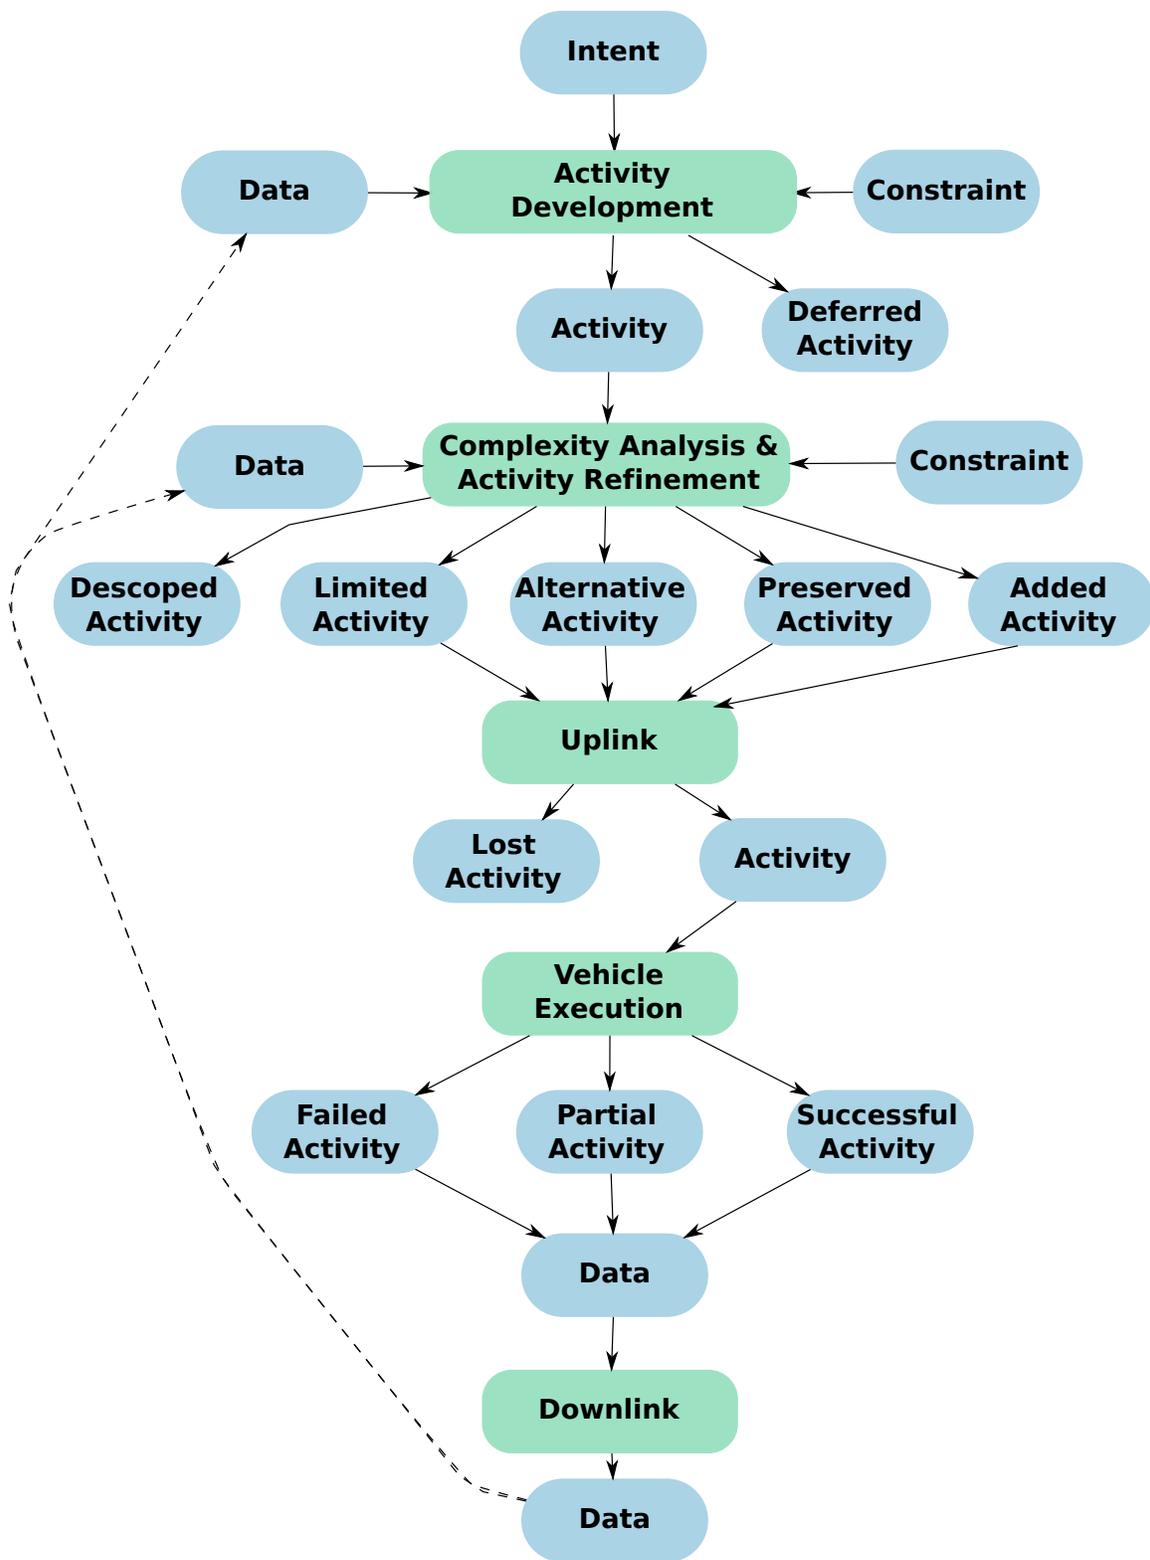


Figure 9: Data collection schema.

Constraint: A constraint on the plan activity

- Unique identifier
- Description of the constraint
- Type of constraint (e.g. Power, Time, Data Volume)
- Sols for which the constraint applies
- Link to source documenting the constraint

Activity: An activity identified to achieve an objective. This includes activities that were considered by the team but, for some reason, did not end up being executed on the vehicle. As such, some of the fields may not be applicable to all activity instances.

- Unique identifier
- Description of activity
- Intent(s) this activity supports
- Activity name in plan (if applicable) (automated)
- Sol in which activity was developed
- Link to source documenting the activity and the intent(s) is supports
- Data items that were used to develop this activity (e.g. images used to plan a drive, images used to target a remote sensing observation) (automated)
- Predicted resource allocations (energy, duration, data volume) (automated)
- Planning outcome specifying what happened to this activity during the planning process. This can be: preserved (activity was retained in final plan), descoped (activity was dropped from the plan), deferred (activity was dropped from the plan with the intention of performing in a future plan), limited (activity was restricted in some way), alternative (activity was created to replace a different activity that was descoped)
- Stage in the planning process in which the activity outcome changes (for activities that were not preserved)
- Reason for the activity's outcome (for activities that were not preserved)
- Link to source documenting the activity's outcome (for activities that were not preserved)
- Sequence identifier (for sequences that were preserved)
- Uplink status to document cases where an activity was uplinked but did not make it onboard the vehicle
- Execution outcome indicating success, failure or partial success (for activities that made it to the vehicle)
- Link to source documenting execution outcome
- Data items that this activity generated during execution (automated)
- Actual resource used (energy (estimated), duration, data volume) (automated)

Data: A piece of data generated by the vehicle. All data information was automatically collected.

- Unique identifier (automated)
- Activity that created this data (automated)
- Priority of this data (assigned by the operations team) (automated)
- Time at which data was created (automated)
- Time at which the data was received on Earth (automated)
- Size of data (automated)

Misc: Miscellaneous item not fitting into pre-defined schema. This was intended to capture items of interest that were found while gathering data but did not fit the established schema. These were typically elaborations from the team on decisions that were made (e.g. strategies for planning a drive or selecting targets).

- Unique identifier
- Description
- Link to source documenting the item

The MSL mission provides a variety of sources for the above information and the project generously made these sources available to our team. We made use of the following sources of information:

Activity Plans: All of the mission’s activity plans are stored in a database and accessible through the MSL InterfaCE (MSLICE) integrated planning tool. These include the look-ahead plans used in the supratactical process and plans for each stage of plan development through the tactical timeline. The plans contain detailed information about the activities that were developed during a tactical shift included predictions for the activities’ timing and energy and data volume use. The plan also includes the sequence identifier for each activity which provides a link from the activity as represented in the activity plan to the sequence that implements that activity. Because it is the sequence that is uplinked and executed on the vehicle, the sequence identify provides a primary link to vehicle telemetry associated with the activity. For targeted activities (e.g. many ChemCam observations) the plan includes a link to the image data product that was used to target the observation.

MSL Reports: MSL uses a web-based reporting facility, called MSL Reports, for documenting operations. MSL Reports contains a vast amount of information spanning the full scope of operations from strategic to tactical. It is used extensively throughout operations to facilitate communication during the development of operational products and to serve as historical documentation for future reference. Information about tactical and supratactical operations is organized by sol and, within each sol, by operations role. One of the main uses of MSL reports for our case study was to determine the intent for each sol of the campaign and what activities were considered for achieving the intent. While the activity plans contain a wealth of detailed information about planned activities, they provide only a limited amount of information for why the activity was performed. Fortunately, this intent information was usually captured in MSL Reports. The reports were also a valuable source of information about alternative activities that were considered but descope. The reports also provide important information about challenges the team faced during operations, background for the decisions that were made and assessments of the vehicle’s execution of the plan. As such, a large part of our data collection process was reading through MSL Reports to gather relevant information for each sol of the selected campaigns.

Vehicle Telemetry: The rover collects a massive amount of information as it performs activity. The information is collected in a variety of forms. Data products are a collection of related information used for both science results (e.g. images, spectra) and engineering data (e.g. statistics on performance of a subsystem). EHA (Engineering Housekeeping and Accountability) provide channelized telemetry on vehicle state (e.g. power loads, temperatures). EVRs (Event Records) document discreet events (e.g. the beginning and ending of a sequence, dispatch of commands), similar to print statements in a programming language. We made use of this telemetry to generate many of the resource analysis plots in this report. In particular, telemetry provided the source of information for actual resource usage.

The actual process of data collection was performed with a combination of manual and automated processes. Because the documentation of intent, decision rationale, and operations challenges was stored as free-form text within MSL Reports, a manual process was used to collect this information. The sols of each campaign were distributed among members of the team. Each team member worked sol by sol, reading through MSL Reports to identify the intent for each activity in the sol’s plan. The team member would also look for documentation about alternative activities that were considered or ways in which the selected activities were constrained in some way (e.g. a drive that was not as long as desired or contact science target that was not on a preferred location).

We developed several scripts to automate a large portion of data collection. The data supporting the resource analysis of each campaign was primarily collected through this automation. Our script development was facilitated by MSL’s Elasticsearch database which indexes the activity plans and all vehicle telemetry.

2.4 Caveats

It is important to note some limitations in the case study. First and foremost, as noted previously, MSL is a vast and varied mission. While we believe we have selected campaigns that reflect issues common throughout the mission and

are relevant to future surface missions, the campaigns certainly do not capture all aspects of the mission. In particular, we did not include sampling campaigns in our study. This would be a valuable area to consider in future extensions of this work.

One of the biggest challenges we faced in the data collection process was determining the intent for a given activity. Operations personnel were generally very good about documenting in MSL Reports the specific reason an activity was included in the plan. However, it was not always clear how this rationale related to specific campaign objectives. For example, a given campaign may have a number of questions to be answered. Given an activity in the plan, it was usually straightforward to determine if the activity pertained to the campaign or not. But it was not as clear which specific campaign question, or set of questions, it was intended to address.

We eventually concluded that, for our objective of understanding productivity factors affecting achieving objectives, it was sufficient to perform a “high level” categorization of activity intent. As such, we defined three high level intent categories for activities:

Campaign Science: Activities that directly contributed toward the campaign’s objectives

Other Science: Activities that contributed to science objectives unrelated to the campaign under study

Engineering: Activities carried out primarily for monitoring and maintaining the health of the vehicle

Not that “Other Science” can also include activities contributing to a campaign, but a campaign other than the one we selected for study. For example, during the Pahrump Hills Walkabout campaign, the team was conducting a parallel campaign to study Comet Siding Spring.

Even this high level categorization of activity intent presents room for interpretation. In particular, it was not always obvious if the intent of an activity should be engineering or one of the science categories. For example, should an drive activity be considered science, because it is moving the rover to a location where science data can be collected, or engineering, because it is focused on mechanical operation of the vehicle?

Because MSL is a science mission, almost everything performed by the vehicle is ultimately for the purpose of accomplishing science. However, we felt it was more informative, from the perspective of understanding mission productivity, to make a distinction between activities that are primarily focused on monitoring and maintaining vehicle health and those that more directly enable the collection of science data.

Based on this guideline, we considered a drive activity to be achieving science intent if the purpose of the drive was to reach a location to conduct science. In contrast, a drive that was performed to monitor wheel health was categorized as engineering.

Another important limitation of the study is that our information about what activities were considered during a given planning session is limited to what was documented in MSL Reports. While the reports provide information about activities that were part of an earlier plan but removed before uplink, and they often document activities that were considered but not included in a plan, they do not document all activities that the team considered. By this point in the mission the team had developed significant experience on the activity that could be included in a plan due to complexity limits and activity constraints. As such, there is a significant amount of activity pruning that is never documented. So, while we recorded information about deferred and descoped activities that were documented in the reports, it must be recognized that there is a significant amount of deferred and descoped activities for which we were unable to collect information.

3 Interviews with Operations Personnel

3.1 Interview Design

The objectives of the interviews was to gain a broad perspective on mission productivity that would complement the more detailed investigation of productivity in our campaign studies. Toward that end, we selected members of the MSL operations team that performed roles that had a strong connection to productivity. The operations roles we chose were:

Scientist: The science team identifies the science objectives of the mission and the activities to accomplish them. The science team is organized into two major theme groups, environmental and geo-chemical. We chose to interview members of the science team due to their direct involvement with identifying and achieving science objectives.

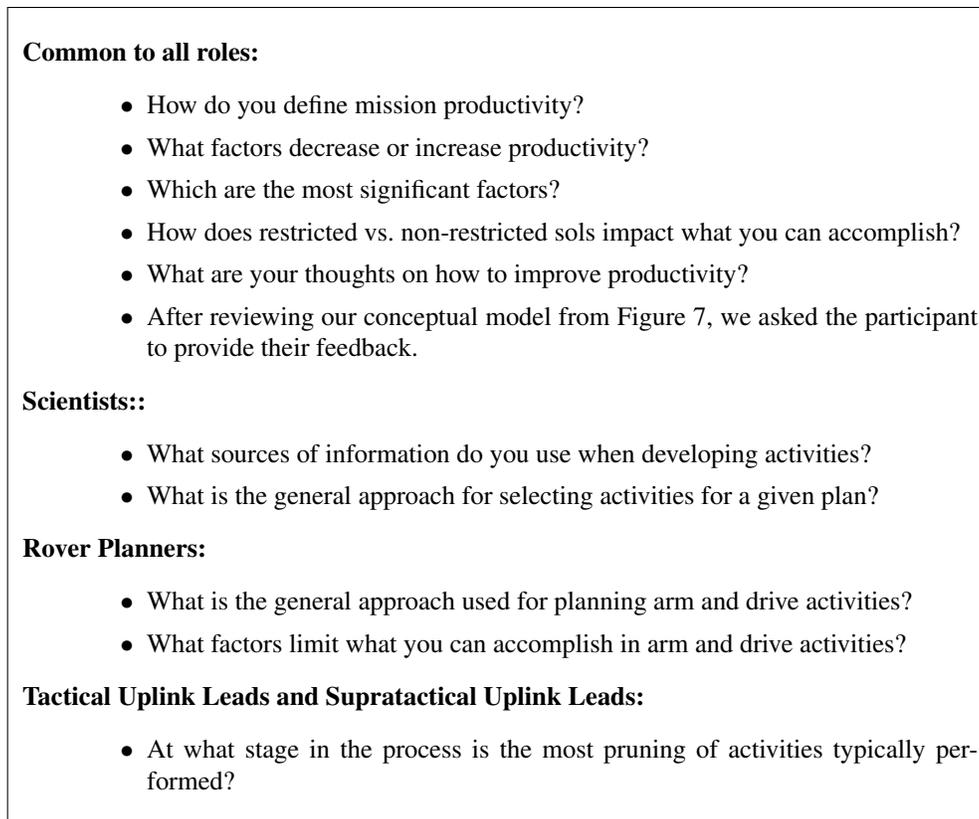


Figure 10: Interview questions.

Rover Planners: Rover Planners are members of the engineering team who create plans and sequences for operating the rover’s arm, sampling and mobility subsystems. We chose to include Rover Planners in our interviews due to the high degree of complexity in planning these subsystems and the strong dependence on terrain interactions.

Tactical Uplink Leads: The Tactical Uplink Leads oversee and coordinate tactical uplink operations. They are responsible for ensuring that the developed command products achieve intended objectives while observing vehicle health and safety. We included Tactical Uplink Leads in the interviews as they are responsible for managing the tactical timeline, including determining how much activity can be accommodated within tactical timeline and vehicle resource limits. In addition, because their role spans the full scope of uplink operations, they have a broad perspective on operations.

Supratactical Uplink Leads: The Supratactical Uplink Leads oversee the supratactical operations. They coordinate science and engineering objectives over a horizon of roughly 7 to 10 sols. We interviewed Supratactical Uplink Leads because, similar to Tactical Uplink Leads, they also make decisions regarding how much activity can fit within the scope of the tactical timeline and predicted vehicle resources as well as having a similarly broad perspective of operations.

We conducted a total of thirteen separate interviews, each interview running between 45 and 90 minutes. We selected 5 members of the science team (three with a geo-chemical emphasis and two with an environmental emphasis), 3 Rover Planners, 1 Tactical Uplink Lead, 1 Supratactical Uplink Lead and 3 team members who performed both Supratactical Uplink Lead and Tactical Uplink Lead roles.

The interviews were structured as guided conversations. We had a set of questions to serve as starting points and to ensure a common set of area was covered in each interview, but each interview proceeded according to the responses of the different participants.

Below we summarize the outcome of the interviews.

3.2 Defining Mission Productivity

The majority of participants described mission productivity as a measure of how much science is accomplished each sol, or how many sols it takes to achieve science objectives. This description is consistent with the definition of productivity as a measure of output per unit of input ([Mankiw(2014), Reid & Sanders(2012)]) with sols taken as an aggregate measure of input.

Two participants described productivity in terms of accomplishing objectives better, faster, and with less “churn”. This definition may be more directly related to efficiency, a measure of the amount of input required to produce output. Productivity and efficiency are closely related terms, with productivity defined as an average measure of the efficiency of production.

One participant also took into account how much science output could be maintained over multiple days or over the course of a campaign without burning out the team. Team longevity is a major concern for long-term operations. For example, team longevity is one of the primary reasons why surface operations does not maintain Mars-time operations (Section 2.2.1) for an extended period of time. Overall mission output may be increased with Mars-time operations over a short window of the mission, but overall a larger window, would result in a drop in output due to team fatigue. This emphasizes the importance of considering productivity over an appropriate duration of time.

One participant noted that studying mission productivity is challenging as one must take into account how much activity was being requested of the vehicle over a given time period. For example, the rover may not be very active on a given sol because the operations team needed more time to make a major decision regarding future activity (e.g. whether to stay in the local area for additional investigations or drive off to a more distant location). In these situations, the rover’s low level of activity may be a result of the operations team not requesting as much from the vehicle at this time. Given that even these lower levels of vehicle activity still require mission resources, and some rover resources (in particular MMRTG output) degrade over time, it seems reasonable to consider such situations as representing lower productivity, but it is important to identify the root cause of this lower productivity appropriately.

3.3 Identifying Productivity Factors

We asked the participants to describe the factors that contribute toward and detract from productivity. Following is a summary of their responses. These are presented in no particular order.

Ability to Exploit Supratactical Work: Several participants observed that MSL productivity is heavily dependent on how well tactical operations can exploit supratactical work. A significant amount of work is performed as part of supratactical operations that can feed into the tactical timeline, thereby increasing the capacity of activity that can be achieved tactically. This work includes creating the basic structure of the plan, in the form of the skeleton, that provides guidance organizing activity and allocating resources. The science team performs preliminary activity development based on guidance from the skeleton prior to the start of tactical operations. Supratactical Rover Planners may be able to begin development of the next plan’s drive or contact science activities. The interview participants noted that this is especially the case if the vehicle is in the same place performing contact science, as the supratactical Rover Planners will be able to develop complete sequences for upcoming activities. All of this work results in a virtual extension of the tactical timeline, increasing the capacity of the activity that can be achieved.

However, the extend to which the tactical timeline can benefit from the supratactical work is dependent on how closely the actual tactical plan follows the plan that was anticipated by supratactical operations. If recently received data prompts the science team to change their mind on their objectives or problems are observed in the previous plans execution, that tactical plan will deviate from the skeleton. This can require the tactical team to re-work resource allocations and may prevent the team from using pre-developed activities. The result is that the team is not able to accomplish as many objectives.

Overhead in Command Product Generation: Many of the participants discussed the effort involved in developing and validating command products. They noted that ground tools have been developed that provide support for this process. However, they noted that when problems occasionally occur with the tools, it can cause delays in the tactical timeline. Others noted that the flight rule checking process can be burdensome. The mission contains many flight rules that must be checked manually and many false or intentionally violated flight rules that must be dispositioned and documented. The time required for this work reduces the amount of activity that can be brought through the tactical timeline.

Managing Tactical Timeline Complexity: Several participants brought up the challenge of managing tactical timeline complexity. This refers to evaluating how much effort will be required to bring different types of activities through the tactical process. Timeline complexity is managed both during supratactical planning, when the Supratactical Uplink Lead distributes objectives across the look-ahead plan, and during tactical planning, when the Tactical Uplink Lead makes decisions about activity to include in the days planning.

Through experience, the team has developed complexity guidelines to help judge the complexity of different types of activity and provide consistency in complexity management across the team. While the participants noted the value in the guidelines they also noted its limitations. Some of the Tactical Uplink Leads and Rover Planners noted that the guidelines do not take into account variations in experience of people staffed for a particular tactical shift. It is natural that more experienced team members will be able to accomplish more in the tactical timeline. However, supratactical planning does not make assumptions about specific staffing as staffing could change on the day of the tactical shift. It was noted that the tactical team always has the option of making changes, potentially adding more activity. However, as noted above, this additional work may not have the benefit of preliminary development from supratactical operations. It was also noted by participants that the complexity guidelines do not factor in terrain-dependent complexities for Rover Planner activities, as such the guidelines may underestimate how complex a planned activity will be.

Predicting Available Vehicle Resources: Almost all of the Tactical Uplink Leads and Supratactical Uplink Leads we interviewed commented on the impact that predictions of available vehicle resources had on productivity. They observed that inaccuracies in power modeling often results in unnecessarily restricting planned activity. The predictive model may indicate that the battery state of charge dips below the allowed minimum limit or does not reach the battery state of charge guideline for handing over to the next plan that was requested by supratactical. However, it is usually the case that actual telemetry from the vehicle shows that the actual battery state of charge was higher than predicted which would have allowed for additional activity. Some of the Tactical Uplink Leads noted that this issue will become more significant as the MMRTG output decreases.

We spoke separately with a representative from the MSL Power subsystem to understand the challenges in predicting battery state of charge. There are a few factors involved. The predictions make use of models for the energy loads of individual activities. The engineers providing these models often build in conservatism to avoid the risk that their model would result in an under-estimate of energy use which may lead to plans that over-deplete the battery. One of the challenges in refining these models is that the telemetry from the vehicle provides only the aggregate load of all currently running activities on the vehicle. If multiple activities are occurring in parallel (e.g. CPU running, camera imaging, survival and other actuator heating, system pumps, etc.) it is difficult to determine which portion of the aggregate load is due to each individual activity.

The predictions of power consumption also depend on models for predicting how long an activity will take, as if the activity runs shorter than predicted, it will often use less energy than predicted. Thus, models to predict activity duration provide another source of error in power modeling. Related to duration modeling is the challenge of modeling the thermal environment of the rover in order to predict the duration of heating that will be required. As with models of energy load, the thermal models are also conservative to avoid the risk of using an actuator with insufficient heating. Further, it requires significant effort to produce thermal tables used for the different Martian seasons, as such, the mission uses a small number of tables, e.g. a table for winter and another for summer. These tables reflect the coldest periods for the sol range that they cover. As such, much of the mission's operations is conducted using thermal tables for colder temperatures than the rover is currently experiencing. This impacts power modeling because the conservative thermal tables results in over-modeling of how much duration is needed to warm up devices, and as a consequence over-modeling how much energy is required to do so.

A third factor is the model used to derive the battery state of charge value itself. The value is an inexact approximation and additional noise is introduced as the value is propagated across the horizon of a plan causing further deviations from the actual state of charge. It was noted that the accumulation of error stop at points when both the predicted and actual telemetry indicate the battery is at 100% state of charge. At this point, the predictive model agrees with the actual battery. As such, the largest deviations in predicted versus actual state of charge occur when modeling long periods (e.g. multi-sol plans) in which the battery is never fully charged.

The interview participants noted that there are inaccuracies in other predictive resource models, as well, including duration and data volume. The uncertainty in predicting how long an activity will take to execute prompts

the team to allocate margin for activities. Duration margin allocates idle time in the plan to allow planned activities to run longer than modeled without interfering with future portions of the plan. Additional duration is also reserved in the event of activities taking longer than the margin and needing to be forced to stop, referred to as clean-up. The downside of margin and clean-up is that it often results in unused time when activities run within predicted durations.

Two of the Tactical Uplink Leads brought up the issue of predicting data volume in a plan including predicting the amount of data an activity will generate and the amount of data that will be transmitted during a communication window. If the predictive models overestimate the amount of data an activity will produce or underestimate the amount of data that can be downlinked in a window, it could lead to the team restricting activity (e.g. acquiring fewer images or fewer frames in a panorama) or using high levels of lossy compression than necessary. One of the Tactical Uplink Leads indicated that the mission tends to do a good job of estimating how much critical data an activity will produce, but non-critical may not be as good. Typically, it is the critical data that is most important for prediction accuracy as this is the data that is used to make decisions for the next plan, so much be downlinked in time for the next shift.

Of the various resource predictions, the Tactical Uplink Leads were in agreement that the inaccuracies in power modeling resulted in the biggest negative impact to productivity.

Two of the participants suggested that one way of addressing resource prediction inaccuracies is for the rover to perform onboard resource management. As they stated, the rover would have up to date knowledge of how long activities are actually taking, including how long it took to heat devices, and actual available energy. This has the potential for enabling it to make use of extra time and energy.

Templates and Re-usable Sequences: Tactical Uplink Leads and Rover Planners discussed the value of having templates and re-usable sequences for the activities being planned. Templates and sequences encapsulate information about how to perform activities. Because the templates and re-usable sequences have been previously used, or previously vetted through strategic review, their use reduces the operations overhead associated with performing them, such as creating the activity plan fragment representing the activity and reviewing the contents of the sequence. The participants noted that there is a trade-off between making use of pre-defined templates and sequences. While it can reduce the time to implement activities, it may limit the flexibility of the team if they are limited to a fixed set.

Interpersonal Communication: Effective communication was identified as a significant productivity factor by many of the participants in each of the operations roles we interviewed. The participants discussed several types of communication. Communication between the science and engineering teams is important so that science can make clear their objectives for the plan and engineering can relate the constraints. Early and effective communication of this information can reduce churn during the day and avoid excess discussions. Communications between the science teams and Rover Planners was specifically called out by some of the participants, especially for plans that include contact science. In these situations there is often a complex dynamic between the types of contact science of interest to the science team and the constraints the Rover Planners have for accessing these targets with the rover's arm. Effective communication of science preferences and rover constraints is essential to identify targets that best meet the interests of science while satisfying the Rover Planner's constraints. Scientists and Rover Planners in the interview noted that the role of campaign scientist can be a significant assistance in this area. The important properties identified included someone who was a) local to JPL, to facilitate interactions with the Rover Planners, b) had a strong understanding of both the science interests and the Rover Planners constraints, to serve as an effective bridge between the two teams, and c) is staffed over a long duration to provide continuity.

It was also stated that effective communication from the supratactical process to the tactical process is important. The team members performing supratactical development are not always the same people performing the tactical implementation. Therefore, it is important that supratactical needs to effectively communicate the intent and priorities of the objectives so that the tactical team can ensure that the activity is being performed in a way that meets the intent of the request, or what activities should be descoped if necessary. It was noted that the mission does not currently have an effective way of relaying intent from supratactical to tactical but it was something being investigated.

The distributed nature of MSL operations was also noted as a challenge. While the engineering team and some science members are located at JPL, many of the scientists are located around the world. Interactions with the Rover Planners was called out as particularly challenging for team members working remotely. The current tools make it difficult for the science teams and Rover Planners to collaborate remotely on the selection of targets.

Science Team Engagement: Scientists and Tactical Uplink Leads emphasized the importance and challenge of engagement, especially for the science team members where it is especially important to be aware of the current and past context of the mission to make informed science decisions. A significant challenge to engagement is that many team members work on the mission part-time, with external obligations that require them to be away from the mission for a week or more. When they return it is difficult to keep up with collected data and results, and the rover may be in a new area.

Restricted Sols: We explicitly asked the participants to comment on the impact that restricted sols had on productivity. Almost all participants stated that restricted sols resulted in the biggest loss of productivity when they occur during periods in which the rover is driving. They noted two major areas in which productivity is lost. First, remote sensing is significantly impacted because the mission is not able to perform targeted observations following a drive, due to lack of imagery for ground to select targets and accumulated errors in the rovers position to support accurate targeting. Second, the amount of drives that can be performed is reduced as the mission cannot effectively perform multiple drives without ground-in-the-loop.

While the rover does have autonomous driving, aka autonav, capability, which could be used on restricted sols, there are a few reasons it is not done. First, due to wheel wear issues, the use of autonav has been restricted to limited types of terrain. Second, the current performance of autonav is not as good as manual drives, aka “blind” drives. The rate at which the rover drives with autonav is significantly slower than with blind drives. Thus, the use of autonav does not address the limitation that the mission cannot perform targeted remote sensing after a drive. If the project were to drive on the restricted sol of a plan, it would then mean that they would be unable to perform targeted remote sensing on the subsequent planning shift, which would make the negative impact to remote sensing even larger.

As a result of these limitations, some participants observed that the team is often unable to make use of all of the time allocated for science during restricted periods. The extra time can be used for performing maintenance activity, but not the high value science activities the scientists would prefer.

One of the Supratactical Uplink Leads noted that restricted periods while driving also reduces the amount of supratactical work that can be done to support tactical. For example, there is usually limited work that can be done on the supratactical timeline to develop a future drive. This is actually relevant for in many cases the rover is driving, not only during restricted periods.

The participants described a different type of productivity impact that restricted sols have on contact science operations. Unlike drives, the mission is better able to chain many types of contact science operations without ground-in-the-loop because the positional knowledge of the rover’s arm is maintained well across activities. One of the scientists stated that there can be a benefit to science when conducting contact science in restricted sols as the mission tends to stay in the same place longer, providing more time to collect data on the area. However, multiple participants noted that, because the mission uses multi-sol plans during restricted periods, plan complexity can limit how much activity is performed across the multi-sol plan. The total amount of activity across the multi-sol plan is often higher than what the team would achieve in a single-sol plan, but usually less than the total activity achieved across a comparable number of single-sol plans. For example, the amount of activity in a typical two sol plan may be higher than that in a typical single sol plan, but not within the sum of activity in two single sol plans. A Tactical Uplink Lead noted that the higher level of activity in the multi-sol planning shift poses a challenge for training as more experience is needed to manage the higher volume of activity within the tactical timeline.

The participants noted that restricted sols negatively impact contact science when ground-in-the-loop is required. This happens when the vehicle enters a new location and ground-in-the-loop is required to perform a stability assessment to ensure it is safe to unstow the arm and place it close to the surface. There are also various points during sample acquisition that require ground-in-the loop before moving forward.

One of the Supratactical Uplink Leads pointed out a side benefit of the overall reduced activity during restricted sols. It is often easier to fit in activities that require high levels of energy because the plans are naturally lighter

on overall activity. Scientists noted that it is often beneficial for environmental science as there is usually more time to fit in these activities, which do not typically require targeting.

While restricted sols are predominantly a factor of the difference between Mars and Earth times, as explained in Section 2.2.1, one of the Tactical Uplink Leads observed that problems with downlink, such as problems transmitting data from a relay orbiter to the ground, can cause a plan that was expected to be non-restricted to effectively become restricted due to lack of data from the rover.

Some of the participants discussed the potential for onboard autonomy to address the productivity challenges of restricted sols. Scientists pointed out that onboard autonomous science capabilities, such as the AEGIS system [Estlin et al.(2012), Francis et al.(2015)], has the potential for collecting targeted data without ground-in-the-loop. Some noted that there is also possibilities for performing onboard stability and safety assessments to reduce the productivity loss due to current ground-in-the-loop requirements for contact science and sampling. As one Tactical Uplink Lead stated it, the more you trust the rover to perform activity on its own, the less decisional data you require.

Time to Analyze Data, Make Decisions: A few of the scientists in the interview pointed to the time provided in the tactical timeline for analyzing data, making decisions and developing plan fragments for their activities as a significant factor in productivity. One of the scientists stated that this is particularly challenging when the mission is in drive mode. There are cases where there is not sufficient time to analyze and respond to the newly received data before the vehicle drives away. This can result in lost opportunities or the need to return to previous locations, which costs time.

Impact to Science while Driving: A Rover Planner and Scientist discussed impact to science when the mission is in a driving mode. There are two primary reasons noted. Due to previously discussed limitations for mid-drive and post-drive targeted remote sensing, the time to perform targeted remote sensing is restricted to a pre-drive science block. This duration is typically limited to allow time for the rover's drive. Therefore, the science team has only a small amount of time for activities that collect data on the rover's current location. Further, some of the data that would help to make decisions on what to target, such as higher-resolution images, often is not available for several days after it is acquired. By then the rover has moved to a new location. The scientist stated that because of these challenges, there are gaps in our understanding of Mars when driving.

Onboard Efficiency of Vehicle Activity: One of the Tactical Uplink Leads brought up the efficiency of how the rover carries out activities. In particular, activities that require the CPU to be powered on carry the overhead of powering the CPU. MSL recently enabled a capability that allows certain devices to be heated while the CPU is powered off. This has reduced the cost of performing certain activities. The Tactical Uplink Lead noted that if more instruments could be operated without requiring the CPU to be on, it would have similar improvements in resource usage.

3.4 Thoughts on Improving Productivity

We asked the participants their thoughts on how productivity could be improved. They provided a wide range of ideas.

- Improved resource modeling, e.g. more accurate models, to enable better prediction of available vehicle resources.
- Better integration between Rover Planner tool and activity plan development tool to avoid effort required to re-create Rover Planner activity within the activity plan.
- Simplified tools to make it easier to access science data without having to be an expert in each type of data (e.g. should not have to be an expert in APXS data to access APXS results).
- More intuitive methods for developing science plan fragments.
- Tools to assist developing Rover Planner activities (identified as important because Rover Planner activities tend to require the most time to develop). A specific suggestion was made for contact science in which the Rover Planner would specify the objectives and the tools would work out the details to optimize the sequences for executing the activities.

- Support for documenting intent, to facilitate communication between supratactical and tactical operations.
- Allow the vehicle to insert activities onboard if extra resources are available.
- Reduce the duration of the tactical timeline to reduce the number of restricted sols.
- Develop a capability to enable a rover trajectory following mode for drives (i.e. rover follows path defined by Rover Planner, compensating for slip as necessary to stay on path).
- Create a tool to help scientists find and explore collected data.
- Increase orbiter relay bandwidth and reduce relay latencies (the time between the overflight and when the data is available to the operations team on the ground).
- Staff the mission with a co-located, fully engaged team that is on the mission full time.
- Have an extra, mid-day, ground-in-the-loop cycle, would allow an extra command cycle on drive sols. It was acknowledged by the person suggesting this that it was not necessarily a feasible suggestion, but it stresses the importance of ground-in-the-loop cycles for science productivity during drives.

4 Pahrump Hills Walkabout Pass 1 Campaign

Curiosity arrived at the Pahrump Hills outcrop in September of 2014. Figure 11 shows the route the rover took to reach Pahrump Hills from the landing site [Jet Propulsion Laboratory Press Release(2014a)]. Pahrump Hills was a significant destination for Curiosity because HiRISE imagery indicates that the outcrop represented the rover's first encounter with the Murray formation, the strata recognized as lower Mount Sharp [Stack et al.(2015)].

To help provide context for the Pahrump Hills Walkabout campaign, the HiRISE map in Figure 12 highlights the contact between the crater floor, which Curiosity had been exploring since landing, and the Murray formation, forming the basal layer of Mount Sharp [Jet Propulsion Laboratory Press Release(2014c)]. As seen in the map, Pahrump Hills lies at the boundary between these regions. Figure 13 provides another view of the layers of Mount Sharp including a cross-sectional view illustrating the transition from the crater floor, to Murray formation and the Hematite ridge further up Mount Sharp [Jet Propulsion Laboratory Press Release(2014b)].

Figure 14 shows a HiRISE image capturing Curiosity at Pahrump Hills [Jet Propulsion Laboratory Press Release(2015b)]. The bright features in the landscape are sedimentary rock and the dark areas are sand. The image was acquired on December 13, 2014, Sol 836, while Curiosity performing the second Pahrump Hills Walkabout. At this time, Curiosity was located near a target called Whale Rock, one of the science targets for the Pahrump Hills campaign.

4.1 Overview of Campaign

The predominant practice for Mars surface exploration, for both MER and MSL, has been to employ a linear approach where the rover examines sites in the order in which they are encountered and rarely backtracks. This is in contrast to the typical approach used by geologists on Earth in which geologists begin by walking the field site, performing a "walkabout" that provides context for the exploration and helps prioritize targets for a more detailed follow-up [Yingst et al.(2015)].

When the team arrived at Pahrump Hills, they decided to employ the walkabout approach to study the formation. They organized the exploration of Pahrump Hills into three passes. The first walkabout would employ the rover's remote sensing instruments, Mastcam, ChemCam and MARDI, to quickly characterize the the formation. Based on the data collected from the first pass, a second pass would be made to perform more detailed study on high priority targets using the DRT, MAHLI and APXS. Finally, based on the data from passes 1 and 2, a site would be selected to collect samples using the drill.

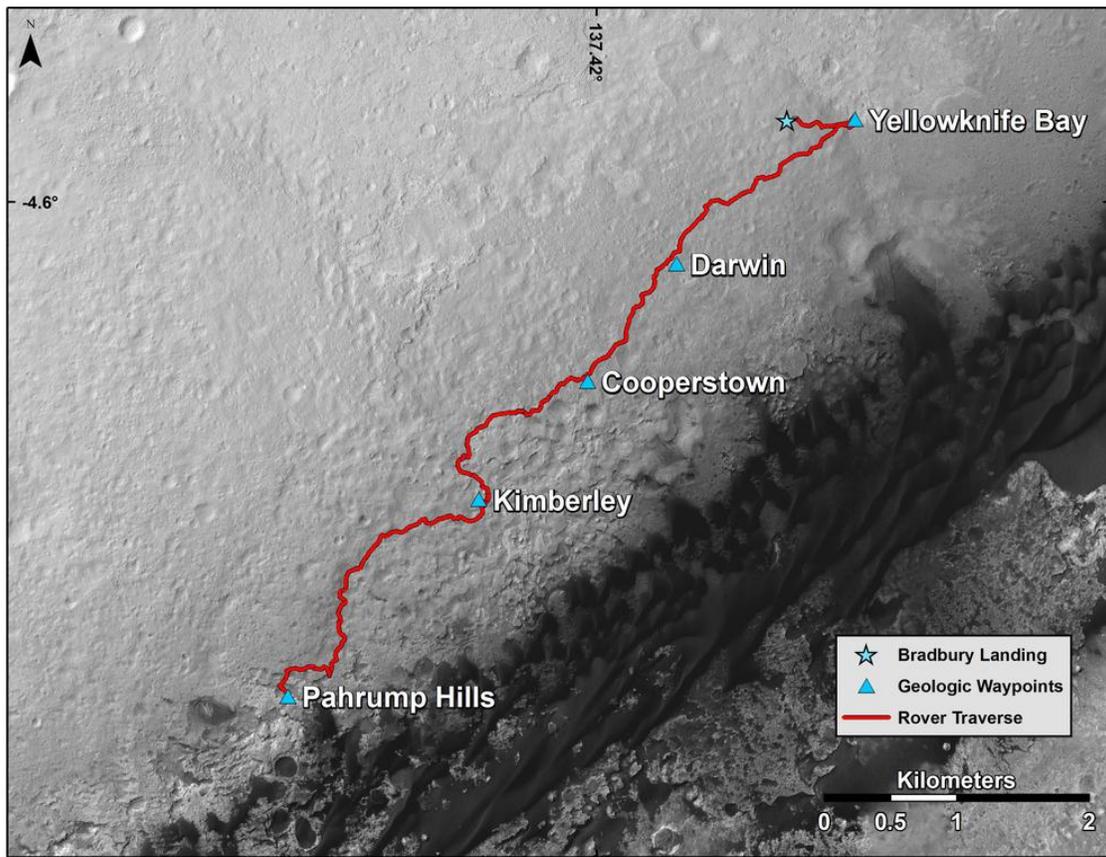


Figure 11: Curiosity's route from landing to Pahrump Hills (NASA/JPL-Caltech/Univ. of Arizona).

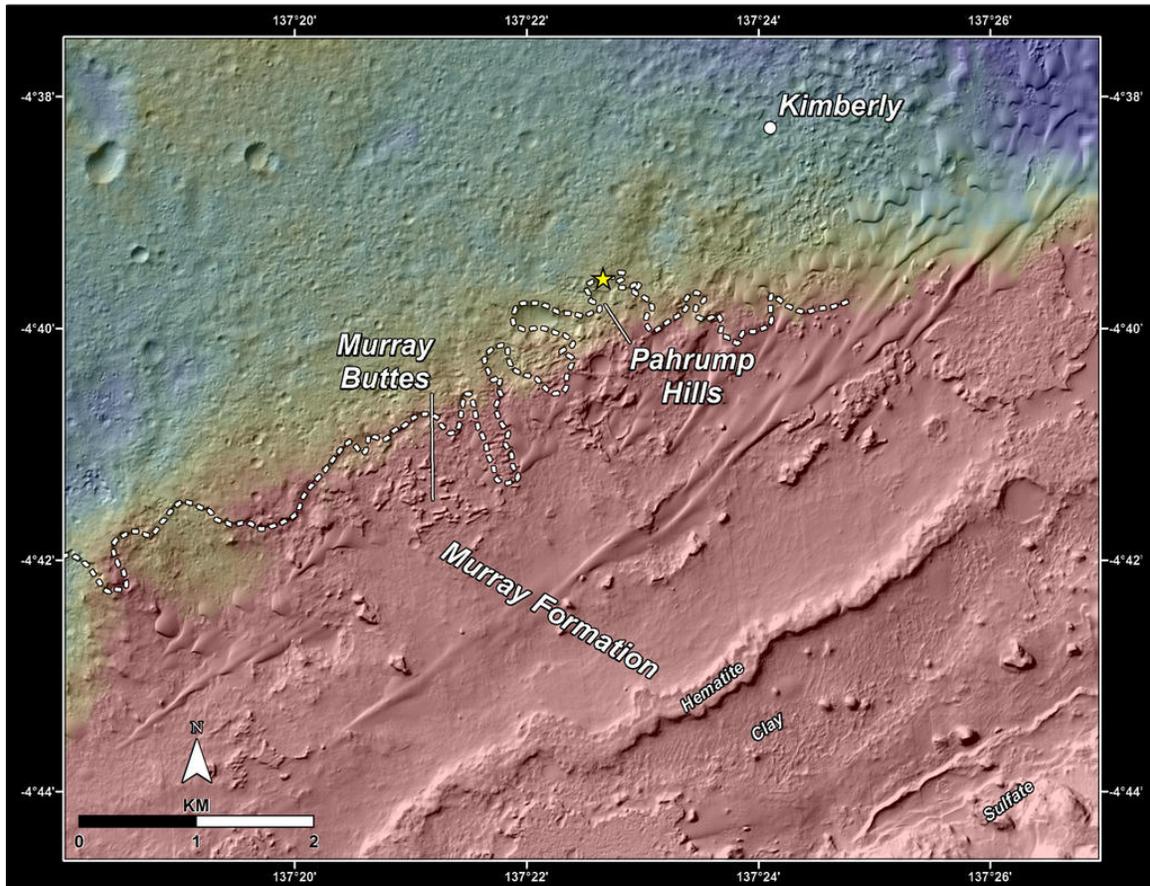
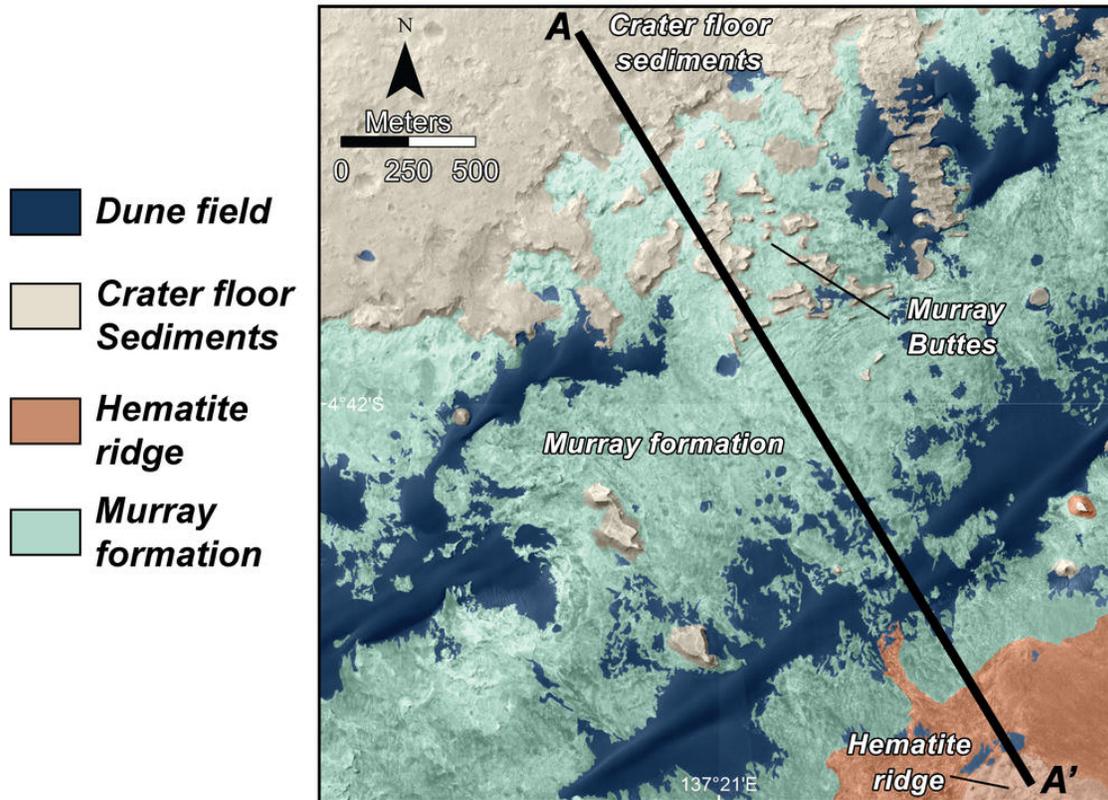
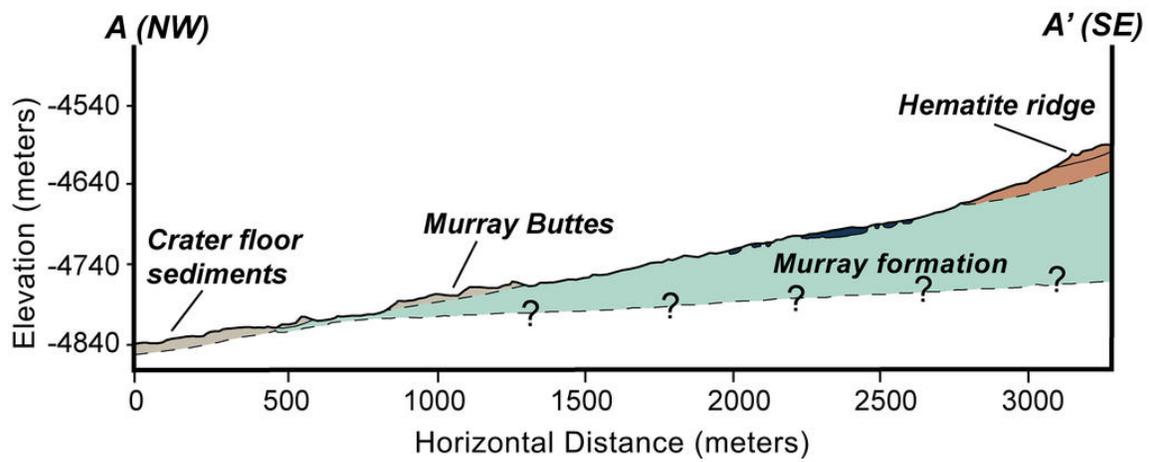


Figure 12: Contact (dashed line) between Gale Crater plains (Aeolis Palus) and Mount Sharp (Aeolis Mons) showing Pahrump Hills at the start of Murray formation. (NASA/JPL-Caltech/Univ. of Arizona).



(a) Map of lower Mount Sharp on Mars, showing the major geologic units identified from orbit



(b) Cross-section view of the geologic sections of Mount Sharp

Figure 13: Geologic Layers of Mount Sharp (NASA/JPL-Caltech/Univ. of Arizona).



Figure 14: Curiosity's at Pahrump Hills (NASA/JPL-Caltech/Univ. of Arizona).

4.1.1 Science Objectives

The specific objective of this initial pass of the Pahrump Hills Walkabout was to perform a reconnaissance imaging survey using the rover's remote sensing suite of instruments to identify and characterize the lithologies present in the formation [Minitti et al.(2015)]. The geography of Pahrump Hills was well suited for developing a strategic plan for the initial walkabout. As shown in Mastcam mosaic in Figure 15, the formation was laid out in front of the rover with an upward slope affording an excellent view of the exposed outcrop and potential paths for the rover to traverse.

The scientists identified specific outcrops of interest in the Mastcam mosaic, white boxes in Figure 15, that they wished to visit for targeted remote sensing. The scientists worked with the Rover Planners to select traverse paths to these areas of interest. The strategically developed traverse plan included defining the expected number of drives that would be needed along with identifying point along drives, mid-drive stops, to collect Mastcam imagery of the surrounding terrain. At the end-of-drive locations, the scientists would select targets for ChemCam and Mastcam observations. Throughout all of the drives, the MARDI instrument would be used to collect a "sidewalk" video of the terrain that was traversed.

Figure 15 represents a high-level, strategic plan of the activities. The specific route the team would take, and the number of drives required to achieve the goals could deviate from this plan as additional data is acquired and detailed, tactical, plans are formed. As an example, at this stage in the planning, it was unclear if the team would be able to drive to the highest white box, above Whale Rock, in Figure 15 given the slope of the terrain and nature of the rocks in the area. Additional analysis would be require to make this assessment. In addition, it is important to note that end-of-drive locations are approximate locations based on the imagery available at the start of the campaign. The science team refined these locations throughout the campaign as additional imagery was collected in order to pick appropriate end-of-drive locations to meet science objectives.

4.1.2 Summary of Activity

Following is a summary of the team's and rover's activities during each tactical planning cycle of the campaign.

Sol 780: Sol 780 marked the beginning of the Pahrump Hills Walkabout campaign. The decision to begin the walkabout on this sol came relatively late. The rover experienced an anomaly with the arm during execution of the Sol 774 plan and the team needed additional time to investigate the problem. A significant implication of the

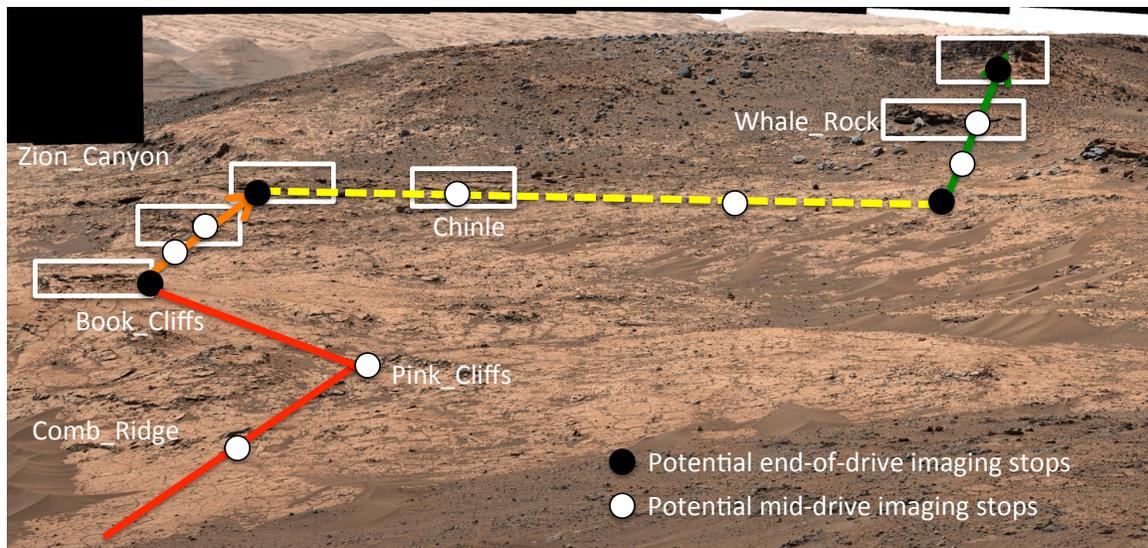


Figure 15: Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of campaign.

ongoing arm investigation is that the use of the MAHLI instrument was restricted to only being used in the arm stowed position. This was to address the concern that there was an increased risk of a recurrence of the arm fault could result in leaving the MAHLI dust cover open for an extended duration. Given the restrictions in MAHLI use, the team decided to begin the Pahrump Hills Walkabout as the first pass was focused on remote sensing instruments, and thus did not require use of MAHLI outside of the arm stowed position.

The major science objectives of the sol were to drive to the Book Cliffs target with two mid-drive stops to image Comb Ridge and Pink Cliffs. See Figure 15 for reference.

One of the features that made this drive, as well as all the drives during the first Pahrump Hills Walkabout, was the acquisition of MARDI video, aka sidewalk video, throughout the drive in order to help characterize the sedimentology of the outcrop.

The MARDI sidewalk video provides an example of how managing the details of an activity can add complexity to operations. The MARDI instrument flight software is designed to detect difference in frames of the video and only save frames when the detected differences exceed a threshold. This capability provides a dramatic reduction in the data volume that would otherwise be required for the video by filtering out redundant frames. However, it adds uncertainty into what the actual data volume will be for the activity since the number of images that will be saved depends on the duration of the drive as well as the terrain being traversed. The instrument operators developed a spreadsheet to make a refined estimate on data volume based on the tactical plan. These estimations were then incorporated into the MSLICE plan.

The sidewalk activity created additional complications related to when the video is stopped. The team used the standard MARDI cleanup sequence that is invoked as part of the Master cleanup at the time all drive and post-drive activities are expected to be completed. A consequence of this is that the MARDI video would continue to run during post-drive imaging and any additional margin the team allocated to the drive and post-drive activities. Although the rover would not be traversing terrain during the post-drive imaging and margin, lighting conditions would be changing and this could cause the MARDI change detection to save additional frames. Thus, the time allocated for post-drive imaging and margin needed to be factored into the estimated data volume for the sidewalk videos.

Finally, there can be conflicts with performing certain CPU intensive activities in parallel with stereo Mastcam activity which can result in loss of the Mastcam images. As a result, it was important that stereo Mastcam not be acquired during the time in which the MARDI video was being collected.

The drive toward Book Cliffs was complicated by the rough terrain in the area and limited stereo data due to occlusions in the viewshed. The platy bedrock contains exposed, jagged ridges which can be damaging to the

rover's wheels. As such, the Rover Planners and Surface Properties Scientist have to work together to try to avoid the excess wheel damage while still achieving the drive objectives.

This particular drive was not able to attain the intended end drive objective of reaching Book Cliffs due to insufficient navigation imagery near the target. A 22m drive was planned that would stop about 7.5m from Book Cliffs. The post-drive imagery from this position should provide the data needed to complete the traverse to Book Cliffs on the next drive.

The plan also included some pre-drive Mastcam and ChemCam on targets of interest to collect information on conglomerate morphology and layering.

The Sol 780 plan also included observations for another significant science campaign that happened to coincide with the Pahrump Hills Walkabout. Comet Siding Spring (Comet C/2013 A1) was approaching Mars and would be passing closer to Mars than any known comet flyby of Earth or Mars. There was significant interest in observing this approach. The comet campaign was a major endeavor involving observations from the Mars rovers Curiosity and Opportunity, the Mars orbiters, MRO, Odyssey, MEX and MAVEN as well as Earth-based telescopes. For MSL, the observations for the comet approach spanned several planning sols and required careful timing to appropriately observe the comet.

The comet's closest approach was on Curiosity's Sol 783, but Sol 780 included important observations that would be used to refine the observations being planned for the closest approach. The observations for Sol 780 plan were in the early hours of Sol 781. Night-time observations such as these can be challenging for operations due to the significant heating needed to use instruments and devices during the colder night hours. Fortunately, this campaign was occurring during Martian Spring for the southern hemisphere and energy was not a significant constraint at this time.

However, managing data for this activity was a concern as the team needed the results from this observation in time for planning the observations on Sol 783. This required the team to carefully consider which specific images should be acquired with a critical priority to support subsequent planning, but avoid marking too many critical as that could interfere with other decision-critical imagery acquired in the Sol 781 plan.

Sol 781: The Sol 780 drive completed as expected, leaving the rover near Book Cliffs. The team took advantage of the location with a long science block to acquire Mastcam and ChemCam observations of targets in the area.

A significant portion of the Sol 781 plan consisted of dumping a sample obtained from Confidence Hills back on Sol 759. The rover has the ability to carry the sample with it as it drives, referred to as a "cached sample", which enables the team to periodically deliver additional samples to the SAM and CheMin instruments. When the team decides it no longer needs the sample, it is dumped on the surface. The team takes advantage of the dump activity to acquire additional observations of the sample, including high-resolution imagery and spectroscopy, to further study the sample.

In this instance, the team was restricted in the high-resolution imagery that could be obtained due to the continuing restriction on the use of MAHLI while the arm anomaly from Sol 774 was still being investigated. As such, the team limited the imagery observations to high-resolution Mastcams. Multiple observations of the dump pile were made with APXS including an overnight integration.

The planning of the sample dump illustrated the significance of communication among the team, as identified as a significant productivity factor by the interview participants. On this sol, there were complications in communicating the location of the sample dump from the Rover Planners, who would be selecting the location, and the science team members who would be targeted the location for their observations. The teams make use of a target database to store the names and locations of significant targets. In this case, there were distinct targets which similar names to the dump location which lead to some confusion in which target represented the actual dump location. The problem was identified and resolved during the Master/Submaster meeting.

The plan included targeted observations with ChemCam and Mastcam. Targets selected for ChemCam included background outcrop with a small vein observed in the Navcam, a target with a light-toned "halo" feature, and a target with a dark-toned, raised feature with variable texture. The Mastcam target exhibited curvilinear features in the bedrock that intersect with fractures.

Observations of Comet Siding Spring continued in the Sol 781 plan with additional Mastcam imagery in the early morning hours of Sol 782. Because Sol 783, closest approach of the comet, would be planned as part of

the Sol 782 plan (part of a 3-sol weekend plan), the results from this observation would not be available in time to support the planning of the closest approach.

The tactical timeline was interrupted for about 45 minutes to allow the JPL operations team to participate in The Great Shakeout earthquake drill.

The team ended up descope some potential Mastcam and ChemCam observations due to overall complexity given the level of activity already in the plan.

The timing of a Navcam 360 observation was placed earlier in the sol, rather than the preferred time later in the afternoon. The later time was taken by the arm activity. The added complexity of inserting the unrelated Navcam observation at a particular time within the arm activity was not worth the added benefit as the team could accept the earlier timing of the Navcam 360 observation.

During planning for this sol, the team investigated concerns regarding the High-Gain Antenna (HGA) reaching a soft stop during the Sol 781 uplink window. The team needed to determine if they would still be able to uplink the command products for Sol 781 and what the consequence of the HGA reaching the soft stop would be. After consultation with various experts, the team was able to conclude that there would be sufficient time to uplink the command products and that the window would not have negative consequences on the vehicle (e.g. would not cause a safing event).

Sols 782, 783, 784: Sols 782, 783 and 784 were part of a 3-sol weekend plan developed on a Friday. Early in the mission, the team would work Saturdays and Sundays, to increase the number of ground-in-the-loop cycles with the rover. However, this pace is not sustainable in the long run, so by this time in the mission, the team created multi-sol plans on Fridays giving the rover activity to perform across the weekend without requiring the team to perform tactical operations on Saturday and Sunday.

The major objectives for the weekend plan were to continue the drive toward Book Cliffs and to make observations of Comet Siding Springs which would be making its closest approach to Mars on Sol 783. Given the rare occurrence of the comet approach, it was declared by the science team as the highest priority activity of the multi-sol plan.

The comet observations in this plan would include ChemCam observations, in addition to Mastcam and Navcam. Use of the ChemCam instrument adds additional complexities due to the extra precautions required to ensure the instrument stays safe from damage that could occur if the sun enters its optics. Even though the comet observations would occur at night, when of course the sun is not visible, the precautions are still in place to protect against the possibility that the mast could incur a motor fault leaving the ChemCam instrument at a location in which the sun could eventually pass through its optics. As such, the combined ChemCam and Mastcam observations of the comet needed to be carefully designed to ensure appropriate timing to observe the comet and mast motions to ensure the safety of ChemCam.

As the tactical day progressed, ChemCam instrument operators observed that the expected heading and tilt of the drive planned for Sol 782 could impact the sun safety of the subsequent observations of the comet's closest approach. Updating the planned observations to be robust to the end of drive attitude would have required additional time within the observation which would disrupt the careful timing of the activity.

The team made the decision to descope the drive in order to preserve the higher priority observation. The team also descope the follow-up comet observations on Sol 784 due to similar concerns about ensuring sun safety.

Given the late decision, there was insufficient time in the tactical timeline to replace the descope activities with alternative observations. As such, the weekend plan had, overall, relatively low activity.

The plan included targeted observations with ChemCam and Mastcam. The team planned targeted ChemCam and Mastcam on the Confidence Hills dump pile to further assess the sample. This was a valuable observation as it allowed the team to compare analysis on this post-sieved dump pile vs. a previous pre-sieved dump pile and allowed ChemCam observation on the same sample that CheMin had processed.

In addition to the dump pile observations, the team considered additional ChemCam targets. These includes a target with a rosette-like aggregate that appears similar to a previously studied outcrop, a dark layered rock, and a target that appeared to be planar fracture-fill that had been split-open. Due to complexity involved with targeting the dump pile, the team decided to descope these additional ChemCam targets and instead retain the higher priority dump-pile observations. Instead, a targeted Mastcam was acquired of a target with fine laminate.

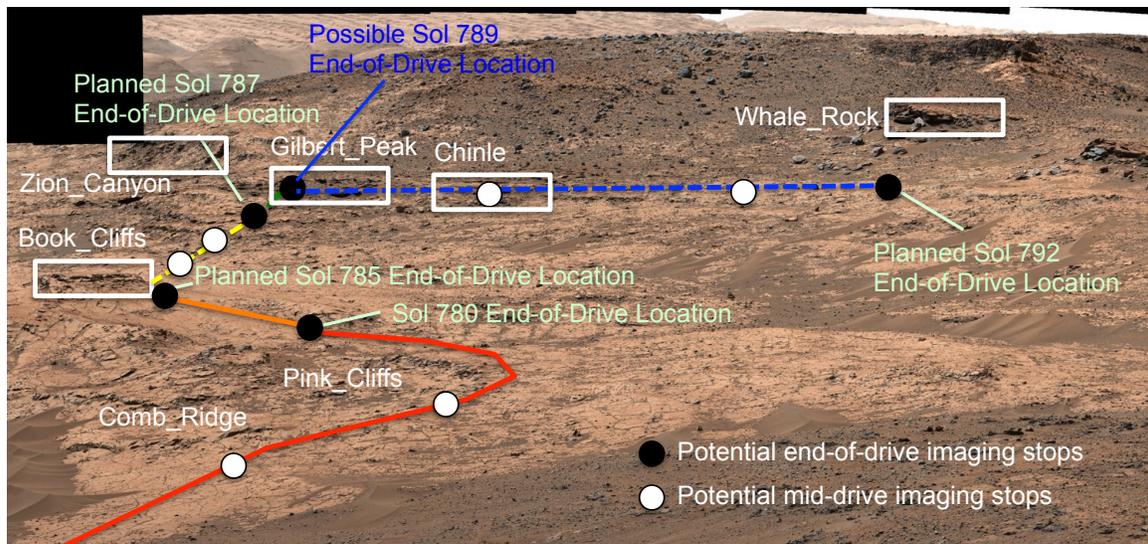


Figure 16: Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of Sol 785 planning.

As mentioned above, the ChemCam targets were limited due to issues with targeting the dump pile. The issue was that the target did not have associated location information. This was due to the dump pile being smooth, with insufficient texture for stereo computation of the target coordinates. Instead, the ChemCam team used a nearby location with sufficient stereo and performed a raster across the pile from this location.

2-sol plan Restricted" Full 9486 6 6 No

Sols 785, 786: After successful comet imaging over the weekend, the team returned Monday to plan two sols, Sols 785 and 786. This was a 2-sol plan due to the relative timing of the Mars and Earth days as described in Section 1.3. The team had moved back into restricted sols and, as such, would be planning a series of multi-sol plans.

The plan included a pre-drive science block consisting of ChemCam observations of three targets and supporting Mastcam documentation images. The team took advantage of the fact that the drive was descope in the weekend plan, leaving the rover in the same location. This gave the team the opportunity to consider the ChemCam targeted observations that had been descope in the weekend plan. The team opted to include the dark-layered rock target and target with the rosette-like aggregate which had been previously descope. They also added a ChemCam observation on a possible Calcium vein that is resistant. Reports indicated that the science team was interested in two additional targets for ChemCam, but limited themselves to three observations. The reports do not indicate the reason for this limitation. The additional targets included the apparent planar fracture-fill descope in the weekend plan and fracture filling ridge.

The team planned a drive for Sol 785 that would continue the drive toward Book Cliffs that began on Sol 780. Figure 16 shows an updated path from the one depicted in Figure 15 taking into account the actual end of drive location from the drive on Sol 780 and the planned end of drive location for the drive on Sol 785.

The drive for Sol 785 was a short drive, aka "bump", of about 6m to get into a better position for subsequent observations of Book Cliffs. As with the Sol 780 drive, the path was selected to avoid jagged rocks that could contribute to wheel wear.

The team was limited in the activity that could be performed on Sol 786 due to being in restricted sols. Since the rover drove on Sol 785, the team would not know its location nor detailed information about the terrain in the rover's vicinity on Sol 786. This means the team could not plan targeted observations on Sol 786. The only activity planned on Sol 786 directly related to the Pahrump Hills Walkabout campaign was a ChemCam "blind" targeting. This activity points ChemCam in a fixed location relative to the rover that has been demonstrated strategically to be safe for the instrument and to avoid hitting the rover with the laser. However, given that it is taken at a fixed pointing and without the team knowing the actual location in which the rover will be at the

time of the observation, it is not guaranteed that the observation will be of outcrop or soil. The Sol 786 blind ChemCam observation happened to hit sedimentary rock.

A large amount of Sol 786 was allocated towards an routine engineering maintenance activity to collect thermal data.

Sols 787, 788: The drive from Sol 785 brought the rover to the desired location at Book Cliffs. The plan for Sol 787 was to perform pre-drive ChemCam and Mastcam observations of Book Cliffs and then drive toward Gilbert Peak (see [Figure 16](#) for the planned end of drive location on Sol 787).

The pre-drive science block included targeted ChemCam and Mastcam observations. The ChemCam observations were of a laminated vertical face and a parallel-sided fracture, to document variations in texture and chemistry. The Mastcam observations include Mastcam documentation of the ChemCam targets, a mosaic to document stratigraphy of Book Cliffs and a separate mosaic to document stratigraphy and geology of upcoming outcrops in future drives.

The drive for Sol 787 was through terrain that was more benign than prior drives but there were still some loose rocks in the area that were avoided to reduce wheel wear. This would be a drive of about 18m. However, the expected end of drive location would not be the desired location from [Figure 16](#) due to concerns about wheel wear. The planned drive would include two stops for mid-drive Mastcam imaging as well as MARDI sidewalk video throughout the drive.

Similar to Sol 786, the plan for Sol 788 was limited due to lack of information about where the vehicle would be at the end of the Sol 787 drive. As such, the only campaign-specific observation on Sol 788 was another ChemCam blind target observation. As in Sol 786, the blind target ChemCam hit rock again.

The plan for Sol 788 did, however, include atmospheric observations which do not require targeting. The observations included monitoring of dust in the atmosphere.

Sols 789, 790, 791: This was another three-sol weekend plan with the objective to continue with the Pahrump Hills Walkabout. Because the team was in restricted Sols, the drive would be on the second sol, Sol 790. This gave the team all of Sol 789 for targeted observations and ensured the post-drive data from Sol 790 would be available for Monday's planning of Sols 792 and 793.

[Figure 17](#) shows the updated traverse plan for the start of Sol 789 planning. The Sol 787 end drive location was updated from [Figure 16](#) reflecting that the rover was not close enough to enable ChemCam observations of targets near Gilbert Peak. The ChemCam operation range is about 2-5m and the rover ended up about 18m away from Gilbert Peak.

The Sol 787 drive was intended to contain two mid-drive imaging stops. However, due to how the mid-drive imaging was sequenced and interaction with the terrain, the second mid-drive imaging observation did not take place. The mid-drive imaging was sequenced using a feature of flight software called Observation Tables. Observation Tables allow the team to sequence a drive and specify conditions in which the drive should pause and run a sequence, in this case to a sequence that would perform mid-drive imagery. Various criteria can be used to select when an Observation Table sequence is invoked. In this case, the Rover Planners specified the sequence should be acquired every 5.9m for at most two invocations. This would result in two mid-drive imaging stops with 5.9m of separation. The portion of the drive in which the observation table was active was about 12 meters, which in ideal conditions would be just enough distance for the two mid-drive imaging stops. Due to slippage, the rover did not traverse sufficient distance to trigger the second invocation.

Prior to the drive, the team planned targeted observations to characterize the outcrop at the rover's current location. ChemCam targets were selected to characterize a laminated vertical face, document variations in texture and chemistry of a well-laminated outcrop, and characterize a bright vein. Targeted Mastcam observations provided documentation for the ChemCam observations and to collect extended color coverage of the nearby outcrop to characterize the laminae.

The objective for the Sol 790 drive was to continue the traverse to get within ChemCam range of two ridges that are in front of Gilbert Peak (These ridges would later be named Alexander Hills). The challenge was to try to get close enough to the targets to be within the operation range of ChemCam but to avoid excessive wheel damage from sharp rocks in the vicinity of the target. The planned drive was expected to get the rover to within 4.2m of the intended ChemCam targets.

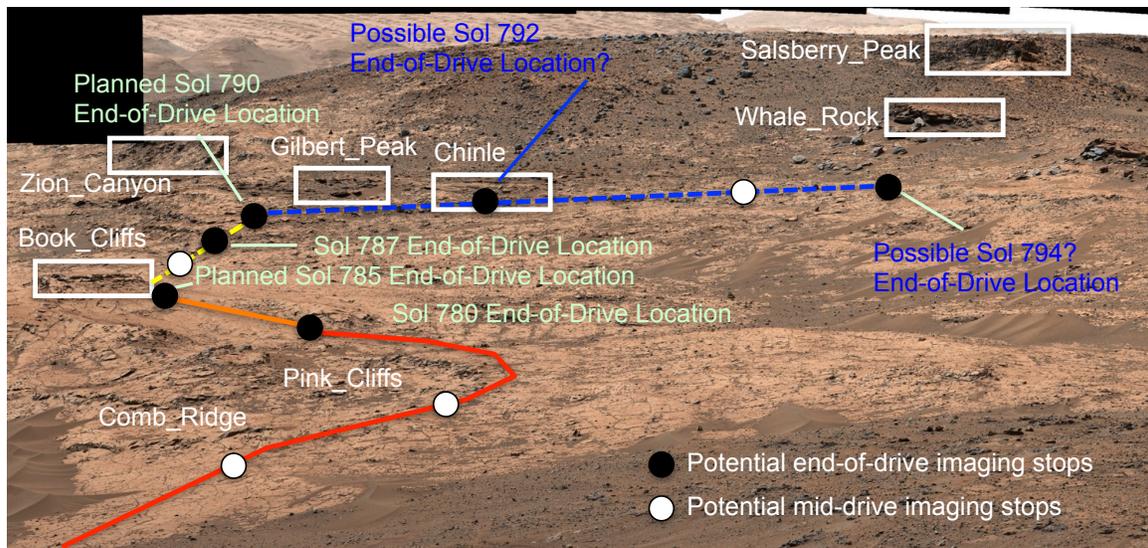


Figure 17: Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of Sol 789 planning.

Sol 791 included another ChemCam blind target and environmental observations. The resulting raster hit a mix of soil and grains.

Sols 792, 793: Figure 18 shows the updated traverse plan for the start of Sol 792 planning. The rover reached its intended end of drive location from Sol 790's drive placing the rover within range of Alexander Hills, the ridges near Gilbert Peak that the Science Team wanted to observe with ChemCam.

The targeted pre-drive science block was used to collect ChemCam observations on the ridges near the rover's current location. The targets were selected to characterize the ridges and a nearby layered rock with a white vein. Targeted Mastcam observations provided documentation for the ChemCam observations and documentation of variability of layering in the foreground. A Mastcam mosaic to the right of the rover was acquired to continue the documentation of the traverse.

The objective of the drive for Sol 792 was to get the rover within 3 meters of Chinle to enable ChemCam observations of the area. This would be a drive of about 16 meters. The terrain was challenging from a wheel wear perspective and some rocks could not be avoided. The second mid-drive imaging stop was planned with the rover straddling a 20cm rock.

As with the other drives in this initial pass of Pahrump Hills, MARDI sidewalk video would be acquired to document the traverse. In addition, this drive was planned to have two mid-drive stops for Mastcam imaging.

The team needed to keep an eye on data volume during planning. They wanted to ensure that they would get sufficient data down from this first walkabout of Pahrump Hills in time to inform decisions about where to go during pass 2 of the Pahrump Hills investigation. As a result, a few observations were descoped from the plan. These included the ChemCam blind targeted observation and its accompanying Mastcam documentation observation that would have been performed on Sol 793 and slip assessment Hazcam images, which were not needed as the team would not be performing contact science.

Sols 794, 795: Figure 19 shows the updated traverse plan for the start of Sol 794 planning. The rover arrived at its intended location near Chinle. The science team was excited by the layered outcrop at Chinle. The team was still closely watching data volume to ensure sufficient data from the first pass of Pahrump Hills could come down in time to inform the second pass.

Pre-drive targeted ChemCam observations were planned on an outcrop with a fracture, and targets to survey layering of the nearby outcrop. Targeted Mastcam observations were planned to document the ChemCam observations and to document fine laminations and cross bedding.

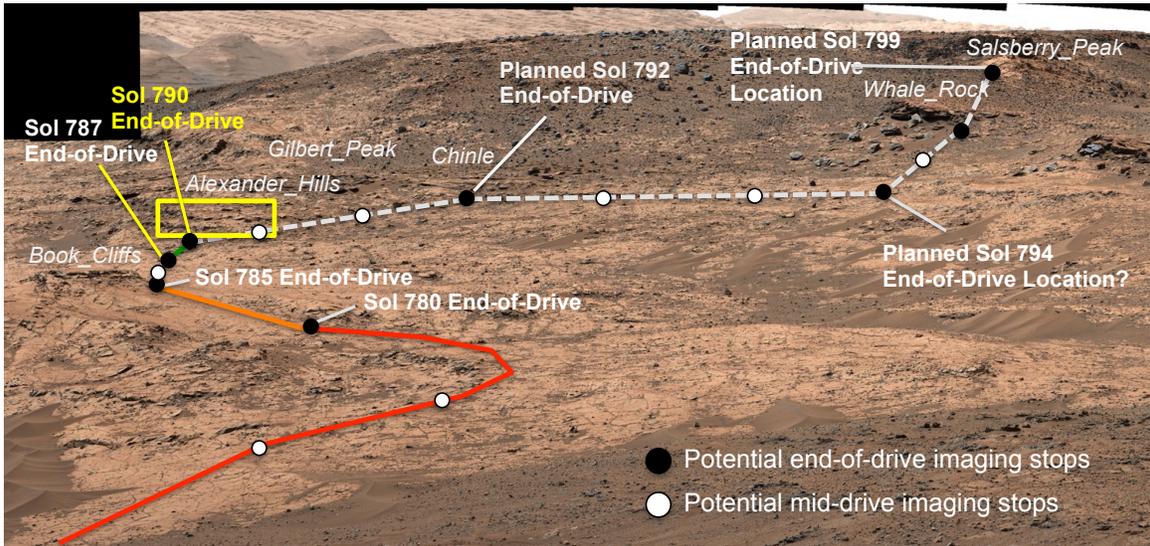


Figure 18: Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of Sol 792 planning.

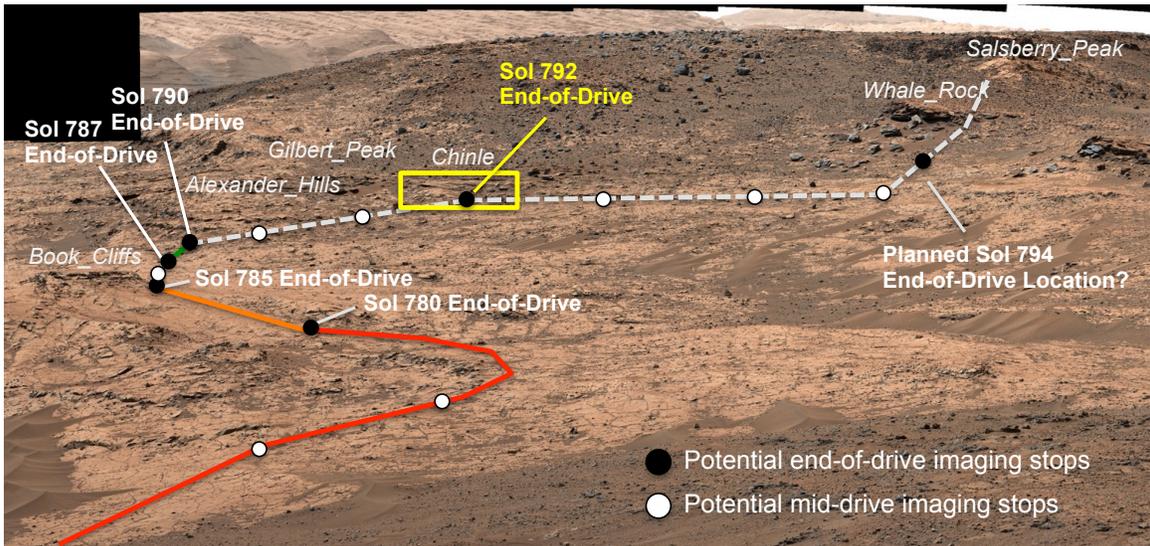


Figure 19: Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of Sol 794 planning.

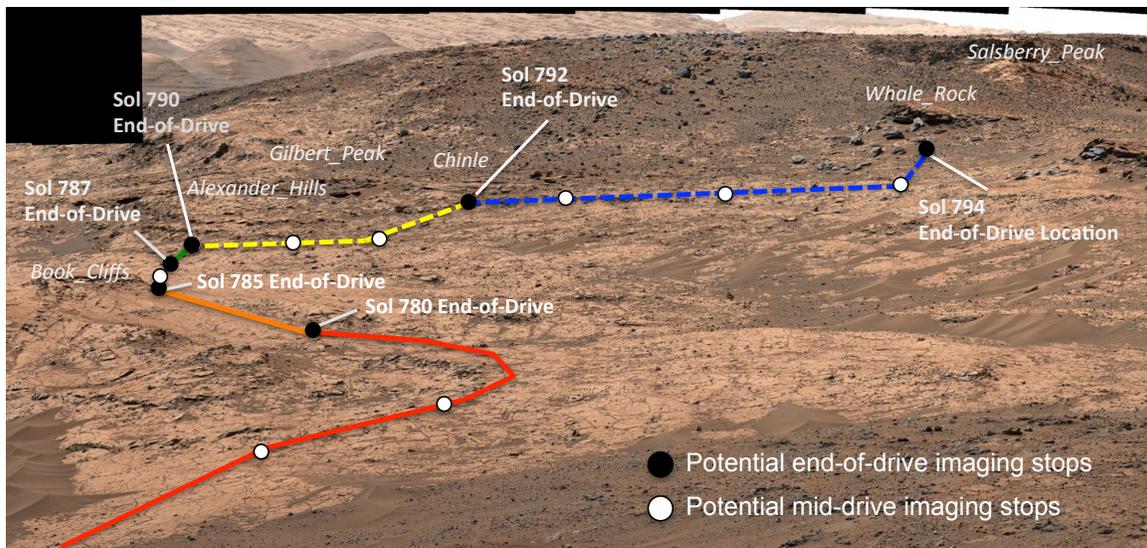


Figure 20: Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of Sol 796 planning.

The objective for the Sol 794 drive was to get the vehicle within range for ChemCam to target Whale Rock and to target Mastcam on Whale Rock and Salsberry Peak. The terrain was challenging as there were large, loose rocks over rough terrain. The 30 meter drive avoided hazards to the rover, but would result in some wheel wear damage. The drive included three mid-drive stops for Mastcam imaging and a MARDI sidewalk video throughout.

Activities for Sol 795 included a ChemCam blind targeted observation and environmental observations. The ChemCam blind targeted observation ended up hitting a mixture of pebbles and soil.

As with the previous plan, the team decided to remove activities from the plan to reduce the volume of collected data, again with the main intent of trying to ensure the data from the first pass of the walkabout that is needed to inform the second pass would come down in time. In this case, the team removed a ChemMin activity and slip assessment imagery from the plan. The team also limited the ChemCam blind target activity to a 5x1 raster rather than a 10x1 and desopped a ChemCam calibration activity, again to reduce data volume.

Sols 796, 797, 798: The final stage of the first Pahrump Hills Walkabout was a three-sol weekend plan, which happened to be planned on Halloween. Figure 20 shows the updated route.

As shown in in Figure 20, the 30 meter drive from Sol 794 executed well, leaving the rover at the expected end of drive location.

The first sol of the plan included several targeted observations with ChemCam and Mastcam. ChemCam was used on Whale Rock to acquire a vertical raster through exposed layers. ChemCam observations were also acquired on two different float rocks with layered material, and background bedrock with veins. Mastcam observations provided documentation of the ChemCam targets along with mosaics of Whale Rock, Salsberry Peak and another target named Newspaper Rock.

A long, 60 meter, drive was planned on Sol 797 to take the rover back near the Confidence Hills drill location in preparation for a scuff test and to prepare for pass two of the Pahrump Hills Walkabout.

Although the weekend plan covers 3 sols, the team chose to restrict the third sol of the plan, Sol 798 to be a "runout" sol. This allows the team to generate sequences to cover the sol with minimal effort but allows for only significantly limited science. In this case, only background REMS and RAD activity would be performed on the runout sol.

4.1.3 Outcome of Campaign

The Pahrump Hills Walkabout Pass 1 spanned 19 sols and included 7 separate drives covering 152 meters, including the final drive back toward Confidence Hills. Figure 21 shows the actual drive path taken during the campaign. As shown in Figure 15, the initial, strategic plan anticipated 5 drives to complete the campaign. Limited viewshed caused the initial drive to Book Cliffs to be broken into two separate drives. Challenging terrain near Alexander Hills resulted in an extra drive to get the vehicle close to desired science targets while reducing wear to the rover's wheels.

The Pahrump Hills walkabout was a highly successful science campaign. ChemCam data collected during the initial walkabout contributed to a chemostratigraphic characterization of the formation [Forni et al.(2015)]. A continuous record of the outcrop was obtained by creating a mosaic from the MARDI sidewalk video that was acquired through each of the drives [Minitti et al.(2015)] which enabled tracking the distributions and relationships of sedimentary structures observed in the formation. The collected reconnaissance survey from all of the remote sensing instruments informed follow-up observations with MAHLI and APXS in the second pass [McBride et al.(2015), Thompson et al.(2015)].

The data collected from Mastcam, MARDI, MAHLI from the first two passes was localized with HiRISE maps and Navcam stereo to produce a stratigraphic column of Pahrump Hills, identifying five main sedimentary facies [Stack et al.(2015)]. The team also observed variations in the occurrence of diagenetic features (e.g. concretions, veins, crystal laths) across the column which provided valuable information regarding the deposition and alteration of the formation.

After the completion of Passes 1 and 2, the team selected Pink Cliffs for the location of the next drill sample. The team attempted to drill at this location on Sol 867, but the initial drill location proved to be unsuitable for drilling as the rock broke apart during preliminary testing of the potential drill site. A nearby location was identified and a successful drill sample was acquired, using a lower percussion level for the drill, on Sol 883 on a target named Mojave [Jet Propulsion Laboratory Press Release(2015f)].

4.2 Analysis of Resource Usage

We performed a series of analyses to understand how mission resources were used during the campaign. This helps shed light on how effectively the team was able to make use of resources during the execution of the campaign. Most of the plots have a similar layout, showing sols across the X axis and resource usage along the Y axis. The plots indicate multi-sol plans by grouping sols with vertical bars. The first sol of a multi-sol plan is indicated with a bold font. For example, Figure 22 shows that Sols 782, 783 and 784 were part of a three-sol plan with Sol 782 the first sol of the plan.

When appropriate, the plots organize the rover's activities are coarsely grouped into several, high-level categories. See Section 2.4 for further information on how activities were organized into Campaign Science, Other Science and Engineering categories.

Campaign Science: Activities that directly contributed toward the campaign's objectives

Other Science: Activities that contributed to science objectives unrelated to the campaign under study

Engineering: Activities carried out primarily for monitoring and maintaining the health of the vehicle

Comm: Communication windows used for uplink and downlink

Heating: Warming up devices and, for energy plots, maintaining temperatures

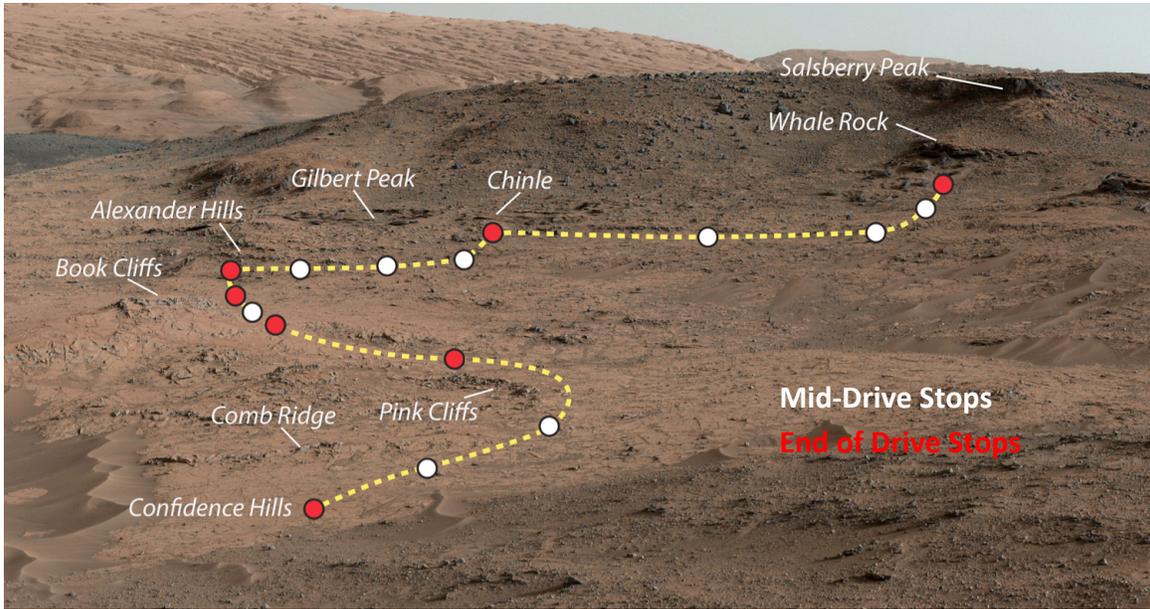
Margin: Time reserved in case activities run longer than expected (RCE duration plots)

Cleanup: Time reserved to force activities to stop in the event that allocated margin is exceeded (RCE duration plots)

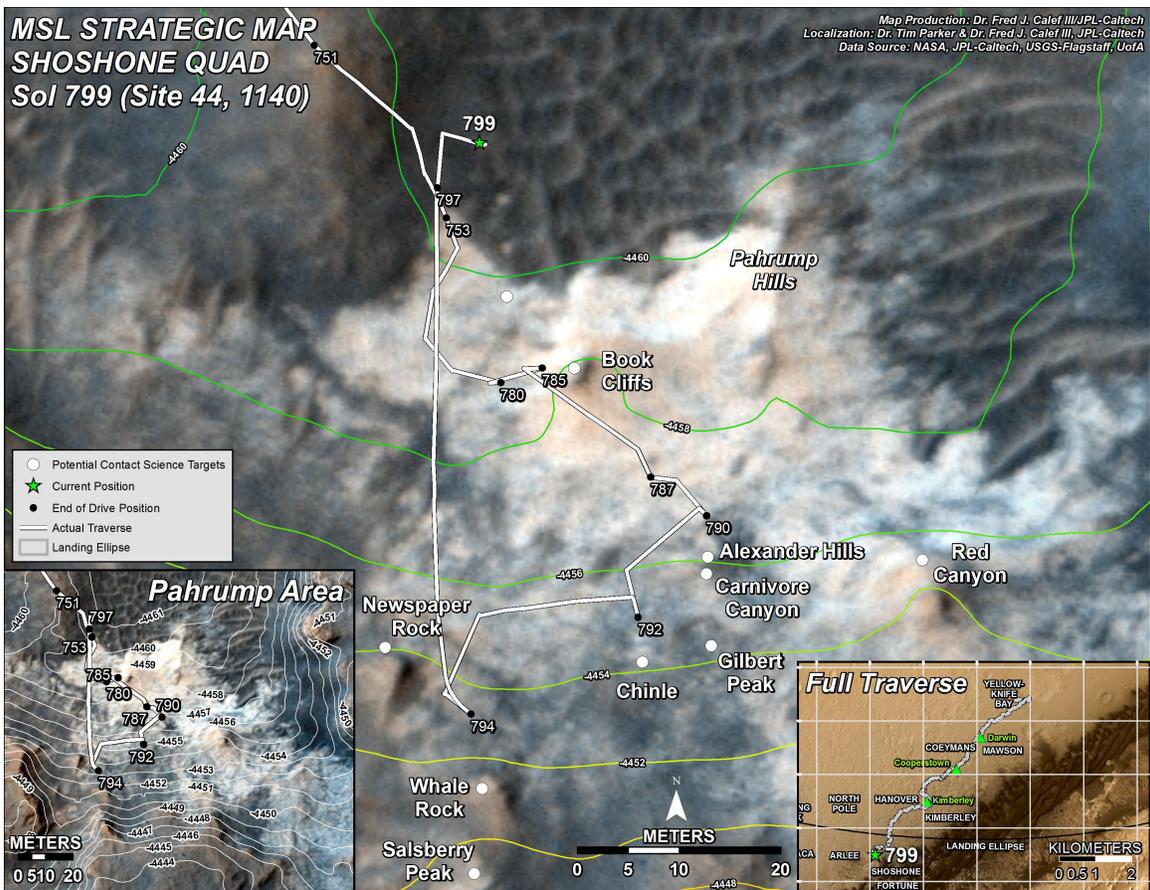
Wakeup/Shutdown: Time the RCE takes to boot, including initializing flight software, up and shut down, including ensuring flight software activity is cleanly stopped (RCE duration plots)

DAN Passive Only: Time spent with the RCE on performing DAN Passive with no other activity (RCE duration plots)

Idle: Time spent with the RCE on, but not performing activity (RCE duration plots)



(a) Traverse route projected onto Mastcam mosaic



(b) Traverse route projected onto HiRISE imagery

Figure 21: Traverse map of Curiosity during Pahrump Hills Walkabout Pass 1 (NASA/JPL-Caltech/Univ. of Arizona).

CPU: Quiescent energy used by the rover when the RCE is powered on (energy plots)

Asleep: Quiescent energy used by the rover when the RCE is powered off (energy plots)

We begin with an analysis of how the team made use of the RCE (Section 1.2) during the campaign. We analyzed RCE duration allocation because this resource is a significant focus of uplink planning and it often a limiting factor in how much activity can be performed. While the rover is capable of performing some activities while the RCE is off, including REMS, RAD, APXS, SAM, and CheMin analyses, much of what the rover does requires the RCE. For example, the RCE is required for driving, using the arm and conducting remote sensing operations.

As with a conventional computer, the RCE is capable of performing multiple activities in parallel, but due to resource conflicts, there are limitations on what activities can occur in parallel. For example the DAN instrument can operate in passive mode in parallel with any activity but a remote sensing activity cannot be performed at the same time as a drive because both activities require use of the rover's mast. For cases in which the RCE was performing multiple types of activities in parallel, we used the following ranking to determine how to categorize this time:

1. Wakeup/Shutdown
2. Campaign Science
3. Other Science
4. Comm
5. Engineering
6. Heating
7. Cleanup
8. Margin
9. DAN Passive Only
10. Idle

The ranking was selected to emphasize the most important activities for the campaign. For example, if the rover was conducting remote sensing for the campaign at the same time it was heating steering and drive actuators in preparation for a subsequent drive, we classified this time as Campaign Science.

Figure 22 shows the allocation of RCE duration for each sol of the campaign. Each sol shows the predicted allocation of RCE duration, according to the activity plan, and the actual allocation based on vehicle telemetry received after the plan was executed.

The first observation we draw from Figure 22 is the decrease in overall RCE usage across multi-sol plans. For example, comparing the sum of RCE durations for the two-sol plan spanning Sols 792, 793 to the amount of RCE duration for a corresponding number of single-sol plans, Sol 780 and sol 781, shows that in general less RCE activity is performed during multi-sol plans.

There is one exception to this pattern for the two-sol plan spanning Sols 785 and 786. There are a few reasons the second sol of this plan had a higher level of activity. There was an unusual increase in Comm activity for this sol. Most sols have four orbiter relay windows, but this sol happened to have five. In addition, the X-band window in this plan was an hour longer than a typical X-band window. The mission was preparing for an upcoming flight software update and used the long X-band window to upload portions of the next version of flight software. The sol also included a periodic engineering maintenance activity that required several, short-duration activities on the RCE at various times throughout the sol. This results in an increase in the Engineering and Wakeup/Shutdown durations for this sol.

This pattern of decreased overall activity during multi-sol plans is likely due to two reasons. First, timeline complexity limits how much activity the team is able to include in a single tactical shift. This productivity factor was discussed in our conceptual model of mission productivity, in Section 2.2.2, and reported by our interview participants in Section 3.3. A second factor is that the team is limited in the types of activity that can be performed after a drive without ground-in-the-loop. This constraint has a strong influence on the distribution of RCE duration allocation

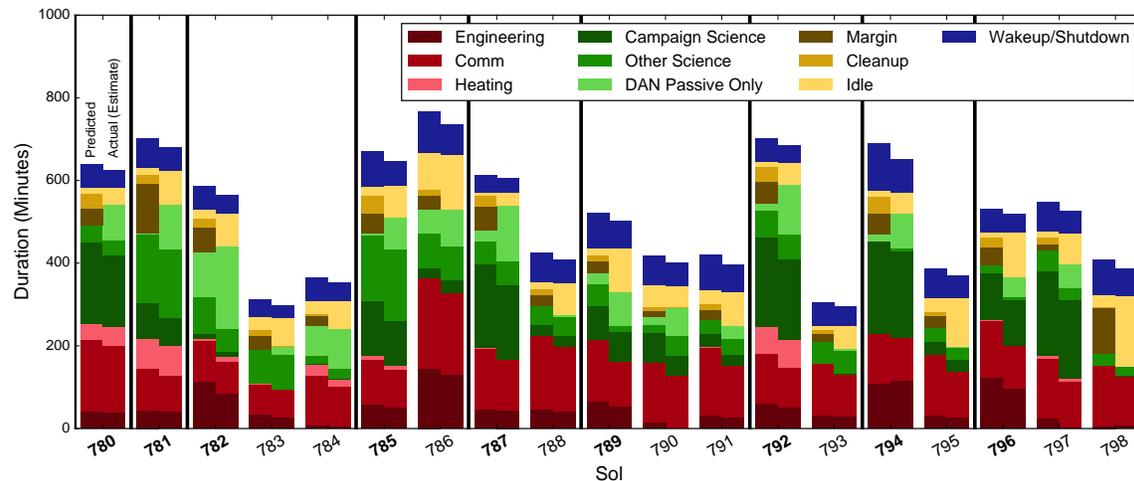


Figure 22: Allocation of RCE duration for Pahrump Hills Walkabout Pass 1 campaign.

within a multi-sol plan. For example, a significant drop in RCE duration is seen on Sols 788, 793 and 795 following drives on Sols 787, 792 and 794.

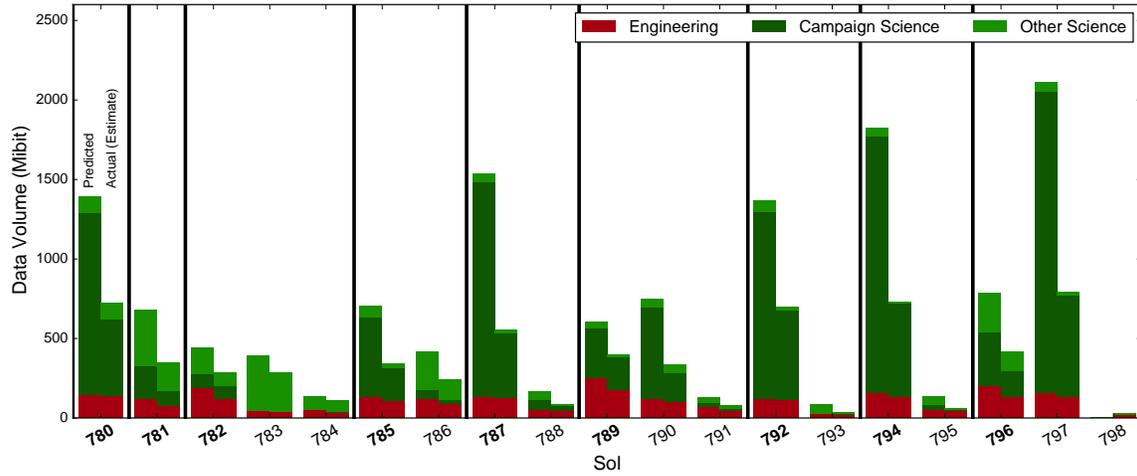
Another constraint related to ground-in-the-loop is that any activity that is needed to make decisions for a subsequent plan must be performed prior to the decisional pass for the plan. This is another reason that RCE duration allocation tends to be higher on earlier sols of a multi-sol plan and drops off in the later sols.

The second significant observation we make from Figure 22 is the decrease in RCE allocation toward campaign objectives in multi-sol plans following drives. This is because ground-in-the-loop is required to perform the majority of the activities needed to accomplish the campaign objectives. The exception to this is that the ChemCam blind targeted activity can be performed following a drive, without ground in the loop. This activity represents the small amount of Campaign activity on Sols 786, 788, 791, and 795. This emphasizes how the team depends on ground-in-the-loop to achieve high levels of productivity when conducting these types of campaigns.

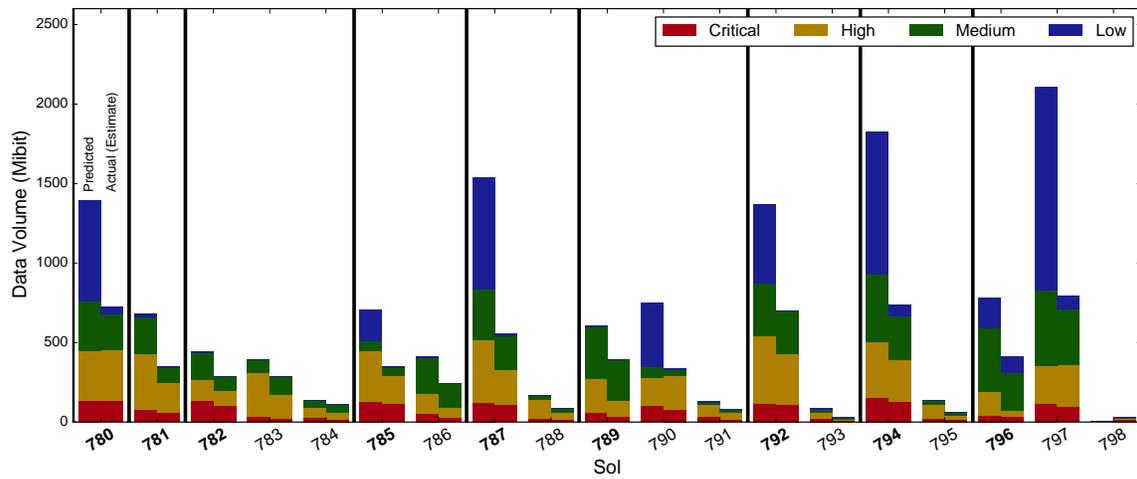
A third observation from Figure 22 is the amount of time the RCE awake but not conducting productive activity. We consider duration categorized as Idle and DAN Passive Only as unproductive time as this is time the RCE could be performing additional activity toward achieving science objectives. Note, this is not intended to discredit the value of DAN passive observations, which do contribute toward science observations. But, because DAN passive can be done in parallel with other activities, this is time that does not take full advantage of the RCE.

Comparing the predicted vs. actual RCE allocations, it can be seen that in most cases the team's intent was to have less allocation of this non-productive duration. There is relatively little Idle and Dan Passive Only duration in the predicted allocations. The increase in these allocations in the actual durations is largely due to activities taking less time than predicted to execute. The plot shows that, overall, the prediction of activity duration is fairly good in that there are not drastic decreases in actual run times. Thus, the increase in unproductive duration due to duration over-modeling is relatively small. The bigger increase is due to the predicted allocation to include Margin and Cleanup durations. Margin would be used if activities ran longer than expected. But this case is rare and, as a consequence, the time allocated for Margin becomes unproductive time. Similarly, Cleanup is needed if Margin is exceeded. But, exceeding Margin is even more rare, so this time also becomes unproductive time.

Figure 23 shows how the acquisition of data was allocated. Figure 23 (a) shows how collected data was allocated across intent categories and Figure 23 (b) shows allocations across priority bins. Priority bins are used to prioritize the order in which collected data comes down in communication passes. The mission has 99 different bins, but they are coarsely grouped into Critical, High, Medium and Low. During tactical operations, the team primarily focuses on what data goes into the Critical bins. Any data needed to make decisions on the next plan will be prioritized as Critical and the team carefully checks the predicted data volume of communication passes to ensure decisional data will fit within the passes up to and including the decisional pass. At certain times of the mission, when the team has collected a large amount of data in a short period of time, the team is also concerned with onboard data storage. However, during this campaign, there was sufficient onboard storage that this was not a concern.



(a) Allocation of data volume by intent category



(b) Allocation of data volume by downlink priority

Figure 23: Allocation of data volume for Pahrump Hills Walkabout Pass 1 campaign.

Figure 23 (a) shows a similar pattern to Figure 22 in that the allocation of data to campaign objectives drops significantly after drives in multi-sol plans (e.g. Sols 788, 791, 783).

The data volume allocation plots show very large discrepancies between predicted and actual volumes on some sols: e.g. 780, 787, 792, 794, and 797. These correspond to drive sols and are largely due to the uncertainty in how much data volume would be collected for the MARDI sidewalk videos that were acquired throughout each traverse. As discussed in the description of Sol 780 above, the MARDI instrument's use of change detection to determine when frames are saved coupled with uncertainty on the duration of the drive and post-drive activities posed a challenge in predicting the data volume that would result from the video. These plots indicate that the predictions were highly conservative and the instrument acquired far less data.

Recall from the discussion of Sols 972 and 973 above, the team was keeping an eye on overall data volume collected to ensure that they would receive the data they needed in time to make decisions for the second walkabout pass. As such, the team descoped several activities in the plans for Sols 792, 793, 794, and 795. Although Figure 23 shows that the team significantly over-modeled the amount of data that would be collected during these plans, it is not clear that they made an incorrect decision about descoping data on these sols. It is the high priority data that is of interest in this case, and Figure 23 (b) shows that the predicted vs. actual high priority data was much closer. The team did still over-model the high priority data. However, given that the team was just trying in general to reduce high priority data volume in the way of the data they needed to inform the second pass, and the challenge of trying to manage this with priorities on re-usable sequences, it seems it was still a reasonable decision to these activities. Ideally, there would be a simpler way for the team to inform the vehicle about which data needed to be received and by when.

While there are cases of large predict vs. actual data volume discrepancies in Figure 23 (b), the predicts for critical data volume are relatively accurate. This may be partly due to the fact there is less overall critical data volume collected than other priority bins. It is likely also a reflection of the importance of getting critical data volume predicts accurate to support operations. This is the data that must be carefully examined to ensure it is received on decisional passes. The accuracy of other bins are typically less important as the team can usually wait until vehicle telemetry is received providing the actual data volume that was collected. The main exception to this are situations when onboard storage is close to full and the team must restrict overall data collection in their plans.

Figure 24 shows the predicted and actual performance of the orbiter relay communication windows in terms of number of Mebibits transferred. Handover-to-Decisional represents the communication passes between the start of the plan through the decisional pass, the last communication window for which the downlinked data will be available in time for the next planning session. The amount of decisional data is particularly important for tactical operations as any data required to make decisions during the next tactical shift must come down prior to or during the decisional pass. Handover-to-Handover shows the amount of data downlinked during a single sol. For multi-sol plans, we also show a total Handover-to-Handover, representing the sum of downlink from all communications windows in the multi-sol plan. The difference between actual and predicted shows that, in general, the pass performance predictions are good.

Figure 25 shows the predicted vs. actual warmup heating duration in the plans. We looked at this heating activity as it has two types of impacts on mission activity. First, heating requires energy and in certain times in the mission the amount of energy needed to heat devices, or at least predicted to be needed, can limit the amount of activity the team can perform. Second, warming up devices takes time, and the device cannot be used until the warmup activity has completed. This can sometimes reduce the amount of time available to perform activities. The team will often look for activity to perform in parallel to warmup heating. For example, pre-drive imaging may be done while the mobility subsystem is being heated, or engineering maintenance activities may be done while the mast and cameras are being heated.

Figure 25 shows that there is often significant over-modeling of heating durations. This was discussed during interviews, in Section 3.3. The need for conservatism and the overhead of generating thermal tables results in over-modeling predicting heating durations. Overall, there was relatively little heating during the Pahrump Hills campaign, compared to other campaigns. There was a total of 11 hours of predicted warmup heating and 4 hours of actual warmup heating. This lower amount of heating is due to the campaign being conducted in a warmer part of the Martian year, as shown in Figure 6.

Figure 26 shows how vehicle energy was allocated during the campaign. Note that the distribution of actual energy across categories is an estimation based on vehicle telemetry. The vehicle reports an aggregate energy load over time. The aggregate load at each time point is the sum of the loads of all currently executing activities, including background and survival activity. The vehicle telemetry does not indicate the specific load that individual activities contributed to

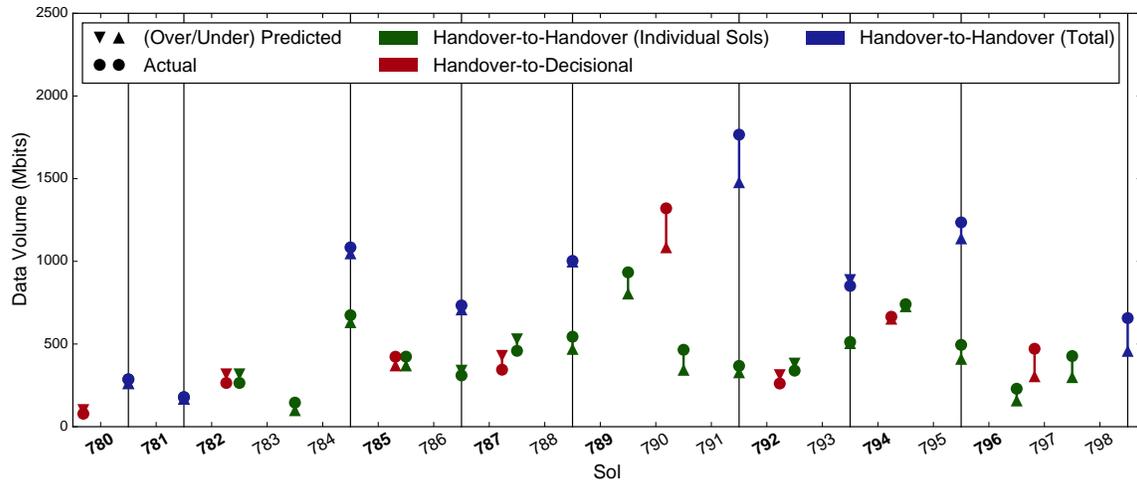


Figure 24: Downlink performance for Pahrump Hills Walkabout Pass 1 campaign.

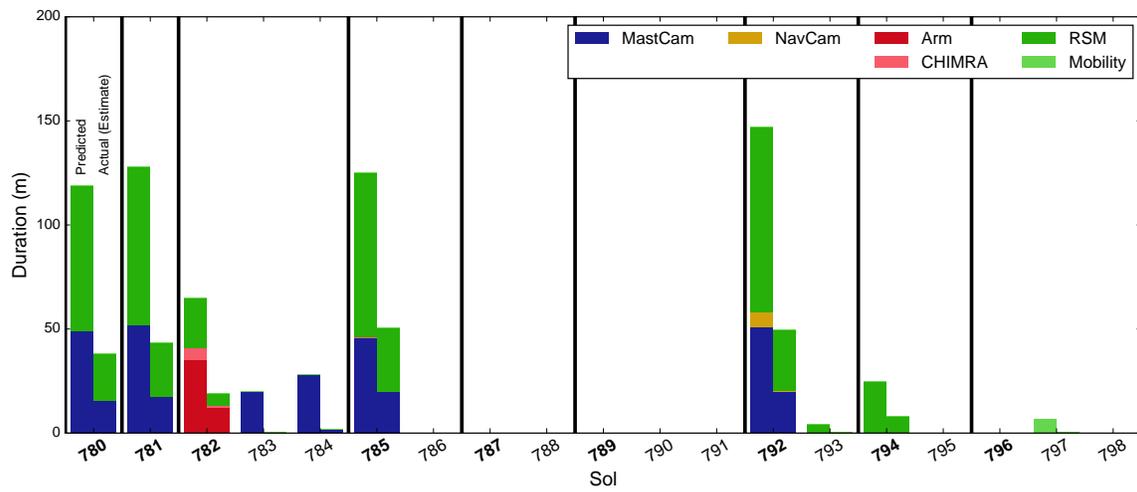


Figure 25: Summary of device warmup activity for Pahrump Hills Walkabout Pass 1 campaign.

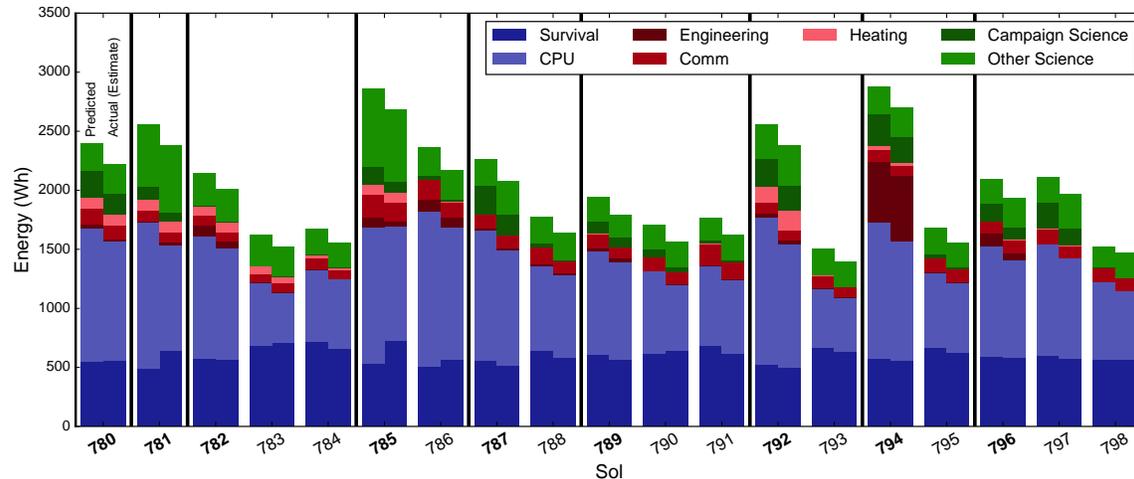


Figure 26: Allocation of energy for Pahrump Hills Walkabout Pass 1 campaign.

this aggregate load. We developed a novel approach to estimate these individual loads using a Bayesian model. The Bayesian model expresses the probability of the activities having a particular load given the evidence in the observed telemetry. The model is seeded with predicted activity loads, obtained from tactical plan modeling, predicted activity start and end times and actual activity start and end times, from vehicle telemetry.

One known limitation of the energy plots is due to how the mission chose to model the “background” energy used by the rover when it is asleep vs. when it is awake. The activity plans use one type of activity to represent the background energy when the CPU is powered on and a separate activity to represent the background energy when the CPU is off. In other words, both of these activities represent the background energy. A third type of activity represents the energy consumed by survival heaters across the full planning horizon. This makes it difficult for the estimation model to differentiate among energy needed for the CPU vs. energy for survival heating vs. energy needed for other background electronics throughout the sol. To simplify the problem, we decided to group the energy needed for background electronics with survival heating into a single category labeled “Survival”. We kept “CPU” as a separate category because it is important to differentiate the energy allocation between times when the CPU is off and when it is on. However, the exact distribution of estimate energy between “Survival” and “CPU” is not as reliable as the estimation of the two combined categories.

The shape of the plot in 26 roughly corresponds to the RCE duration allocation plot from Figure 22. The plots show a similar trend in the allocation of overall activity as well as allocation toward Campaign objectives. This makes sense since most of the activity performed during the Pahrump Hills Walkabout Pass 1 objectives required the RCE to be powered on. The main difference between these two plots is the difference in total resource allocations between Sols 785 and 786. In terms of RCE allocation, Sol 786 was higher. But in terms of energy allocation, Sol 785 was higher. This is because Sol 785 included a CheMin analysis activity which consumed energy while the RCE was powered off.

The plots show that there is a consistent over-modeling of the amount of energy the rover’s activity will take. In considering the difference between predicted and actual energy, it is important to realize that, roughly speaking, about 200 Wh of energy can be used for an hour of remote sensing or for driving. As such, the differences between predicted and actual energy shown in Figure 26 represents a significant amount of additional activity that could have been performed. We will have more to say about this below.

In Figure 27, we look at the predicted and actual state of charge for the rover’s battery during the campaign. The uplink operations team primarily focus on two main aspects of battery state of charge. The minimum state of charge is important as the team must ensure that the state of charge stays above 40% to avoid damaging the battery. The handover state of charge is the available energy in the battery at the end of one sol’s plan and the beginning of the next. The most important handover state of charge happens on the boundary between planning sols. For these sols, the supratactical team provides handover state of charge guidelines. This is the state of charge that the supratactical team requires for the projected activities in the look-ahead plan. If the tactical team is struggling to meet the guideline, they

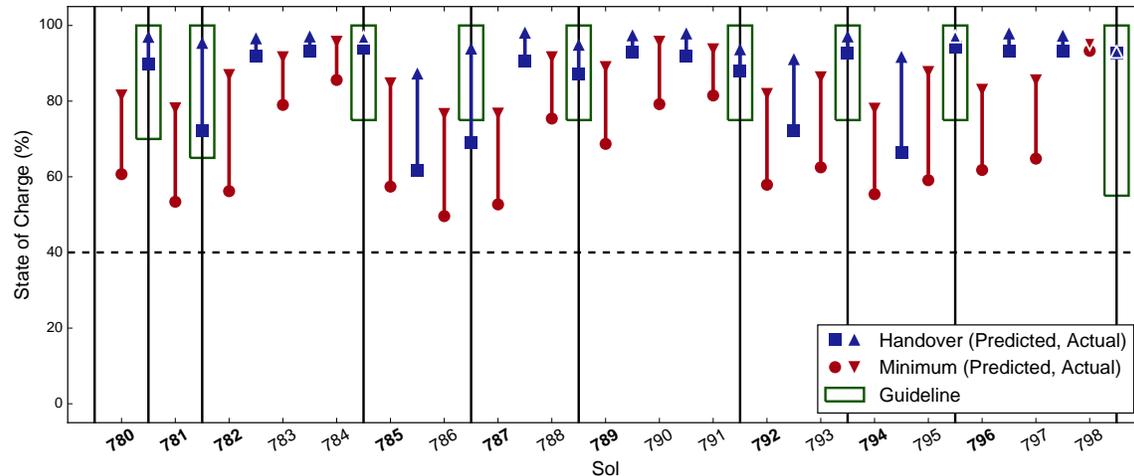


Figure 27: Summary of battery state of charge for Pahrump Hills Walkabout Pass 1 campaign.

may choose to violate the guideline with the risk that they may not be able to accomplish all of the activities in the next plan.

There are a couple important observations to draw from Figure 27. First, it is apparent from the minimum and handover state of charge values that energy was not a limiting factor during this campaign. All of the predicted minimum state of charge values are well above the 40% limit. And all but one handover state of charge did not meet the supratactical guideline and most were well above the guideline. The case in which the team did not meet the supratactical guideline was due to a late correction made to the plan that required an increase in the duration the RCE was powered on. Due to the lateness of the change, rather than try to re-work power in the plan, the team talked with the Supratactical Uplink Lead and agreed to violate the handover guideline.

The reason that energy was not a significant factor is due partly to the Martian season in which the campaign took place as well as the type of activity performed during the campaign. First, the campaign took place during the Martian spring when, due to the relatively warmer temperatures of spring, less energy is needed for heating. Second, for the most part, the campaign did not include high-energy consuming activities. The exception was that two sols included overnight CheMin analysis, in the plans for Sol 785 and 794.

A second important observation from Figure 27 is the large discrepancies between the predicted and actual state of charge values. The actual state of charge is consistently higher than the predicted state of charge. The challenge of predicting battery state of charge was discussed with the interview participants in Section 3.3 as was the impact this imprecision can have on productivity. The interview participant noted that there are often times in which activity is descoped because the predicted minimum state of charge drops below 40% or the handover state of charge does not meet supratactical guidelines. However, subsequent telemetry indicates that the state of charge was higher than predicted and would have allowed the activity. These cases did not occur during the Pahrump Hills Walkabout Pass 1 campaign as energy was generally an issue during this time.

The sols in which predicted state-of-charge is close to actual are those in which the battery was at a higher state of charge. This was explained in Section 3.3. The points in the time where both predict and actual state-of-charge are near 100% results in the model that stop the accumulation of error in the predicted state-of-charge.

Figure 28 shows the how energy from the MMRTG was used during the campaign. The black line shows the energy output from the MMRTG each sol, orange shows the energy used by the rover and gray is the energy that was shunted. Shunt energy is the energy produced by the MMRTG that exceeds the demand of rover activity after the batteries have been charged. It represents unused energy. The plot also shows when there was a net depletion or charge to the battery. Similar to Figure 27, the MMRTG energy plot shows that energy was not a limiter for this campaign. The two sols in which overnight CheMin analysis were performed show up as sols in which there was a net depletion of the battery.

We were interested in estimating how much extra duration could have been used for the campaign's objectives given the amount of unproductive duration from Figure 22 and shunt energy from Figure 28. We used a rough, but

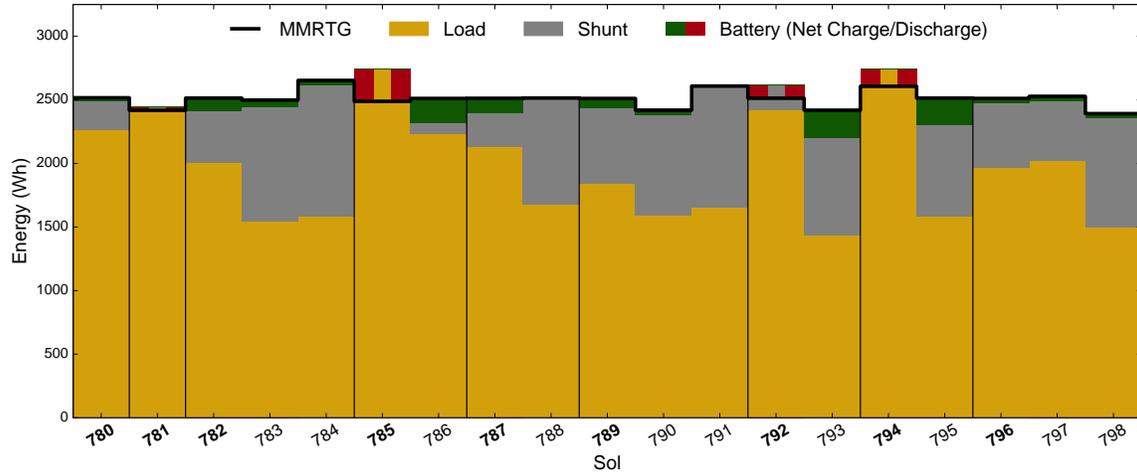


Figure 28: Summary of battery and MMRTG energy for Pahrump Hills Walkabout Pass 1 campaign.

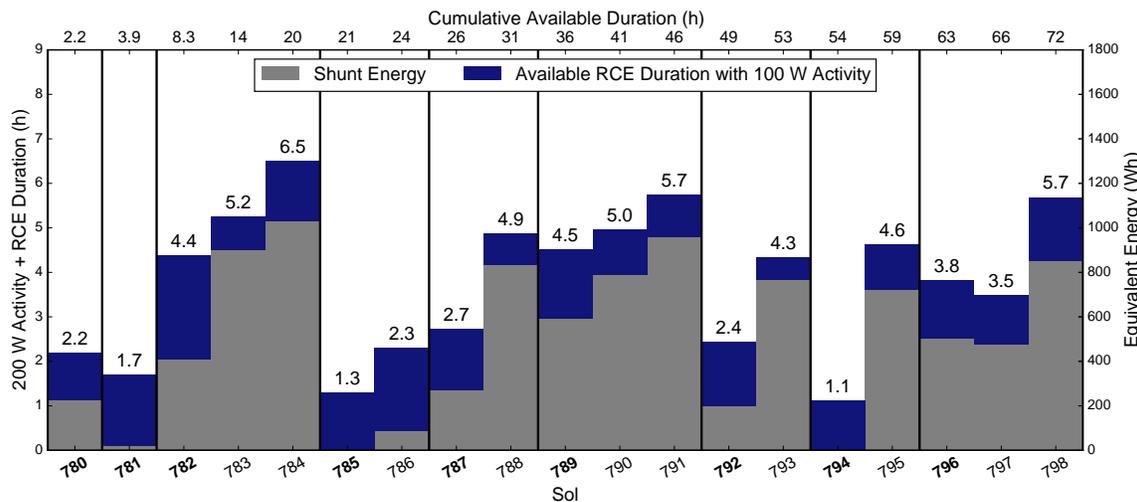


Figure 29: Estimate of extra duration availability for Pahrump Hills Walkabout Pass 1 campaign.

conservative, heuristic that 200 Wh of energy corresponds to 1 hour of remote sensing or driving and that 2 hours of unproductive RCE duration corresponds to 1 hour of remote sensing or driving. Figure 29 shows the results of applying that heuristic to the Pahrump Hills Walkabout Pass 1 campaign. Gray is the extra duration that could have been used from shunt energy and blue is the duration that could have been used from time the RCE was otherwise on but unproductive. A cumulative sum of the estimated extra duration is shown at the top of the plot. There was a total of 72 hours of estimated extra duration throughout the 19 sols of the campaign.

The nature of the Pahrump Hills Walkabout campaign, which consisted of driving and performing targeted remote sensing, required activities during daylight hours. While this analysis shows that there was a significant amount of extra duration available, it is not clear if there was unused time during the Martian day in which to use this extra time for the campaign. Therefore, we analyzed the RCE duration allocations similar to Figure 22 but just for the times between the Martian sunrise and sunset. Figure 30 shows the result. The black horizontal line above the RCE duration allocation stacks represents the total amount of daylight duration available each sol. The number above this line is the number of hours in which the RCE was off or the RCE was on but unproductive (Idle or DAN Passive Only). A cumulative sum of unused daylight duration is shown at the top of the plot. The plot shows that there was in fact enough unused daylight hours to make use of the extra available duration from Figure 29.

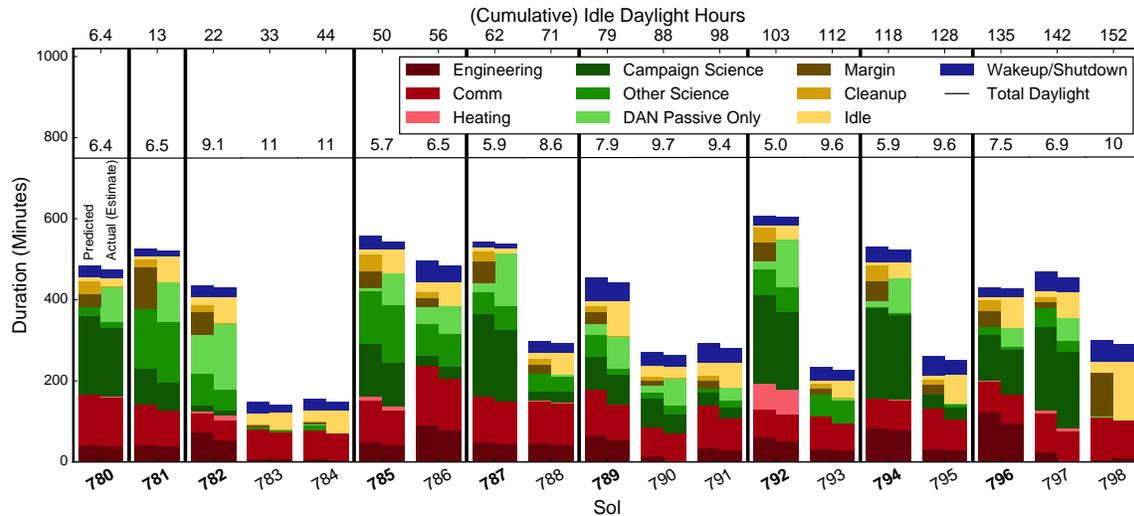


Figure 30: Allocation of RCE duration during daylight hours for Pahrump Hills Walkabout Pass 1 campaign.

Figure 29 can be viewed as partial measure of the campaign’s productivity. This is unused capacity of the rover to accomplish the campaign’s objectives. A more productive campaign could have made use of this extra duration and, potentially, reduced the number of sols required to complete the objectives.

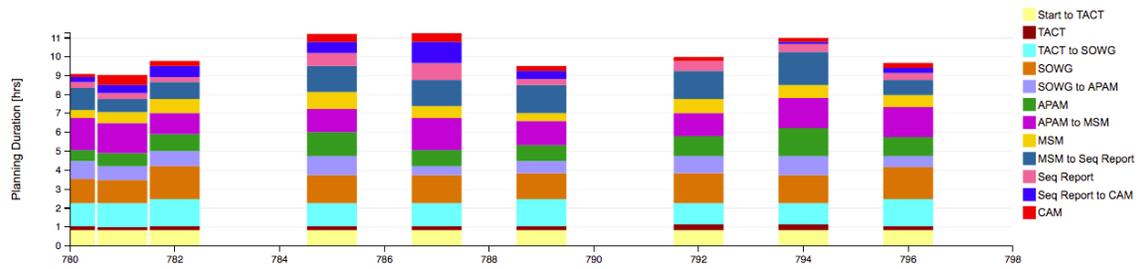
Of course, the vehicle resources are not the only constraint on mission productivity. As noted previously, timeline complexity plays an important role. The mission does an excellent job of tracking the time spent at different stages during the tactical timeline. This can be roughly interpreted as a measure of the tactical team effort required to create each plan’s command products. Figure 31 shows the tactical timeline durations for the Pahrump Hills Walkabout Pass 1 campaign. During this time of the mission, the nominal tactical duration was 11 hours for weekend plans and 10 hours other. The plot shows that in all but three cases, the team finished within the nominal duration. In two cases, the team finished an hour earlier than the nominal duration. It is interesting to note that the two single-sol plans resulted in the shortest tactical timeline duration.

There were three cases in which the team exceeded the nominal timeline duration. The team prefers to avoid exceeding the nominal timeline as it results in wear on the team and can increase the risk of errors. The APAM meeting ran a little longer than usual for Sols 785 and 794 shifts. On Sol 785, APAM ran long due to correcting several adjustments to the plan during the meeting and discussions related to concern regarding a potential HGA stop stop issue, as discussed for Sol 781 in Section 4.1.2. On Sol 794, APAM took extra time to effectively communicate the drive plan with the team. This is an example of the importance of communication expressed by interview participants in Section 3.3. All three shifts that ran long encountered in the latter part of the shift with various problems with tools used to assist preparing and validating command products. This was discussed in Section 3.3 under “Overhead in Command Product Generation”. The Tactical Uplink Leads for these shifts also noted that the shifts had several trainees. This was also discussed during interviews as a challenge for multi-sol plans as the larger volume of activity creates difficulties for team members learning the tactical process.

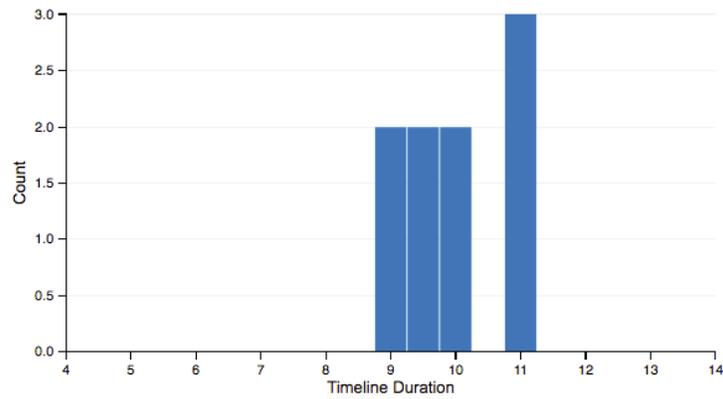
4.3 Significant Productivity Factors

The Pahrump Hills Walkabout campaign benefited from extensive strategic planning. In discussions with the team, it is often described as one of the most organized campaigns of the mission. This was partly due to the favorable geography, which enabled a single panorama to provide extensive coverage of the area to be explored, and to the substantial strategic planning that went into designing the multi-pass exploration strategy. This provided a well-structured plan for the team to follow when developing tactical plans.

During the case study interviews, one of the scientists pointed out that one of the challenges the team faced when starting the Pahrump Hills Walkabout was the abrupt shift in mission activity. For several months leading up to this campaign, the major focus of the mission had been maximizing drive distance to quickly reach the base of Mount



(a) Tactical durations per planning sol



(b) Tactical duration histogram

Figure 31: Tactical timeline durations for Pahrump Hills Walkabout Pass 1 campaign.

Sharp. The arrival at Pahrump Hills required an abrupt shift from long-range driving to performing a detailed study of an area. The scientist reported there was a learning curve in preparing for the strategic planning that was needed. This included determining what types of discussions were needed and anticipating the needs that the supratactical team would need weeks in advance.

We used the Tom Sawyer Perspectives visualization tool to generate an overview of the campaign data we collected. Figure 32 shows the visualization for the Pahrump Hills Walkabout campaign. The diagram in Figure 32 (a) shows all high-level intent categories and the sols that contain activities to accomplish these objectives. Figure 32 (b) shows just the high-level Campaign Science intent category to make it easier to see which sols are accomplishing campaign-specific objectives. In both plots, the width of the link from intent category to sol reflects how many activities in the sol are achieving that intent, with a thicker line indicating a larger number of activities contributing to that intent.

Links between two sols indicate that data acquired in the first sol is being used by some activity on the second sol. For example, the team typically uses Navcam imagery from a prior sol to create targeted ChemCam observations on a subsequent sol. Similarly, on a drive sol, the Rover Planners will make use of Navcam and Mastcam imagery from previous sols.

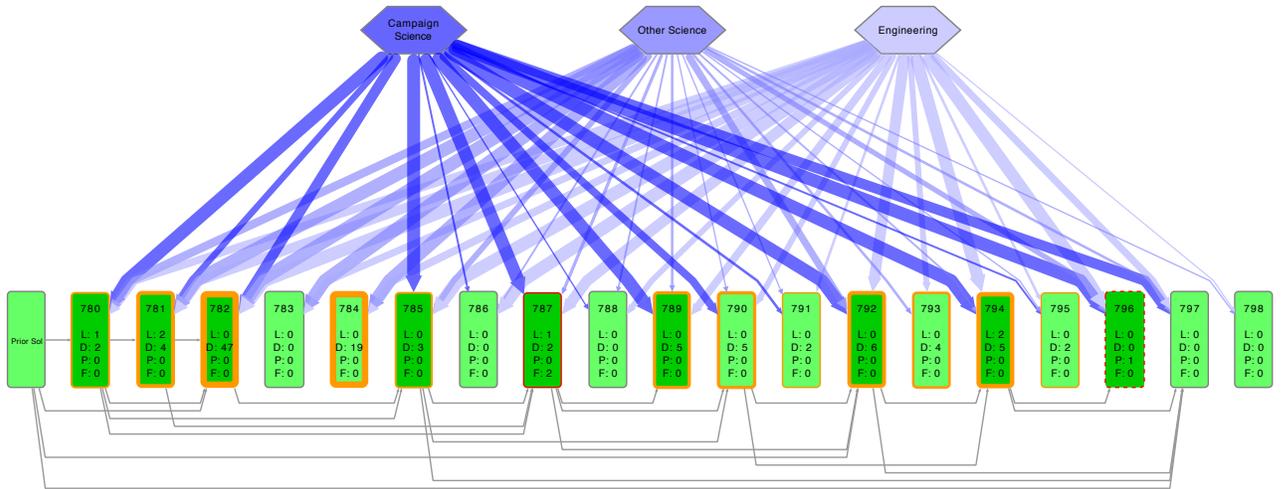
The color and outline of sols in Figure 32 provide a visual indication of properties of that sol. Dark green sols indicate that this is either a single sol plan or the first sol in a multi-sol plan. Light green sols denote subsequent sols in multi-sol plans. A solid orange outline indicates there were activities considered for inclusion in the plan for this sol but they were either deferred or descoped. A dotted orange outline indicates an activity was included in the sol, but the planned activity was limited in scope for some reason. For example, if a drive activity could not reach the desired goal due to limited viewshed, we considered it a limited activity. A solid red outline indicates an activity in this sol failed when it was executed on the vehicle. A dotted red outline indicates an activity had partial success. In all cases, the thickness of the outline reflects the number of activities there were of that type in the sol.

The visualization helps identify periods of both high and low productivity in the campaign. For example, it is easy to identify sols in which a large amount of objectives are achieved. In contrast, it also helps identify sols in which little or no campaign objectives were accomplished, sols in which activities had to be descoped, deferred or limited, and sols in which activities did not have full success during execution.

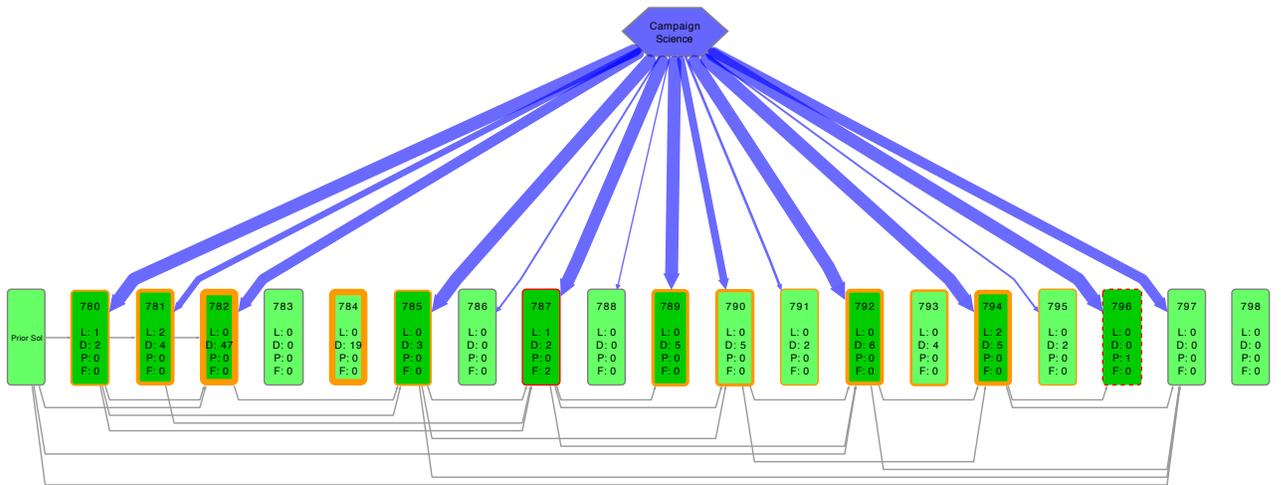
Table 1 provides a breakdown of the sols of the campaign in terms of how they contributed toward accomplishing the campaign objectives. Sols labeled “Campaign” were those that directly contributed to the campaign objectives with remote sensing and/or drives. “Campaign Multi-Sol” sols are those in which significant activity was performed toward the campaign objectives as part of a multi-sol plan, either due to a weekend or restricted planning. The reason for calling these sols out separately is that the presence of the multi-sol plan limited the team’s options for these sols. For example, had there not been a multi-sol plan, the team may have opted for to move up activity that was performed in a subsequent plan (e.g. a drive activity) which would have reduced the overall number of sols required to achieve the campaign objectives. The “Extra Drives” label denotes sols in which unexpected drives were required. “Post-Drive Multi-Sol” sols were those sols in which the team was not able to achieve substantial campaign objectives due to lack of data following a drive during a multi-sol plan. The sols labeled “Deferred” were the sols in which campaign objectives were descoped due to potential conflict with the Comet Siding Springs observations. Finally, “Runout” are sols of very low activity that used in cases the team had to create multi-sol plans but the tactical timeline capacity did not allow for sufficient time to develop activities for all sols of the plan.

It should be noted that Table 1 is a coarse breakdown of sols intended to characterize the productivity factors that impacted achieving the campaign objectives. In particular, all but 4 sols included some activity that contributed in at least a small way toward the walkabout objectives. However, in most cases, such activity on non-“Campaign” sols was due to ChemCam blind targeted observations. While this is useful data, it was not a significant contribution to the walkabout objectives. Similarly, the sols labeled “Extra Drives” were important toward accomplishing the walkabout, but given that they did not correspond to the initial strategic end-of-drive objectives, it is informative to consider them separately to help understand the productivity factors that caused them.

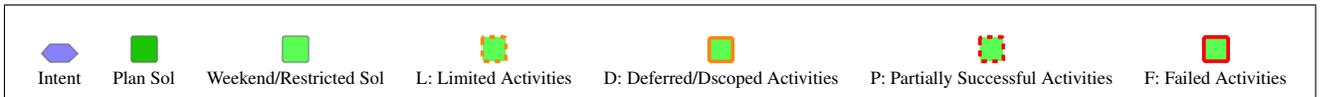
In terms of productivity factors that resulted in the largest delay in completing the campaign objectives, Table 1 shows that transition into restricted sols was likely the biggest negative factor. As noted by several participants in the case study interviews, the mission experiences a significant drop in productivity of drive-oriented campaigns during restricted periods. In the case of the Pahrump Hills Walkabout, after each drive, the team required post-drive imagery in order to perform the desired observations at each of the end-of-drive location. There are two reasons for this. First, although the initial mosaic of the formation was sufficient to identify the major pieces of outcrop the team wished to visit, there was not sufficient detail from this range to identify the specific targets of interest. Second, targeting these



(a) Visualization showing all intent categories



(b) Visualization showing only Campaign Science intent category



(c) Legend

Figure 32: Visualization of the Pahrump Hills Walkabout 1 Campaign.

Sol Type	Sols	Number of Sols
Campaign	781, 785, 789, 792, 794, 796	6
Campaign Multi-Sol	790, 797	2
Extra Drives	780, 787	2
Post-Drive Multi-Sol	786, 788, 791, 793, 795, 798	5
Deferred	782, 783, 784	3
Runout	798	1
Total Sols		19

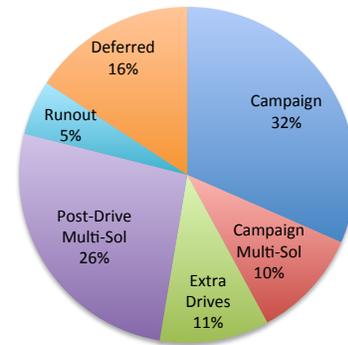


Table 1: Breakdown of sols for the Pahrump Hills Walkabout Pass 1 campaign.

observations, especially the ChemCam observations, requires precise knowledge of the location of the rover relative to the targets. Once the rover starts driving, this precise knowledge is lost due to localization uncertainty that results from driving across the Martian surface. As a result, the team must wait for post-drive imagery to arrive on Earth so that targets can be selected and targeted, and the rover must remain in that same location until it receives the command products in order to maintain targeting accuracy.

The second biggest negative factor from Table 1 was the descoping of Pahrump Hills Walkabout activities due to the potential conflict with the Comet Siding Spring observations. This illustrates how activities can result in unexpected complications due to unanticipated interactions among seemingly unrelated activities. In this case, the fact that a drive might impact the sun safety of comet observations was not taken into account during the strategic development of the comet observations.

This event also illustrates how the options for recourse the tactical teams has for responding to issues becomes increasingly limited as the tactical timeline proceeds. The conflict between the drive and the comet observations was identified during the SOWG meeting. Had the problem been identified earlier in the day, there might have been time to develop a workaround that would enable the drive and the comet observations. Or, at the least, additional, alternate activities could have been proposed to fill in the time made available by descoping the drive. Instead, given the limited time available for developing activities and modeling resources at this point in the timeline, the team was not able to take advantage of the extra resources made available by descoping the drive.

The final negative productivity factor shown in Table 1 is the “Extra Drives” category. The campaign ended up requiring 7 drives (Figure 21) rather than the 5 originally anticipated in the strategic plan (Figure 15). The first event that required an extra drive was the during Sol 780, the first sol of the campaign. Due to occlusions in the terrain, the Rover Planners had insufficient data to drive the full distance to Book Cliffs, instead, they needed to end the drive about 7 meters short. The post-drive imagery acquired from this location would provide the necessary data to complete the drive, which was performed on Sol 785. Note that autonav was not an option for completing this drive in a single sol. For one reason, given the concern regarding wheel wear, the use autonav was prohibited in this region due to the significant hazards to wheel wear. Even if the terrain had been more benign, there likely would not have been sufficient time in the plan for autonav to complete the drive, due to the slower pace at which autonav runs.

The second instance of an extra drive occurred on Sol 787 while attempting to drive toward Gilbert Peak. The terrain in this area had significant hazards to the rovers wheels and an extra drive was used to reach the location the science team wanted while avoiding excessive wheel damage.

The campaign included examples of how the capacity of the tactical timeline and the overhead involved in developing command products can limit activity as discussed in the interviews in Section 3.3. As an example, the team opted to descoped ChemCam observations due to concerns with managing the additional work with the already complex plan.

4.4 Significant Ground-in-the-Loop Cycles

One of the motivations for conducting these case studies was so that we could understand the reasons for ground-in-the-loop cycles. Toward that end, the following summarizes significant ground-in-the-loop cycles during the Pahrump

Hills Walkabout campaign.

- Mastcam and ChemCam observations at each end-of-drive location. As noted in Section 4.3, ground-in-the-loop was required to perform the remote sensing activities at each of the end-of-drive locations. The imagery obtained at the end of the drive provided the information needed by the science team to identify and target the observations that would contribute to the Pahrump Hills Walkabout. Sols in which this type of targeted activity was performed were: Sols 780, 781, 785, 787, 789, 792, 794, 796.
- Specific end-of-drive locations. Although the science team had a general idea of where they wanted the rover to stop for each of the anticipated drives at the start of the walkabout, as shown in Figure 15, the specific end-of-drive locations were refined with the latest imagery from the rover at the start of each drive to select a location that met science objectives.
- Drive planning. The development of each drive in the campaign was based data obtained from the previous drive. The post-drive imagery allowed the team to identify hazards and evaluate routes for reaching the desired end-of-drive locations. Sols in which drives were planned were: Sols 780, 785, 787, 790, 792, 794, 797.
- Refined comet observations. The team used the data from the comet observations in the Sol 780 plan to inform the observations that were performed on the closest approach during the weekend plan for Sols 782, 783 and 784.
- ChemCam observations of dump pile. The team wanted to acquire ChemCam observations of the sample that was dumped on Sol 781. Due to uncertainties in the exact location and shape of the dump pile and the precise pointing requirements of ChemCam, a ground-in-the-loop cycle was required to target the pile. The ChemCam observation of the pile was acquired on Sol 782.

4.5 Conclusions

The Pahrump Hills Walkabout Pass 1 campaign was highly successful. It achieved the objectives of characterizing the region and collect data to prioritize locations for close-contact observations on the subsequent walkabout. The campaign is considered by many on the team as a good model for conducting science campaigns.

The campaign benefited from a well-developed strategic plan at the start of the campaign along with effective communication between the science and engineering teams to refine objectives throughout the campaign. As stated in Section 3.3, several interview participants noted that the campaign scientist role can be particularly important in facilitating coordination between science and Rover Planner activities. This was certainly the case for the Pahrump Hills Walkabout as the campaign scientist worked daily with the Rover Planners to assist with defining drive objectives and mid-drive activities.

While it was highly successful, the Pahrump Hills Walkabout campaign highlights several productivity challenges faced by surface missions. Table 1 shows that there were several sols in the campaign in which the team has highly restricted in their ability to achieve campaign objectives. The restricted planning period had the largest negative impact to the campaign's productivity with 6 sols (or 31% of the campaign sols) occurring after drives while the mission was in the restricted operations.

The large impact of restricted sols emphasizes the reliance operations has on ground-in-the-loop cycle for conducting many types of science campaigns. Although the geometry of Pahrump Hills enabled the team to acquire a panorama of the region and develop a relatively detailed strategic plan, ground-in-the-loop was still required throughout the campaign. Section 4.4 described the types of decisions made with ground-in-the-loop cycles during the campaign. These included drive planning and targeted remote sensing.

The ground-in-the-loop requirement results in two forms of activity restriction. First, the team is limited when they can perform certain activities in a plan. Any data that is needed to make decisions for the next plan (such as end-of-drive imagery) must be completed before the decisional pass in the plan. This can result in the team not being able to make effective use of available duration that follows the decisional pass and can limit the quality of data that is acquired decisionally to ensure that it fits within bandwidth limitations of the communication passes.

The second type of restriction caused by the ground-in-the-loop requirements is that the team is highly limited in the activities that can be performed on the sols that follow the decisional pass if the rover's position has changed.

In order to conduct targeted remote sensing, contact science and effective drives the team requires knowledge of the rover's current position and local terrain.

Operations complexity also resulted in loss of productivity during the campaign. Unexpected interactions between the comet Siding Spring observations and the drive resulted in descopeing a large amount of campaign activity during the multi-sol plan for Sols 782 through 784.

Our analysis of resources during the campaign shows that vehicle resources, including energy and data volume storage, were not limiting factors for the team. In fact, there was sufficient energy available, unused daylight hours to perform an estimated 72 hours of additional activity toward the campaign objectives during this time. This indicates that there are sufficient vehicle resources available to increase mission productivity provided there are methods for overcoming the productivity challenges identified in this case study.

5 Artist's Drive Campaign

The team had spent nearly six months conducting an extensive exploration of Pahrump Hills. The exploration included three walkabout passes and a total of three drill samples at different elevations of the formation.

Further progress was delayed due to the detection of a short during a sample processing activity on Sol 911 [Jet Propulsion Laboratory Press Release(2015l)]. The team temporarily suspended normal arm and drive activities while the anomaly was investigated. Tests performed over the course of several sols enabled the team to determine the likely cause of the fault to be a transient short in the motor for the drill's percussion action [Jet Propulsion Laboratory Press Release(2015k)]. With this understanding of the problem, the team was able to resume driving and most arm-related activities. They performed additional analysis and process development to devise a strategy for performing drills in a way that would be more robust to a recurrence of the short.

In the meantime, the team left Pahrump Hills and explored a nearby area named Garden City filled with two-tone mineral veins [Jet Propulsion Laboratory Press Release(2015g)]. The scientists hypothesized that the veins may have been formed by two different fluids interacting with the rocks.

After the team's lengthy explorations at the base of Mount Sharp, it was time for the rover to move on and explore higher layers of the mountain. Thus, the team embarked on the Artist's Drive route.

5.1 Overview of Campaign

The Artist's Drive campaign differs from the Pahrump Hills campaign (Section 4) in that the goal was to travel through a large region and characterize it along the way, rather than to explore a particular location or contact. The team had been exploring the Murray formation, the material making up the basal layer of Mount Sharp. The strategic drive objective was the reach higher layers of Mount Sharp where HiRISE imagery indicated the presence of a second material, which the team named the Stimson formation. The team was especially interested in exploring the contact between these two regions. Orbital imagery indicated there was a promising contact at a location the team named Logan Pass, as seen in Figure 33.

With the Logan Pass objective in mind, the scientists and engineers developed a strategic traverse route to reach this location. The science team identified strategic imaging locations from which the rover would have a good viewshed for imaging the surrounding topography. Figure 34 shows the route and anticipated imaging locations. These would be refined during the course of the campaign as higher-resolution data is obtained from the rover.

In contrast to the walkabout approach employed at Pahrump Hills, the team chose to traverse Artist's Drive using the more traditional (at least for rover missions) linear approach to exploration. Here, the rover traversed the Artist's Drive Valley along the Lewis and Clark Trail. An initial Strategic Traverse Route was designed prior to the beginning of the campaign to guide tactical planning for the traverse. This route was generally followed until it became necessary to deviate to the north, on Sol 960, to improve visibility of terrain ahead on the path.

A major deviation occurred on Sol 963, when an interesting outcrop to the west was discovered. This feature, called "Logan's Run", included several rock structures poking out from the valley wall. These rocks were close to the boundary between Pahrump materials and the overlying rock, and they could yield new information about both rock layers. The rover diverted to study "Logan's Run" up close over the next 23 sols.

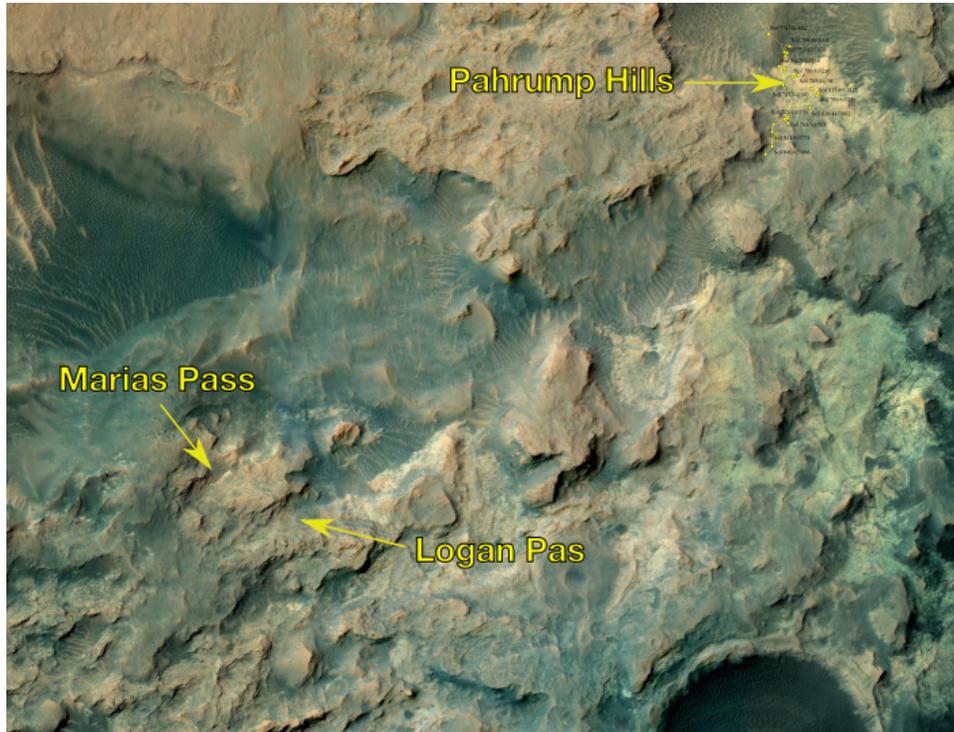


Figure 33: View of Logan Pass from HiRISE (NASA/JPL-Caltech/Univ. of Arizona).

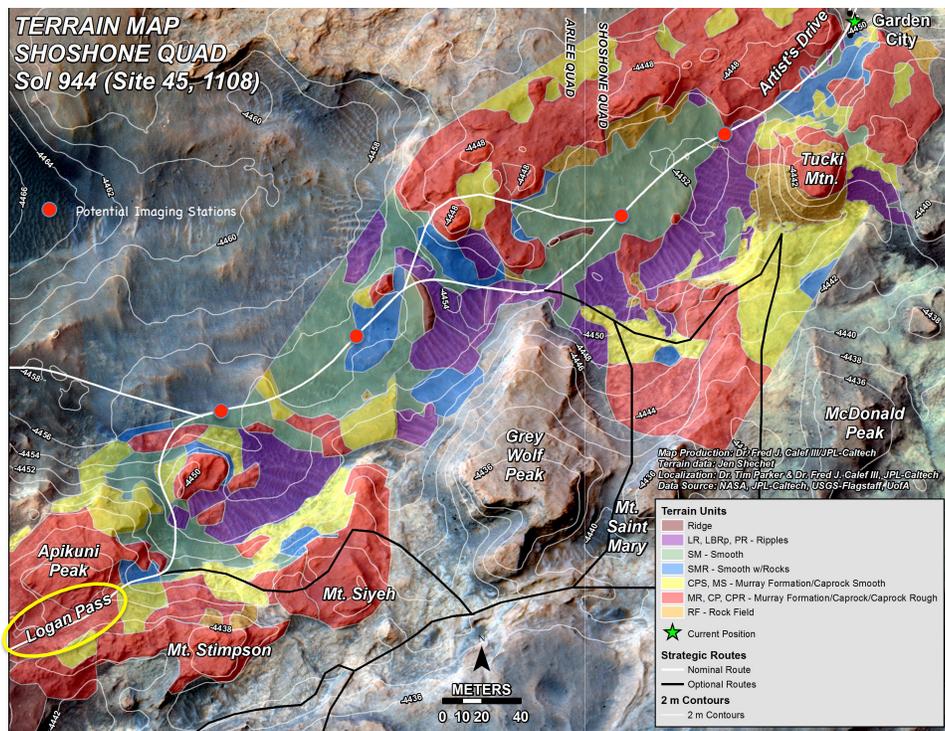


Figure 34: Strategic traverse route for Artist's drive including anticipated imaging locations.

5.1.1 Science Objectives

The high-level goal of this campaign was to reach and inspect higher layers of Mt. Sharp. The team sought to identify another outcrop of Pahrump-like material (Stimson over Murray formation). The traverse through Artist's Drive valley originally aimed for Logan Pass. During the campaign, the team discovered rock lenses sticking out of a hillside (later dubbed Logan's Run) and diverted to inspect them up close.

Although we refer to this period of time as the "Artist's Drive Campaign", it occurred during a part of the mission in which the overall goal was to make progress toward the next major scientific waypoint at Logan Pass. It was therefore not a scientific campaign as such, but it serves as an example of how mission science occurs and evolves along a traverse. The team sought to identify and study the geological contact between the Stimson and Murray formations, an example of which was found to occur at Logan Pass based on orbital imagery. The traverse through Artist's Drive valley originally aimed for Logan Pass, but during the campaign, the team discovered rock lenses sticking out of a hillside (later dubbed Logan's Run) and diverted to inspect them up close.

5.1.2 Summary of Activity

Following is a summary of the team's and rover's activities during each tactical planning cycle of the campaign. Note that the team had decided to work some Saturdays during this portion of the mission. This choice was made to help increase science return by reducing the number of multi-sol plans and increasing the number of sols with ground-in-the-loop. As will be seen below, there was one Saturday planning shift during the Artist's Drive campaign.

Sol 949: Sol 949 wrapped up investigations at the Garden City site and made the first drive of the Artist's Drive campaign. Final observations at Garden City included a targeted Mastcam observation (Anza Borrego), a Navcam dust devil search, and a DAN measurement of mineral veins to look for elevated H content. The DAN observation was considered to be the highest scientific priority for the plan. Additional ChemCam observations were considered, but not performed due to lack of available time before the drive.

The drive was planned to be 23 m using visual odometry. The drive length was constrained by the amount of time available before the decisional downlink pass (i.e., last chance to get data for use when planning the next sol). The drive strategy was influenced by previous experience with slippage on similar terrain, so it was decided to aim for the shallow end of the ripple near the end of the drive and keep one side of the rover on or close to bedrock.

Typical post-drive Navcam images were collected to inform planning for sol 950. Finally, the rover performed a Mastcam clast survey and took a MARDI image.

One element that contributed to complexity in this plan was the acquisition of a mid-drive Mastcam mosaic. This mosaic was of the initial Garden City area but was intended to be collected after the rover had driven 6 m away. The team spent significant time deciding how many images were needed for the mosaic and where they should be pointed.

Another complexity factor was that the team decided to incorporate the science activities into the rover planner sequence to simplify the plan. However, this new approach led to a hiccup due to the state of the rover after a science imaging cleanup and the state needed to commence driving. Specifically, the Rover Motor Controller Assembly (RMCA) was turned off by the cleanup and needed to be turned back on prior to driving.

Later in the planning process, it was determined that the planned Mastcam observations might be scheduled only five minutes apart. Mastcam has a "homing" step that takes at least five minutes, and if a subsequent Mastcam observation is planned before the homing completes, then it will be rejected and would cause a fault that would terminate the drive. This problem was solved by modifying the Navcam sequence to add a 90-second delay between each frame that was acquired, thereby extending the sequence to consume six minutes.

Sol 950: This plan continued the traverse through Artist's Drive valley. This was a tight planning sol and there was the possibility of a fire drill to interrupt and further constrain the tactical timeline.

Pre-drive imaging included Mastcam mosaics of exposed layers in the northwest and southeast valley walls and Mastcam atmospheric opacity imaging. As with Sol 949, there was interest in performing ChemCam observations before the drive, but these were excluded due to lack of available duration before the drive.



Figure 35: Curiosity's planned drive direction along Artist's Drive at the start of planning for Sol 950.

A drive was planned for about 37 m in a southwest direction. Figure 35 shows the rover's view in the drive direction. Given ongoing concerns about wheel wear, the team tried to avoid rocks where possible but accepted the risk associated with several small but unavoidable rocks (the potential for wheel damage was “moderate to high”). At the same time, the team wanted to avoid entirely sandy areas that could lead to slippage. The compromise was to put the right side of the rover on rockier terrain and the left side in the sandy area. The drive endpoint was anticipated to provide a good view further along the valley to aid in strategic traverse planning.

For this shift, the team faced two different challenges related to data volume. The decisional data volume, the amount of data that could be downlinked up to and including the decisional pass, was slightly low for this plan. Thus, the team had a challenge in trying to ensure the post-drive imagery needed for the next plan would fit in the available downlink volume. They opted to reduce the compression quality and reduce the number of frames acquired for some of the decisional post-drive imagery.

In addition to the decisional data volume issue, the team was also starting to get concerned regarding the amount of overall data, including non-decisional data, that was being acquired. The supratactical team was looking ahead at upcoming downlink windows and noticed the mission would be entering a period of somewhat reduced downlink volume. The team had already acquired a large amount of low-priority images from a previous engineering test and, since the rover would be driving into new terrain, they anticipated the desire to start collecting a higher volume of data to document and study the new regions. Note that the issue was not due an immediate concern regarding onboard storage, there was a healthy margin of available space for data products. Instead, the team was concerned about the delay in how long it would take to received acquired data if a large backlog of data built up in combination with the upcoming period of lower downlink volume. And, eventually, onboard storage would become an issue which would force harsher restrictions on science collections, and the team wanted to avoid this. As a result, several potential activities were descoped from the plan. These included mid-drive imaging of ripples, post-drive imaging to support ChemCam targeting and systematic documentation of terrain, and a systematic environmental observation.

Sol 951: Before driving on sol 951, the rover collected two Mastcam mosaics (again observing layers in the Mt. Tucki and the valley walls) plus the clast survey that was removed from the previous plan due to downlink constraints. The team also included targeted observations of Mount Sharp and the target “Joshua Tree,” by Mastcam and

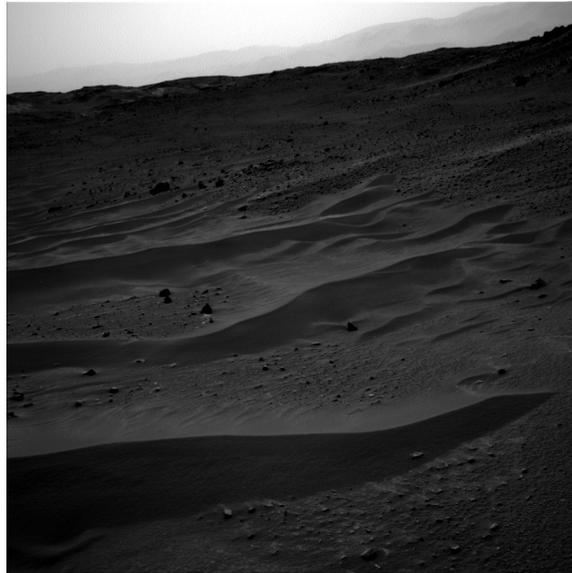


Figure 36: Ripples as seen by Curiosity on Sol 951 for planning Sol 952.

ChemCam's RMI.

Next the plan included a 18-m drive. Initially a longer drive was desired, but the data that came down after the Sol 950 drive did not have enough coverage to allow planning a drive to safely traverse a ripple field that crossed the valley ahead (see Figure 35; the ripple is near the middle of the scene). This sol's 18-m drive was designed to approach the ripples for close-up imaging to determine whether they would be traversable or an alternate route would be needed.

This drive was notable because it pushed the rover's odometer over the 10,000-m mark. The drive was followed by post-drive imaging and a Navcam zenith movie to search for clouds and measure wind direction overhead. The plan was again constrained by the available downlink, and some initially desired mid-drive images were deferred.

A JPL fire drill occurred in the morning and delayed the SOWG meeting by about 20 minutes. The MSL chat server was down for about an hour early in the shift, which the team felt significantly hindered their communication and progress. This plan had a later than usual handover time, which led to a conflict between a rover health "beep" and the uplink window.

Concerns with data volume were a challenge for the team again on this sol due to a low amount of data volume on the decisional pass. As a result the team had to opt for reduced quality and scope of post-drive imaging to try to get the data down needed to make decisions for the next plan. Even with these restrictions, they would not get all the data down they would like, but they would at least get what was necessary to support a drive in the next plan.

Sol 952: The Sol 952 plan began with pre-drive imaging by Mastcam and ChemCam RMI that covered more of the valley walls, including the "Zzyzx Spring" target. Mastcam and Navcam also observed of target "Algodones," a nearby ripple crest, to help assess grain size.

As anticipated, the post-drive data from Sol 951 enabled sufficient terrain assessment to plan a long drive on Sol 952. This drive covered almost 90 m.

Figure 36 shows the ripples in front of the rover. The rover drove to the right around the ripples and then continued on. In planning this drive, the Rover Planners once again had to balance concerns about driving over rocks and damaging the wheels versus driving in sand and possibly becoming embedded. In addition, the final location at the end of the drive was constrained by a requirement from the ChemCam instrument; on sol 953, a ChemCam focus test was planned and this required rocks visible within 3 m of the rover. Determining the

final drive and ending location required significant coordination between the Rover Planners and the ChemCam representatives.

The plan was constrained by the amount of time available. A mid-drive DAN observation was descope to increase the time available for driving by 11 minutes. The Rover Planners reduced the amount of visual odometry used near the end of the drive to save 20 minutes. Several post-drive imaging observations were moved to a later imaging block to allow more time for the drive. In addition, a SAPP sun observation was added to the drive which consumed 10 minutes of duration. This observation is required periodically to update the rover's knowledge of its attitude.

Post-drive activities included MAHLI stowed imaging and a Navcam mosaic for planning the ChemCam focus test as well as the next drive. After the decisional pass, additional data was collected including a Mastcam clast survey, Mastcam workspace survey, Mastcam 360, and Navcam upper tier observation.

Once again, the team was challenged by a relatively limited decisional downlink volume. As with Sol 951, they decided to reduce the quality and scope of post-drive imaging. Also of special importance was collecting enough post-driver imagery to support the upcoming ChemCam focus test. The team adjusted the priorities of some of the images to increase the likelihood of receiving the data needed to plan this test.

Sols 953, 954: This two-sol plan was created on a Saturday shift. Most of the planning shifts were by this point scheduled for weekdays, with Fridays devoted to multi-sol plans that would cover weekend activities. However, some Saturdays were employed to increase the number of nominal planning sols (those for which data was available from the immediately preceding sol). While Saturday plans increased mission science productivity, they also put a strain on the team.

The Sol 952 drive completed successfully. The downlink passes over-performed their data volume predictions, so the team had the data needed to select the targets for ChemCam and the sample dump activity.

The primary goal for Sol 953 was to do a ChemCam focus test. This involves collecting LIBS spectra and RMI images of the same target ("Eaton Canyon") three times throughout the day and assessing the impact of temperature on instrument performance. LIBS is a laser-induced breakdown spectrometer that uses a laser to ablate a small amount of material from a distant target (up to 7 m away), then record the emission spectra produced by the plasma. RMI is the Remote Micro Imager which collects visible images of LIBS (or other) targets. Since LIBS targets range from next to the rover to 7 m away, it is necessary to properly focus the laser for the true target distance. In November 2014, it was discovered that the auto-focus system had failed. For some time, the team proceeded by commanding multiple observations at different focal distances for each target, but this multiplied the time required to observe each target. In tandem, new software was being developed to automatically take nine RMI images, select the best, and use that to specify the focal distance to LIBS. In support of this effort, the focus test was important because ideally the focus of the instrument should not vary with temperature. This test enabled the team to assess focal performance empirically.

Selecting the times for the ChemCam observations required negotiation amongst the team. Initially, the first proposed time (10:30) would have required heating the mast (for articulation) and then cooling the ChemCam instrument, which increases complexity and power consumption. After discussion, a time of 10:50 was selected, which required shifting the second observation later as well (from 13:30 to 13:50). The final observation was scheduled for 17:00. This discussion was held on the supratactical side so it did not hold up tactical planning.

In addition, the team included a large Mastcam mosaic of target "Mount Saint Mary" to further characterize the stratigraphy of the valley walls.

On Sol 954, the rover analyzed a drill sample from target "Telegraph Peak," using a strategy similar to that described above for Sol 781's analysis of a sample from Confidence Hills. Part of the sample was transferred to the SAM instrument. The rest of the sample was dumped onto the surface for later analysis by APXS. Mastcam was used to image the SAM instrument before and after the sample was delivered as well as the dumped pile.

The sequence to command the activities necessary for the sample dump and analysis was created on the strategic timeline. The tactical team focused on selecting an appropriate location to dump the sample and performing a rover stability assessment that was required prior to using the arm. After the sample was dumped, MAHLI was used to image the pile and the APXS instrument was placed onto the pile to integrate overnight and record compositional information.

Sequencing the MAHLI observations of the dump pile was challenging since it was done “in the blind” (i.e., the team had to plan it before seeing what the dump pile looked like and, in particular, without knowledge of its height). In addition, it was discovered late in the day that the brush would be casting a shadow at the time MAHLI would be observing the dump pile. This was resolved by adding a 15-minute delay in the RP sequence. The MAHLI activity was also reduced from the full requests form science. The MAHLI PUL described the day as “semi-chaotic.”

There was interest in an additional MAHLI image before dumping the sample, however it was descope from the plan due to concerns about added complexity and limited time in the plan.

The plan also included a Navcam zenith movie. There was interest from the science team in acquiring an additional environmental observation but it was descope due to concerns regarding complexity in ensuring the safety of the activity.

Sol 955: This sol was devoted to remote sensing and contact science activities. The original supratactical plan included a drive, but it was noted that this would be strongly constrained by the amount of time available before the decisional pass. In addition, inspection of the dump pile from Sol 954 revealed that it was very thinly spread out, so the APXS data might not be usable (i.e., it might not be possible to determine which elements were from the pile versus the underlying bedrock). Therefore, the drive was deferred to Sol 956.

The Sol 955 plan included Mastcam and ChemCam RMI images of target “Eaton Canyon,” plus further Mastcam and passive ChemCam imaging of the sample dump pile that was deposited on Sol 954. Passive ChemCam observations are those in which spectra are collected without employing the LIBS (ablating laser).

After this imaging, the arm was put back down on a bedrock target near the dump pile named “La Brea.” Comparing the compositional data between the dump pile and the rock on which it rested would allow the determination of which chemical elements were unique to the “Telegraph Peak” target.

Halfway through APAM, it was discovered that the rover arm would be in the way of Mastcam and ChemCam imaging of the dump pile. At the beginning of the plan, the APXS instrument was still placed in contact with the dump pile, so it was not visible to the mast instruments. To accomplish the dump pile imaging, the arm had to be first raised out of the way. Once the imaging finished, the rover used MAHLI to image “La Brea” and placed the APXS on it. The imaging activity in the middle required that the arm activities be broken into two blocks, which increased the plan complexity.

In preparation for planning the next drive, the rover collected a Mastcam mosaic in the anticipated drive direction (southwest) as well as a large mosaic of target “Tucki Mountain.”

Activities in the morning of Sol 956 included atmospheric observations with Navcam and Mastcam.

Sol 956: The additional APXS observation of the bedrock near the dump pile was received, so the rover was ready to drive on to its next site.

Unfortunately, the telemetry received from the rover indicated there was a problem with the activity to load a new schedule table to the REMS instrument. There is a known idiosyncrasy with the way the RCE communicates with the REMS instrument that can sometimes cause an error in loading a new schedule table to the instrument. When this happens, due to the way the schedule table upload was sequenced, it puts the instrument into a state such that it will be safed by flight software later in the sol. It is straightforward for the team to recover the instrument at this point and resume its nominal operation. However, at the start of Sol 956 planning, while the team had data showing the schedule table upload failed, the instrument had not yet been safed. This would happen later in the day. The team did not want to send a recovery sequence on Sol 956 without first observing the instrument being safed and confirming it was for the expected reason. As such, there was no REMS activity in the Sol 956 plan.

The plan began with observations of nearby boulders (dubbed “Waucoba” with Mastcam and the collection of Navcam images to complete the 360-degree panorama. A ChemCam calibration target observation was originally planned as a pre-drive activity but later moved to a post-drive activity to ensure there was sufficient time to complete the drive before the decisional downlink pass.

Next, the plan included a drive of 65 meters. There were two possible paths for this drive in the Strategic Traverse Route. The team elected to take the southwest route deeper into the valley (see Figure 37). [KW: I

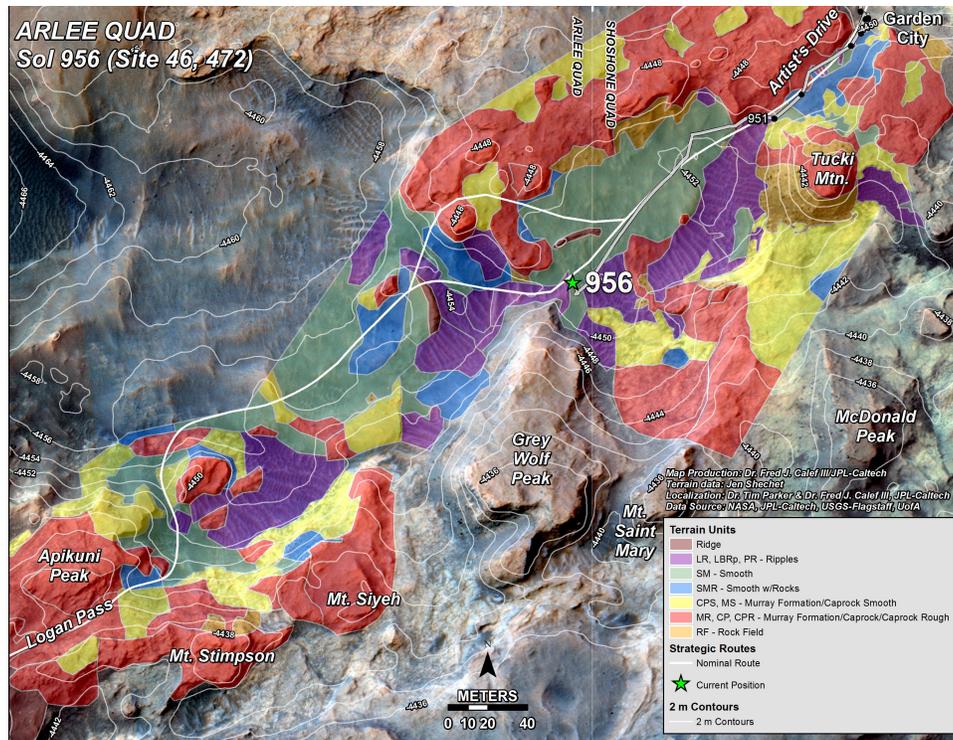


Figure 37: Southwest traverse route chosen at decision point on Sol 956.

could not find any rationale stated for this decision - it is simply attributed to “Science.”] One complication was that the end-of-drive position would not provide sufficient visibility to plan future drives along the route, so mid-drive imaging was employed from a better vantage point. This rendered the mid-drive imaging as important as post-drive imaging usually is.

Planning was constrained by a small downlink volume predicted for the decisional MRO pass, so post-drive imaging was slimmed down. Image priorities were set to ensure that the mid-drive images would come down in the decisional pass, and the remaining images would come down on the next pass, which would be a bit past the start of Sol 957 planning but still usable.

Originally, the plan was to stow the arm (from its position with the APXS on the dump pile), but since it would not be in the way of the pre-drive imaging activities, it was decided for simplicity to combine the arm stow with the mobility block. This also reduced the amount of heating needed.

After the drive, there was a science block dedicated to an observation of Mercury transiting the sun. The rover collected a Mastcam mosaic at sunset. This required precise timing and pointing. The team drew on their previous experience with a Phobos transit to plan this activity. They included multiple images with different exposure settings to maximize the chance of getting good-quality images of the transit.

Sol 957: The primary goal of this plan was to continue driving. Pre-drive imaging included a Mastcam mosaic of the stratigraphy of Mount Sharp and along the planned traverse route as well as ChemCam RMI imaging of target “Lewis”, which exhibits steeply inclined bedding. There was also a Mastcam observation of target “Placid” on the ripple nearest the rover to assess surface grain size (see Figure 38). Finally, the plan included a Navcam movie of Mount Sharp, looking for cloud activity, and another movie looking for dust devils.

The drive was planned to cover 63 m, heading to the west, through a narrow gap north of Grey Wolf Mountain (see Figure 38). Initially, the end of the drive was planned to be positioned with a good view of Logan Pass to the south. However, it was determined that in the time available for driving, the rover would be 20 m short of that location, with a less favorable view for future plans.

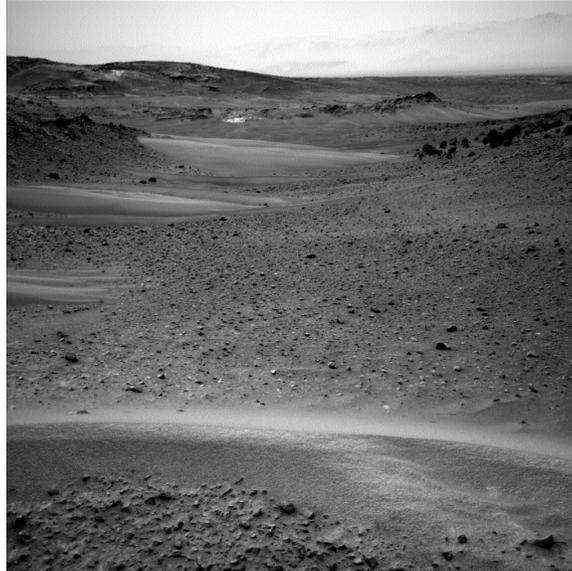


Figure 38: Drive direction ahead for Curiosity for planning Sol 957.

This plan was constrained by the downlink volume available. Post-drive imaging was reduced in quality and volume. The team also prioritized ChemCam calibration data below Mastcam and other imaging data.

Starting with this plan, the team benefited from an updated power model that ran in about one quarter of the time required by the previous version. This kind of efficiency improvement is always welcome and sometimes dictates how many iterations of plan refinements can be done during the tactical planning process. This affects how many alternatives can be considered and/or how optimal the plan will be with respect to resource consumption.

Sol 958: This was another driving sol. The Sol 957 drive concluded with the rover in a position with limited visibility due to a nearby ridge.

Pre-drive imaging included a passive ChemCam sky observation (for atmospheric composition), a Mastcam mosaic of the ridge to the east, and a Mastcam observation of another ripple (target “Libby”) for grain size analysis. The plan also included Mastcam and MAHLI observations of the right wheels to support ongoing monitoring of the wheel health.

In this case, there was plenty of time for a longer drive, but the plan was constrained by limited terrain visibility. The planned drive was for 20 m to approach the ridge and ideally position the rover for a long drive in the next plan.

After the drive, the plan included a ChemCam observation of a calibration target and a Mastcam clast survey.

Sols 959, 960, 961: This three-sol plan was created on a Friday shift. Sol 961 had a limited downlink prediction, so imaging was prioritized for Sol 959, and the drive was planned for the second sol, Sol 960.

The Sol 959 plan included Mastcam images and ChemCam rasters on two targets. Target “Gold” was selected because it appeared to be laminated and similar to Pahrump targets, and target “Espinoza” appeared to be tilted and layered bedrock. Later in the day, the plan included Mastcam mosaics of Logan Pass (the goal of this campaign), a small outcrop near the rover with target “Daughter of the Sun,” and the ridge to the west.

The drive was planned for Sol 960. The Sol 958 drive successfully placed the rover with a view that enabled a much longer drive (104 m). In fact, the view was so good that the current location was used as waypoint 3 for science imaging along the Strategic Traverse Route, thereby precluding a need to stop at the originally planned waypoint 3. The drive path deviated to the north from the Strategic Traverse Route (see Figure 39) to improve visibility for future traverse planning. There was a small ridge to the southwest that would occlude terrain further in that direction; the chosen path navigated north and around this feature.

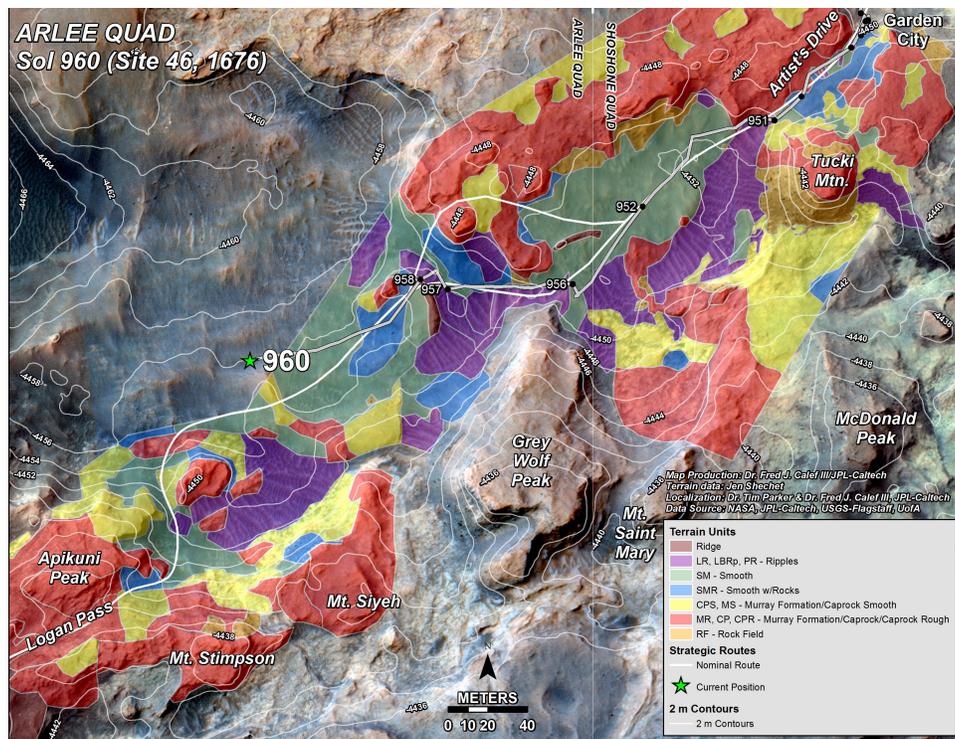


Figure 39: Long drive planned on Sol 960, deviating north of the Strategic Traverse Route.

Sol 961 included ChemCam calibration target observations, a Mastcam atmospheric observation, and Navcam movies of Mount Sharp (clouds) and to the north (dust devils).

Sol 962: The primary objective of this plan was to acquire MAHLI images of the wheels in different positions as part of ongoing monitoring of the rover's wheel wear. This activity is typically done after every 500 meters of driving. Although it had only been about 450 meters since the last wheel images were acquired, the team expected the next drive would take the rover into more challenging terrain and so it made sense to take the images a little earlier before moving into that different terrain.

The wheel imaging activity takes a long time and the time available in this plan prior to decisional downlink communications was limited. The team chose to only do imaging at four of the five desired positions in this plan, leaving the fifth position images to be acquired in the following plan. Doing so would allow a little time for some science observations in this plan.

Science observations prior to the wheel imaging activity included Mastcam atmospheric observations of the sky and the sun, additional Mastcam imaging of the backside of the outcrop "Daughter of the Sun" (previously observed on Sol 959), and a Mastcam mosaic of the Logan Pass region (Viewpoint #4 of the Logan Pass imaging campaign). While they wanted the Logan Pass mosaic to be as large as possible, the number of image frames for that mosaic was chosen (and limited) to fit within the amount of time remaining after accounting for the expected durations of the other science activities.

Since the wheel imaging activity required moving the rover, additional post-drive imaging was required to enable precise pointing at ChemCam science targets in the following plan. This imaging was kept to a bare minimum since the amount of data expected to be received in time for the next planning day was limited. Fortunately, because the wheel imaging activity does not move the rover very far, the team would be able to make use of previously collected data to plan the next drive.

After the downlink communications passes, the plan included a MARDI image and a new NavCam atmospheric observation of the sun just before it sets.



Figure 40: The Mt. Shields area, which appears to a valley carved in bedrock and later filled with sedimentary material. This area is the target of the diversion away from Logan's Pass called "Logan's Run." (NASA/JPL-Caltech/MSSS).

Sol 963: Limited post-drive imaging from the Sol 962 plan did not provide a lot of options for targeting ChemCam in this plan before driving away. A single ChemCam activity was planned for the float rock target, "Apple," with accompanying Mastcam imaging.

By the time Sol 963 was being planned, the scientists had begun discussing the possibility of a detour to investigate some "lens features" on the valley wall northwest of Logan Pass in the direction of Mt. Shields (shown in Figure 40). Specifically, the geologists were intrigued by the fact that the rock layers did not run straight through the hill, instead they formed lenses of resistant rock interrupted by other material. The pattern suggested the area could be an incised valley fill, in which a valley has been cut into bedrock and then filled in with other sediment [Jet Propulsion Laboratory Press Release(2015j)]. The detour to explore this area more closely became known as "Logan's Run."

The drive planned for Sol 963 was intended to get to a position where they could image both the entrance to Logan Pass and the path along "Logan's Run." The drive began with acquiring the fifth set of MAHLI wheel images which didn't fit in the Sol 962 plan. The drive was limited by terrain visibility, as small ripples blocked the view of some areas along the drive path. A drive of 47 meters was planned, with the first part of the drive over visible terrain and the later part of the drive commanded in guarded mode. Guarded mode driving uses a subset of autonomous navigation in which the rover takes images along the commanded drive path and stops if it detects any hazards such as large rocks or steep slopes.

All of the standard post-drive imaging, including a 360-degree NavCam panorama, and a DAN active observation were planned. A NavCam zenith movie to look for clouds and wind direction as well as a MARDI image were also planned. Late night Mastcam imaging of the moons Phobos and Deimos were planned to investigate the opacity of the upper atmosphere and aerosol size.

Sols 964, 965: Although the previous plan's drive was expected to achieve 47 meters, the rover only accomplished 17 meters of that drive and would need to make up the rest on Sol 964. The Sol 963 drive was stopped by the autonomous navigation almost immediately after guarded mode driving began due to a large rock which was detected a little too close to the rover's right side, as shown in Figure 41.

Before resuming driving, the team decided to acquire two ChemCam observations of the "Blackrock" outcrop along with accompanying Mastcam imaging. A drive then was planned to cover the 32 meters now visible in the latest images, but beyond that the terrain was still occluded. The drive was followed by a SAPP sun observation to update the vehicle attitude knowledge, along with the usual set of post-drive imaging and DAN active observations. The plan also included a 10x2 Mastcam mosaic of Logan's Run and a MARDI image.

Sol 965 began with a direct to Earth communication window in which time correlation data would be exchanged. The afternoon's untargeted remote science consisted of a newly improved passive ChemCam atmospheric observation and related calibration activities. The team also considered including additional routine ChemCam



Figure 41: This NavCam image taken after the Sol 963 drive shows the nearby large rock that autonomous navigation deemed a hazard and stopped the drive early (NASA/JPL-Caltech).

observations of the calibration targets, but those were deferred to a later plan due to the ChemCam team's concern about the complexity of the new ChemCam atmospheric observations. An overnight SAM atmospheric observation was also included in the plan.

Sol 966, 967, 968: With a good viewshed from the previous drive, a long 90 meter drive was planned for the three-sol weekend plan. Before driving away from the current location, sol 966 was devoted to remote sensing and imaging of two nearby rock targets, "Pauline" and "Pablo," with ChemCam and Mastcam. Interestingly, both of these targets were difficult to classify from NavCam images. Pauline appeared to be a possible conglomerate, and Pablo appeared to be a layered rock, but both were revealed to be relatively homogeneous caprock material. The latter part of sol 966 was used for atmospheric measurements with Mastcam and NavCam.

Figure 42 shows an outline of the long weekend drive that took place on sol 967 and put the rover about 10 meters from an outcrop of interest in front of the Mt. Shields area. A test of the Visual Target Tracking (VTT) system was run during this drive. VTT works by selecting a target such as an outcrop in a navigation image, then tracking this target as the rover drives towards it. If the rover begins to drift away from this intended target during the approach, the system can correct the rover's trajectory and put it back on course. For the test on sol 967, the system was only run to demonstrate that it could successfully track the target; it did not feed back into the rover navigation. Standard post-drive imaging was conducted, including a MAHLI stowed image and some additional NavCam images of the valley wall of the Mt. Shields area.

During the night of sol 967, CheMin dumped the Mojave and Telegraph Peak drill samples. An analysis of the emptied Telegraph Peak sample cell was performed to provide calibration for subsequent readings. However, the analysis revealed that the sample cell had not completely emptied, and would have to be re-dumped later.

Sol 968 was used for calibration and untargeted science activities including ChemCam calibration and NavCam atmospheric observations.

Sol 969, 970: Developing this two-sol plan was hindered by an MSLICE bug that prevented the post-drive NavCam mosaic from sol 967 from appearing. Without this, scientists were unable to select targets or begin planning subsequent drives. The bug was quickly fixed, but not before the SOWG meeting was nearly over.

There were two possible strategies for this plan, and much discussion about which to pursue. The first strategy was to acquire detailed images and remote sensing measurements of the outcrop in front of the rover at Mt. Shields (Figure 42, left). A second strategy was to plan a short bump drive to perform contact science at

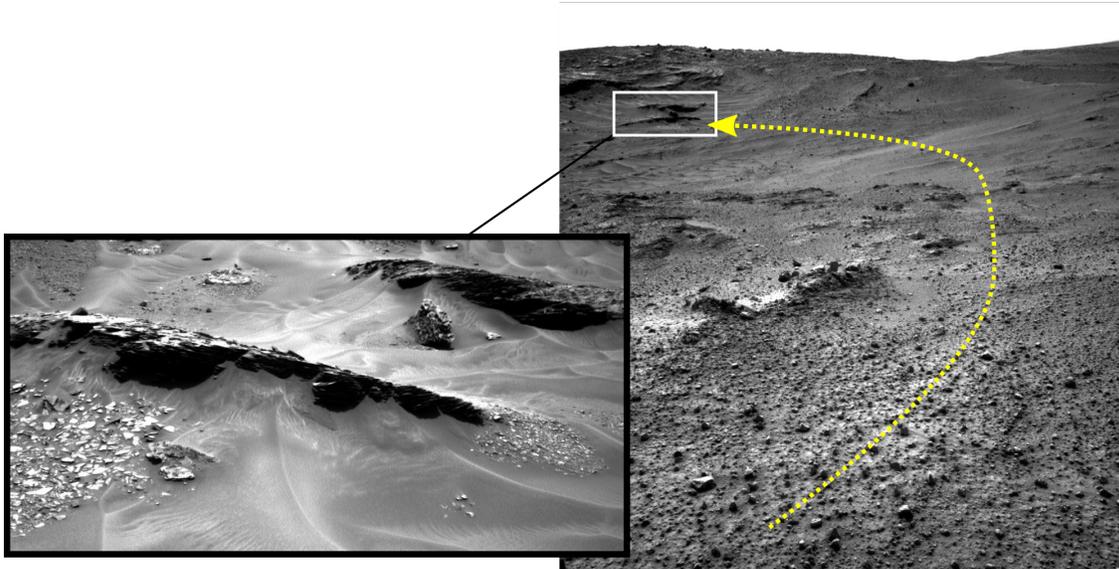


Figure 42: A sketch of the sol 967 drive, approaching an outcrop at the lower part of Mt. Shields (NASA/JPL-Caltech).

this outcrop, without waiting to use the imaging to select a specific location. Eventually, it was decided to only perform detailed imaging, given the extra complexity introduced by the MSLICE bug.

The imaging included a large Mastcam stereo mosaic of the Mt. Shields area, in addition to four ChemCam RMI images of the outcrop in front of the rover to determine its grain sizes in the rocks, which would reveal information about its sedimentary history. Finally, a ChemCam LIBS raster was taken of a bedrock target “Hungry Horse” that appeared similar to the bedrock seen at Pahrump hills. This observation would help inform the geological connection between the Pahrump and Mt. Shields areas.

In the evening of sol 970, Mastcam observed the eclipse of Phobos as it passed into shadow behind Mars. Since the light hitting Phobos just before eclipse passes through the limb of Mars’ atmosphere, such observations are useful for understanding atmospheric composition.

Sol 971, 972: During this two-sol plan, the primary goal was to approach the nearest outcrop in front of the rover (Figure 42, left) with a short 6 meter bump in order to perform some contact science. Before the bump, there were Mastcam and ChemCam LIBS observations of “Helena,” a dark float rock that appeared to be from the upper part of the Mt. Shields stratigraphy, and ChemCam RMI observations of “Lolo,” a block of layered sedimentary material. There were also some Mastcam atmospheric observations.

Prior to the drive, there was some MAHLI imaging to assess wheel health, then standard post-drive imaging afterward. Since the drive was so short, there was no MAHLI stowed image.

During the night of sol 971, there was CheMin analysis of the empty cell that once held the Mojave sample to provide a baseline and ensure that all of the previous sample had been removed.

On sol 972, there were several engineering activities, including ChemCam calibration and preparation of the SAM instrument for a subsequent experiment. The SAM preparation included uploading a new data table and cleaning the instrument. A twilight MARDI image was taken on the evening of sol 972.

5.1.3 Outcome of Campaign

Figure 43 shows the path the rover followed during the course of the Artist’s Drive campaign, including the detour to Logan’s Run. In order to manage the scope of our case study, we decided to end the Artist’s Drive campaign at Sol 972. However, the actual Artist’s Drive activities for the mission continued past this point. After the bump on Sol 971, the team spent a weekend at Logan’s Run performing contact science. On Sol 976 they resumed the drive toward Logan Pass.

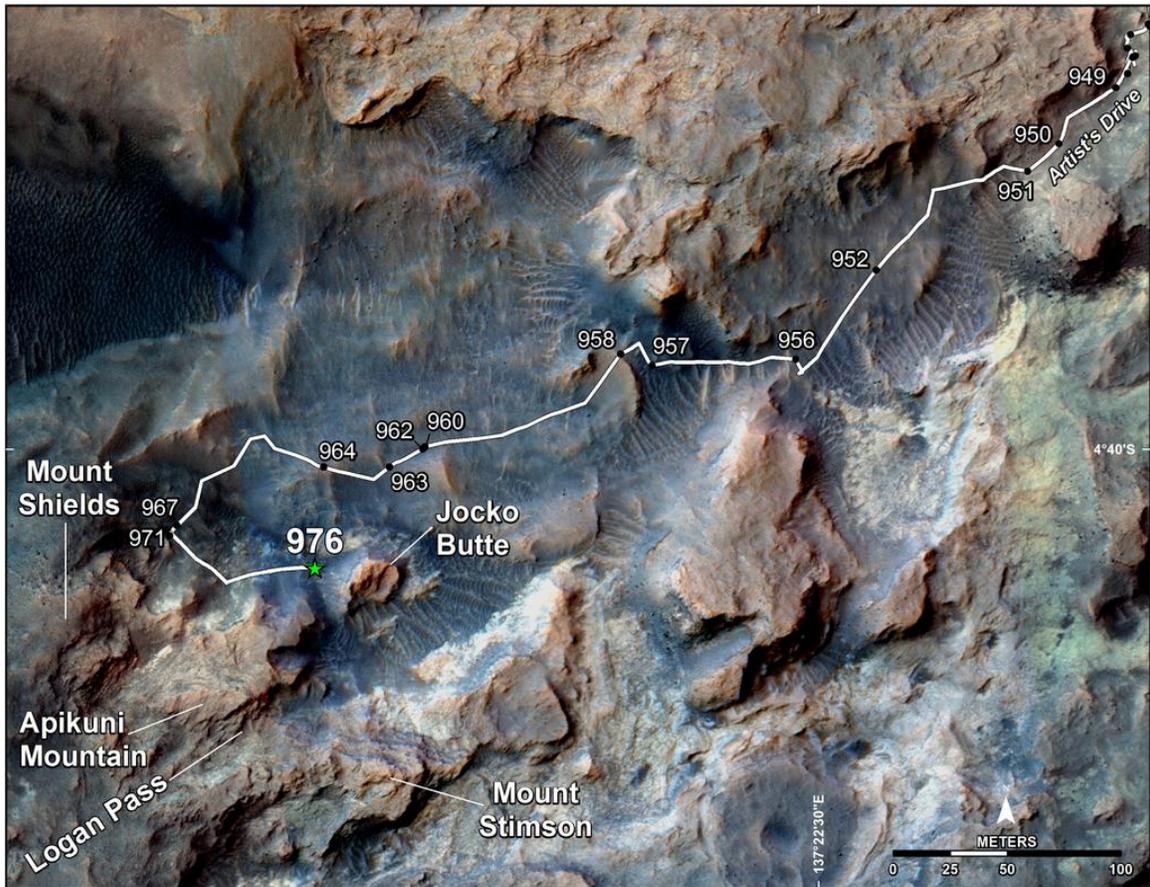


Figure 43: Traverse path during the Artist's Drive campaign (NASA/JPL-Caltech/Univ. of Arizona).

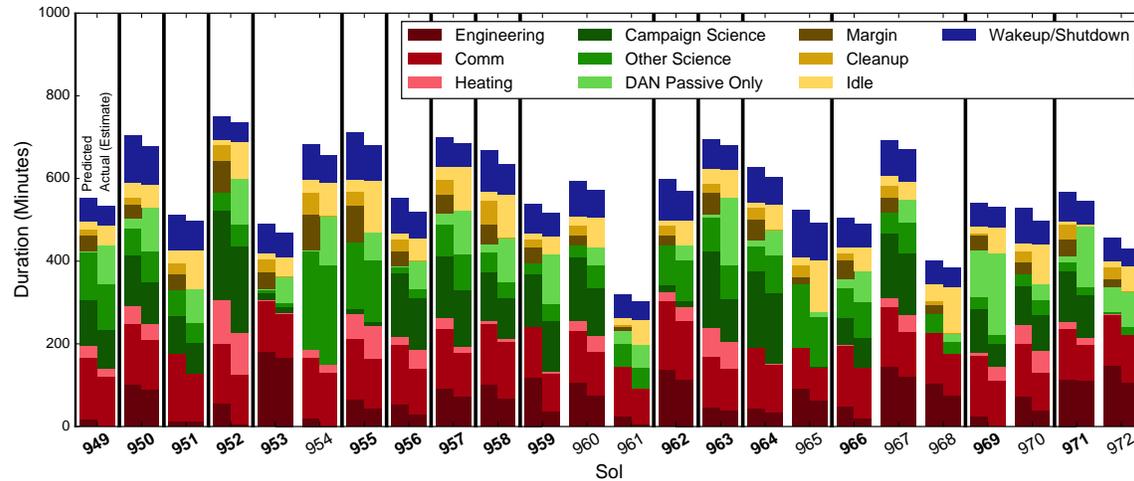


Figure 44: Allocation of RCE duration for Artist's Drive campaign.

Unfortunately, the team was not able to reach the Stimson / Murray formation contact at Logan Pass. In order to reach the area, the rover needed to ascend the steepest slopes it had yet seen on Mars, with slopes as steep as 21° . The team made multiple attempts to climb the slopes but due to excessive slip encountered on the drives, the team decided to abandon the attempt to reach Logan Pass. While some slippage was expected, the team did not expect the amount that was experienced. While they knew that the terrain with sand ripples posed significant traversability challenges, there appeared to be rockier terrain on the side. However, this terrain resulted in significant slip, as well [Jet Propulsion Laboratory Press Release(2015i)].

Fortunately, HiRISE imagery showed another Stimson / Murray formation contact at a nearby location named Marias Pass. This alternate location is shown in Figure 33.

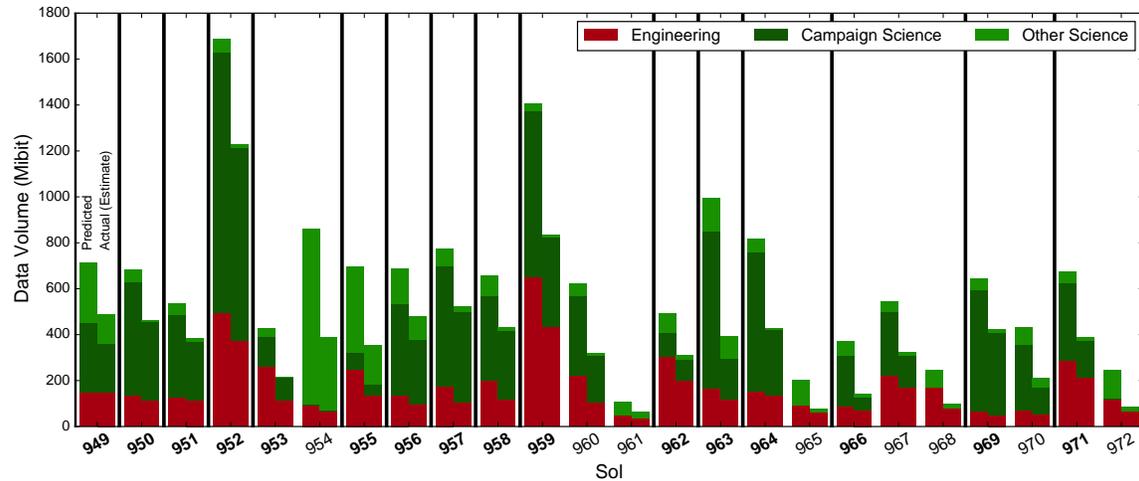
5.2 Analysis of Resource Usage

We applied the same set of resource analyses that we performed with the Pahrump Hills Walkabout Pass 1 campaign. Please see Section 4.2 for an overview of the plots.

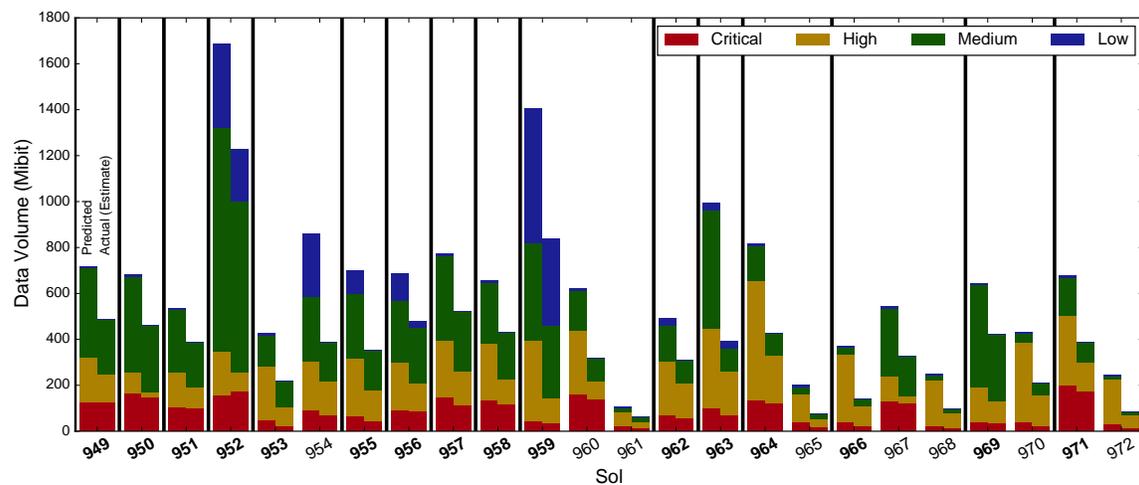
We begin with the allocation of RCE duration shown in Figure 44. One obvious difference between Figure 44 and the similar plot for Pahrump Hills in Figure 22 is that Artist's Drive has several more single-sol plans than Pahrump Hills. This is because the Artist's Drive campaign happened to start soon after the end of a restricted period, which had ended on Sol 947. As a result, the team did not encounter a restricted sol during the Artist's Drive campaign until Sol 964.

The result of the larger number of non-restricted sols can be seen in Figure 44 as there is a higher volume of RCE duration allocations compared to the Pahrump Hills campaign. However, the Artist's Drive campaign RCE allocation plot does show a similar pattern as seen in the Pahrump Hills campaign in that the overall activity during a multi-sol plan is less than the overall activity for a corresponding number of single-sol plans. Similarly, there is a significant decrease in campaign-related activity on sols following drives during the restricted period, e.g. Sols 965, 968 and 972. Note that the two-sol plan for Sols 969 and 970 were used for targeted remote sensing without a drive, so the team was able to perform campaign-related activity on both sols.

Figure 44 shows that the over-prediction in activity duration and the need for margin and cleanup duration resulted in a large amount of non-productive duration in the actual RCE duration allocations. In particular, time spent with the RCE Idle or performing DAN Passive Only could have been spent doing additional activity. The sol activity summary from above indicated several instances where activities were descoped or limited in scope due to concerns about available duration. For example, pre-drive activities were descoped on Sols 949 and 950, a mid-drive activity and post-drive imaging were descoped in Sol 952 and a MAHLI activity was descoped in sol 953 partly due to time constraints. In addition, the drive on sol 949 was limited in how much visual odometry was performed due to concerns about duration. However, Figure 44 indicates that there may have been time for these activities.



(a) Allocation of data volume by intent category



(b) Allocation of data volume by downlink priority

Figure 45: Allocation of data volume for Artist's Drive campaign.

Figure 45 shows how the acquisition of data was allocated with Figure 45 (a) showing data volume allocations to intent categories and Figure 45 (b) showing allocations to priority bins. As with the Pahrump Hills versions of these plots (Figure 23), the largest predict vs. actual discrepancies are with non-critical priority bins. The largest discrepancy was on Sol 959 largely due to data collected in high and low priority data related to Engineering and Campaign Science activities.

Recall from the discussion of Sol 950 above that the team began to have concerns about overall data volume acquisition and descope several activities. Figure 45 shows that actual data volume was quite a bit lower than predictions, meaning the team could have potentially decided to descope fewer activities and achieved the predicted result.

The team also faced challenges on several sols due to limited data volume on decisional passes: Sols 950, 951, 952, 956, 957, 962. On these sols, the team restricted the amount of post-drive imaging that was acquired to better ensure decisional data fit within available passes. Figure 45 (b) indicates that the predictions for decisional (aka critical) data was fairly well predicted. Of these sols, Sol 957 had the largest discrepancy between predicted and actual critical data volume with an over-modeling of 38 Mbits. That is significant in terms of the potential increase in images and image quality. To fully assess how well predictions supported the teams decisions on these sols, one must also consider

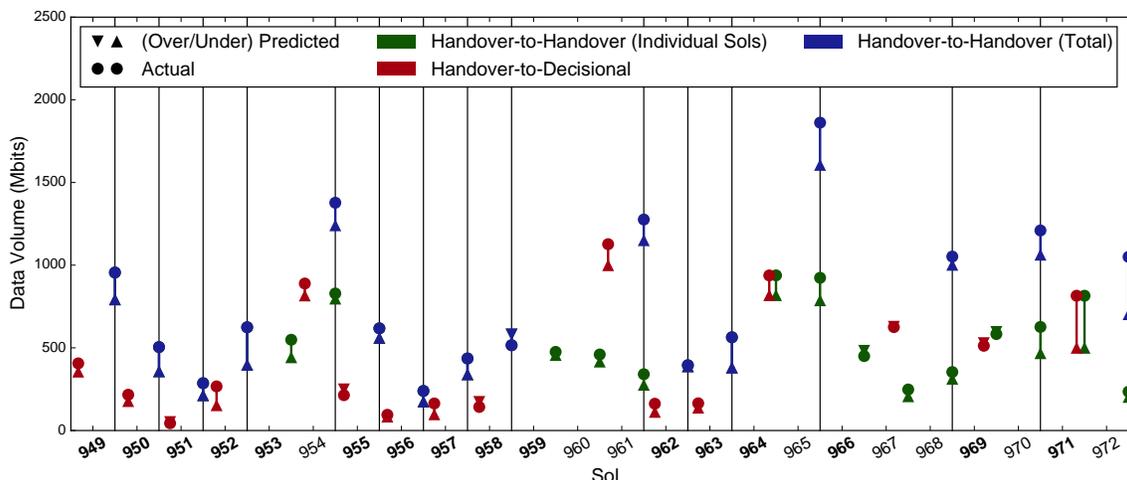


Figure 46: Downlink performance for Artist’s Drive campaign.

Sol	Acquired Critical (Mb)			Downlinked (Mb)		
	Predicted (Mb)	Actual (Mb)	Predict - Actual (Mb)	Predicted (Mb)	Actual (Mb)	Actual - Predict (Mb)
950	156	172	16	178	217	38
951	103	112	10	51	44	-7
952	182	164	-17	153	267	114
956	90	96	6	85	95	10
957	119	157	38	98	164	65
962	62	74	12	113	162	49

Table 2: Predicted vs. actual critical data volumes for decisional limited sos during Artist’s Drive campaign.

downlink performance. We look at this next.

Figure 46 shows the predicted and actual performance of the orbiter relay communication windows in terms of number of Mebibits transferred. As with Pahrump Hills, the actual downlink performance generally matched predicted performance. The biggest discrepancy was on Sol 971 due to multiple MRO windows performing better than expected.

We return now to the question from above regarding how well predicted data volume acquisition and downlink performance supported the team’s decisions to restrict post-drive imaging on sols in which decisional data volume was especially constrained. Table 2 summarizes the predicted vs. actual values for critical data acquisition and decisional data downlink from these sols. In most cases, the predicted acquisition over-modeled how much data would be acquired and the predicted downlink performance under-modeled how much data would be downlinked. This indicates the team could often have acquired more post-drive images and/or at higher quality. It is understandable the team would be conservative on these decisions, as the consequence of not getting the decisional data could result in a significant reduction in productivity in the next plan.

Figure 68 shows the predicted vs. actual warmup heating duration in the plans. The overall amount of warmup heating for this campaign was significantly higher than for Pahrump Hills. Pahrump Hills had a total of 11 hours of predicted and 4 hours of actual warmup heating, whereas Artist’s Drive had a total of 30 hours of predicted and 9 hours of actual warmup heating. As Figure 6 shows, The Pahrump Hills campaign was conducted during the warmest part of the Martian year. The Artist’s Drive campaign was conducted at a slightly cooler time of year and thus required more heating.

As with the Pahrump Hills campaign, Figure 47 indicates there was often large amounts of over-modeling of required pre-heat. In general, the over-modeling is proportional to the total amount of pre-heating. Significant over-modeling examples were on Sols 952, 955 and 963. These sols all had activity either early in the morning or late at night, which accounts for the higher amount of heating in the plans.

Figure 48 shows how vehicle energy was allocated during the campaign. As with the Pahrump Hills, the shape of the plot roughly matches the corresponding RCE duration allocation plot from Figure 44. The significant deviations between the plots are Sols 965 and 972, in which the RCE durations for these sols are relatively low while the energy allocations are relatively high. Both of these sols included overnight SAM activity, which consumes significant energy

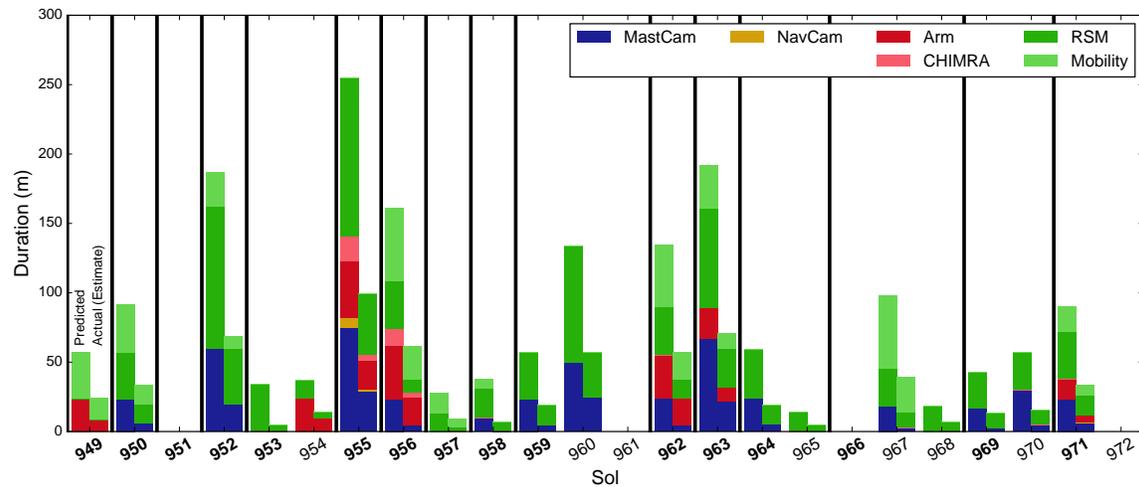


Figure 47: Summary of device warmup activity for Artist's Drive campaign.

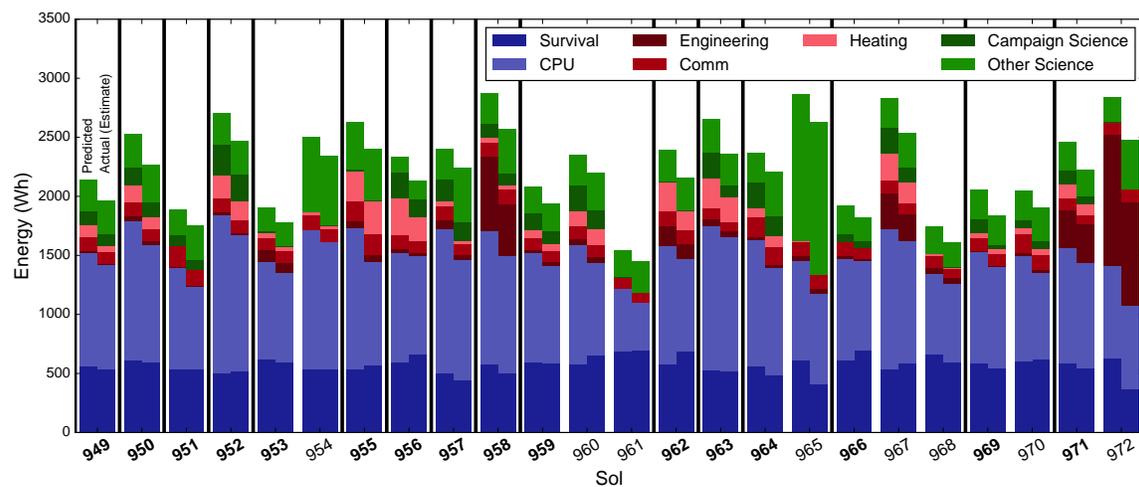


Figure 48: Allocation of energy for Artist's Drive campaign.

and can run while the rover is asleep. As discussed in the Pahrump Hills campaign, the differences between predicted and actual energy shown in Figure 48 represents a significant amount of additional activity that could have been performed. We will have more to say about this below.

In Figure 49, we look at the predicted and actual state of charge for the rover's battery during the campaign. The plot indicates that energy was generally not a limiter for the campaign activity. For most of the campaign, the minimum state of charge was well above the 40% minimum limit and the handover state of charge values were well above the supratactical guidelines. Though temperatures were slightly cooler during these sols than during the Pahrump Hills campaign, it was still a relatively warm part of the Martian year. It is mainly in the colder winter months when energy becomes more of a limiter due to the larger amounts of required heating.

The team did get close to the 40% battery state-of-charge limit in the Sol 949 plan. In addition, the team did not quite meet the supratactical handover guideline for the next sol, though it was within 1% of the guideline. Though energy was tighter on this sol than others, it did not cause the tactical team to descope activity. It is also important to note that all of the predictions were conservative. For example, although the team predicted a minimum state of charge of 44% on Sol 949, the actual state-of-charge was 63%. And the actual handover state-of-charge was 91%. Similar behavior is observed on the other sols. As with Pahrump Hills, the sols in which predicted and actual state-of-charge

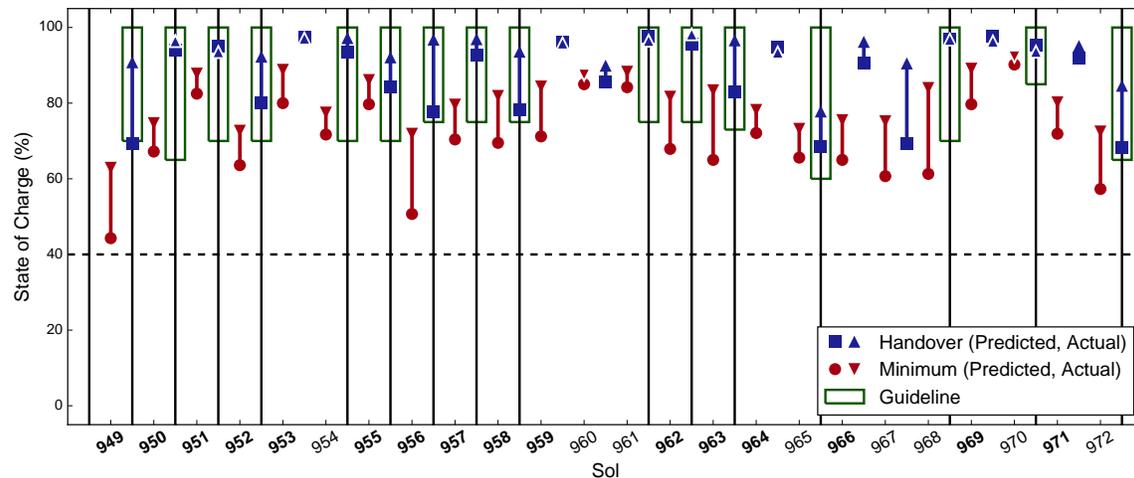


Figure 49: Summary of battery state of charge for Artist's Drive campaign.

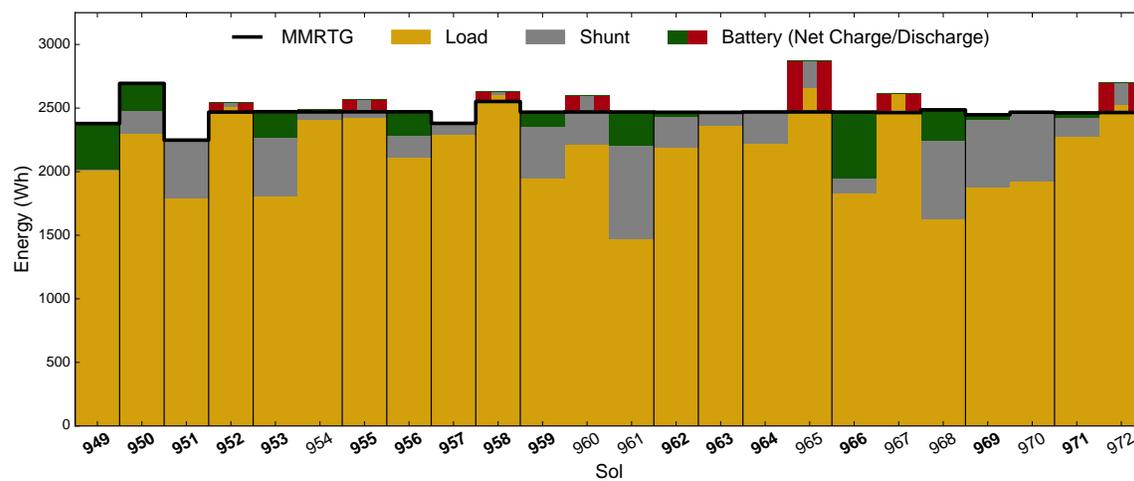


Figure 50: Summary of battery and MMRTG energy for Artist's Drive campaign.

are close are those in which the battery is at a high state-of-charge. This is consistent with the explanation of the challenges of modeling battery state-of-charge from Section 3.3. Points where both the modeled battery and actual battery are near 100% stop the accumulation of error in the prediction.

Figure 50 shows the how energy from the MMRTG was used during the campaign. The plot is consistent with Figure 49, in that energy was largely not a constraining factor for the campaign. Compared to Pahrump Hills, the Artist's drive had more sols in which there was a net drain of the battery, though most cases were relatively small. The largest net drain was on Sol 965. As mentioned in the discussion of energy allocation from Figure 48, this sol included a high-energy SAM activity. Thus, it is not surprising to see a depletion of the battery.

The shunt energy from Figure 50 and the non-productive RCE duration from Figure 44 represent energy and time that could have been used for other activities, including campaign-related activities. As we did with Pahrump Hills, we generated the plot in Figure 51, which estimates the amount of additional remote sensing or driving that could have been performed with these resources. It estimates that there was an additional 62 hours of available activity duration. Figure 52 confirms there was sufficient unused daylight duration to make use of this time.

Figure 74 shows the tactical timeline durations for the Marias Pass campaign. During this time of the mission, the nominal tactical duration was 9.5 hours. There were only two shifts that exceeded this duration, and they only exceeded the duration by about a half our each. The two shifts that ran long were for Sols 949 and 952. The reason

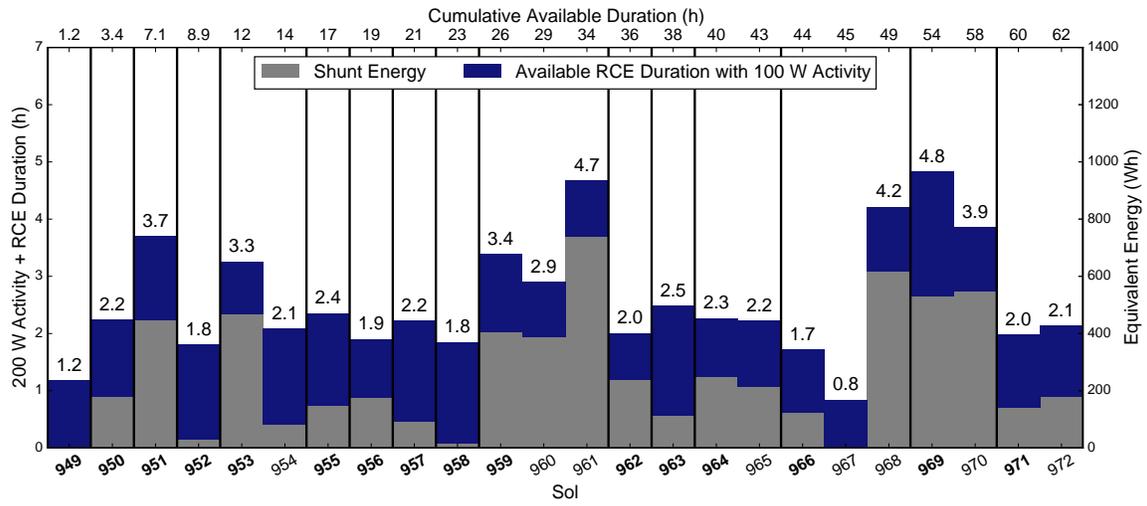


Figure 51: Estimate of extra duration availability for Artist's Drive campaign.

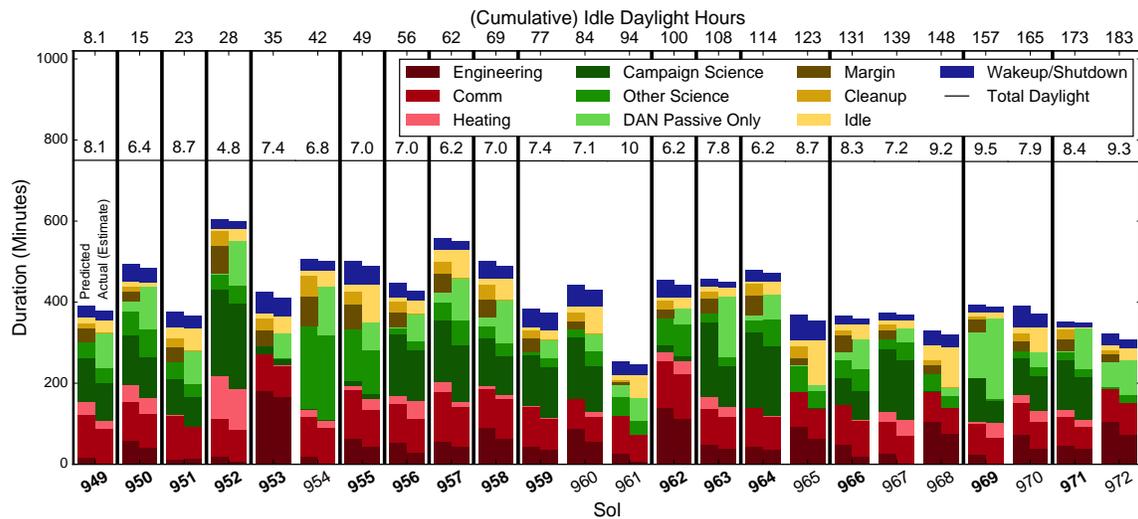


Figure 52: Allocation of RCE duration during daylight hours for Artist's Drive campaign.

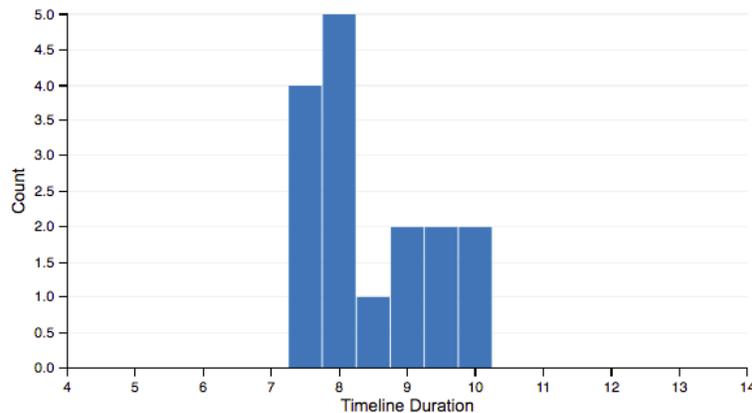
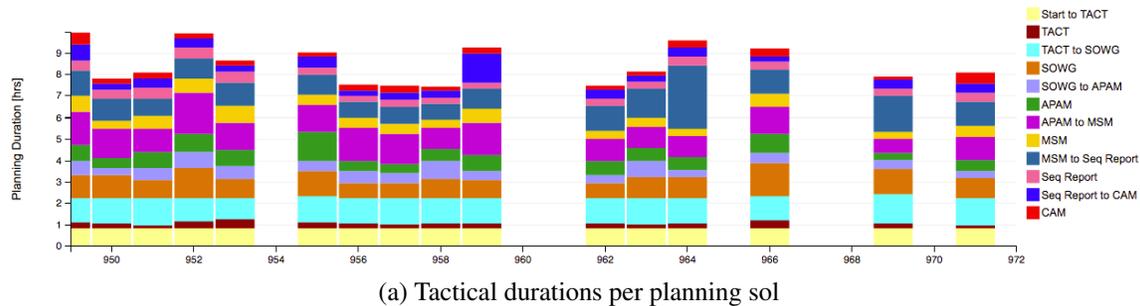


Figure 53: Tactical timeline durations for Artist's Drive campaign.

for the long duration on Sol 949 is due to the sequencing issues discussed in Section 5.1.2. There were issues with the sequencing of the pre-drive imaging block within the mobility backbone and timing issues with Mastcam activities. The long shift duration of Sol 952 was due to complexities in sequencing the drive activity. This can be seen by the larger duration between APAM and Master Submaster Review meetings in Figure 53 for this sol.

5.3 Significant Productivity Factors

Although it was generally successful, there were several factors that negatively impacted productivity during the Artist's Drive campaign.

This campaign focused on traversing Artist's Drive and collecting observations of stratigraphy along the way. However, three of the drives were limited by the visibility available to the planning team in the images collected at the end of the preceding drive. On Sol 951, a ripple near the rover prevented a long drive; on Sol 958, the rover ended its drive near a ridge that likewise occluded the drive direction; and on Sol 963, several small ripples in front of the rover limited the amount of safe terrain that could be identified.

Concerns regarding data volume had two different types of impact for this campaign. The supratactical team observed that the team would be entering a period with reduced downlink volume. As a result, the team decided to reduce the overall amount of data that was being collected in anticipation of a desire to collect a higher volume of data as the rover drove into new terrain. A more common obstacle was the predicted amount of data that could be transmitted back to Earth by the decisional passes (Sols 950, 951, 952, 956, 957, 961, 962, 963). Even if time and power would allow for the collection of more data, if it would not be downlinked by that pass then it could not be used for planning the next sol. For a drive-focused campaign, good imaging coverage of the end-of-drive location as well as the surrounding terrain is critical to enable continued progress. Since predicted data volumes are used to limit rover activities, the team relies on having good predictions (see Section 3.3).

Another common challenge was that the time available before the decisional pass imposed a limit on rover activities

Sol Type	Sols	Number of Sols
Campaign	952, 953, 956, 957, 959, 962, 966, 969, 971	9
Campaign Multi-Sol	954, 960, 967, 970	4
Extra Drives	949, 950, 951, 958, 963, 964	6
Post-Drive Multi-Sol	961, 965, 968, 972	4
Deferred	955	1
Total Sols		24

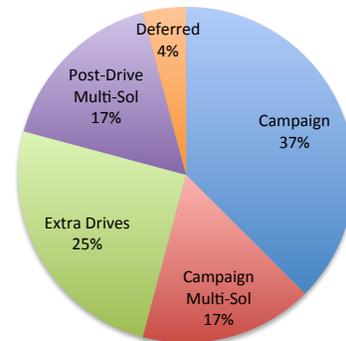


Table 3: Breakdown of sols for the Artist's Drive campaign.

(Sols 952, 955, 956, 957, 962). Although time was available for more activities during the rover's sol on Mars, the team was limited by when in that sol the decision pass occurred. Critical activities that produced products required for planning the next sol had to happen before the decisional pass. On one sol, the rover experienced some limitations due to the amount of power available (Sol 956). In general, power was not a constraining factor.

General complexity was a frequent challenge to the team. It sometimes led to catching mistakes late in the planning process (e.g., RMCA state and Mastcam timing on Sol 949, brush shadow on the dump pile target on Sol 954, arm occlusion of the dump pile on Sol 955). Similarly, an atmospheric observation was descopeed from Sol 953 for complexity reasons.

There were several activities descopeed or limited in the campaign due to concerns about their duration. However, as discussed in Section 5.2 there was additional available duration for activity during the campaign.

There were two significant interruptions to drives that merit comment. First, on Sol 954, the rover dumped a drill sample from "Telegraph Peak" and then collected images of the dump pile as well as an overnight APXS integration. On Sol 955, it was determined that a second integration would be needed to separate the dump pile from its underlying bedrock which caused the team to cancel the drive planned for that sol and defer it to Sol 956. From the perspective of the primary campaign goal of traversing the valley, this was a setback. However, it was deemed worth this cost to acquire a good understanding of the composition of "Telegraph Peak." Second, the drive on Sol 963 was interrupted after achieving only 17 of the planned 47 meters. The autonav system halted the drive when it detected a large rock near the rover's right side. The drive resumed on Sol 964, in which the rover was commanded to take images of that rock and then traverse the remaining 32 meters needed to approach "Logan's Run."

Figure 54 shows the Tom Sawyer Perspectives visualization of the data collected from the Artist's Drive campaign. Please refer to Section 4.3 for a description of the diagram nomenclature. As with the Pahrump Hills visualization, from Figure 32, there is generally a higher amount of campaign activity in non-restricted sols. The failed activity from Sol 955, shown in Figure 54, is due to the REMS schedule table upload that failed on that sol. The other failed activity was from Sol 963 in which an end-of-drive image was not acquired due to the partial failure of the drive on that sol. There were a few cases of partial failures. These included the partial failure of the APXS integrations from Sols 953 and 954, the drive from sol 963 and a failure to fully empty a CheMin cell on Sol 967.

Table 3 provides a breakdown of the campaign sols akin to what was done with Pahrump Hills. The majority of the sols contributed to the overall campaign goal of traversing Artist's Drive. However, the number of extra drives required to achieve the goals was much higher (6) than for Pahrump Hills (2). Restricted planning also had an impact, leading to a total of eight sols that were the second or third sol in a multi-sol (restricted) plan and therefore could not achieve as much as would have been possible with nominal planning.

5.4 Significant Ground-in-the-Loop Cycles

There were three major factors in this campaign that posed significant ground-in-the-loop requirements for planning:

- Drive planning. Each drive relies on data from the previous sol to assess terrain and plan the next end of drive location. Drives were planned on Sols 949, 950, 951, 952, 956, 957, 958, 960, 962, 963, 967, and 971.

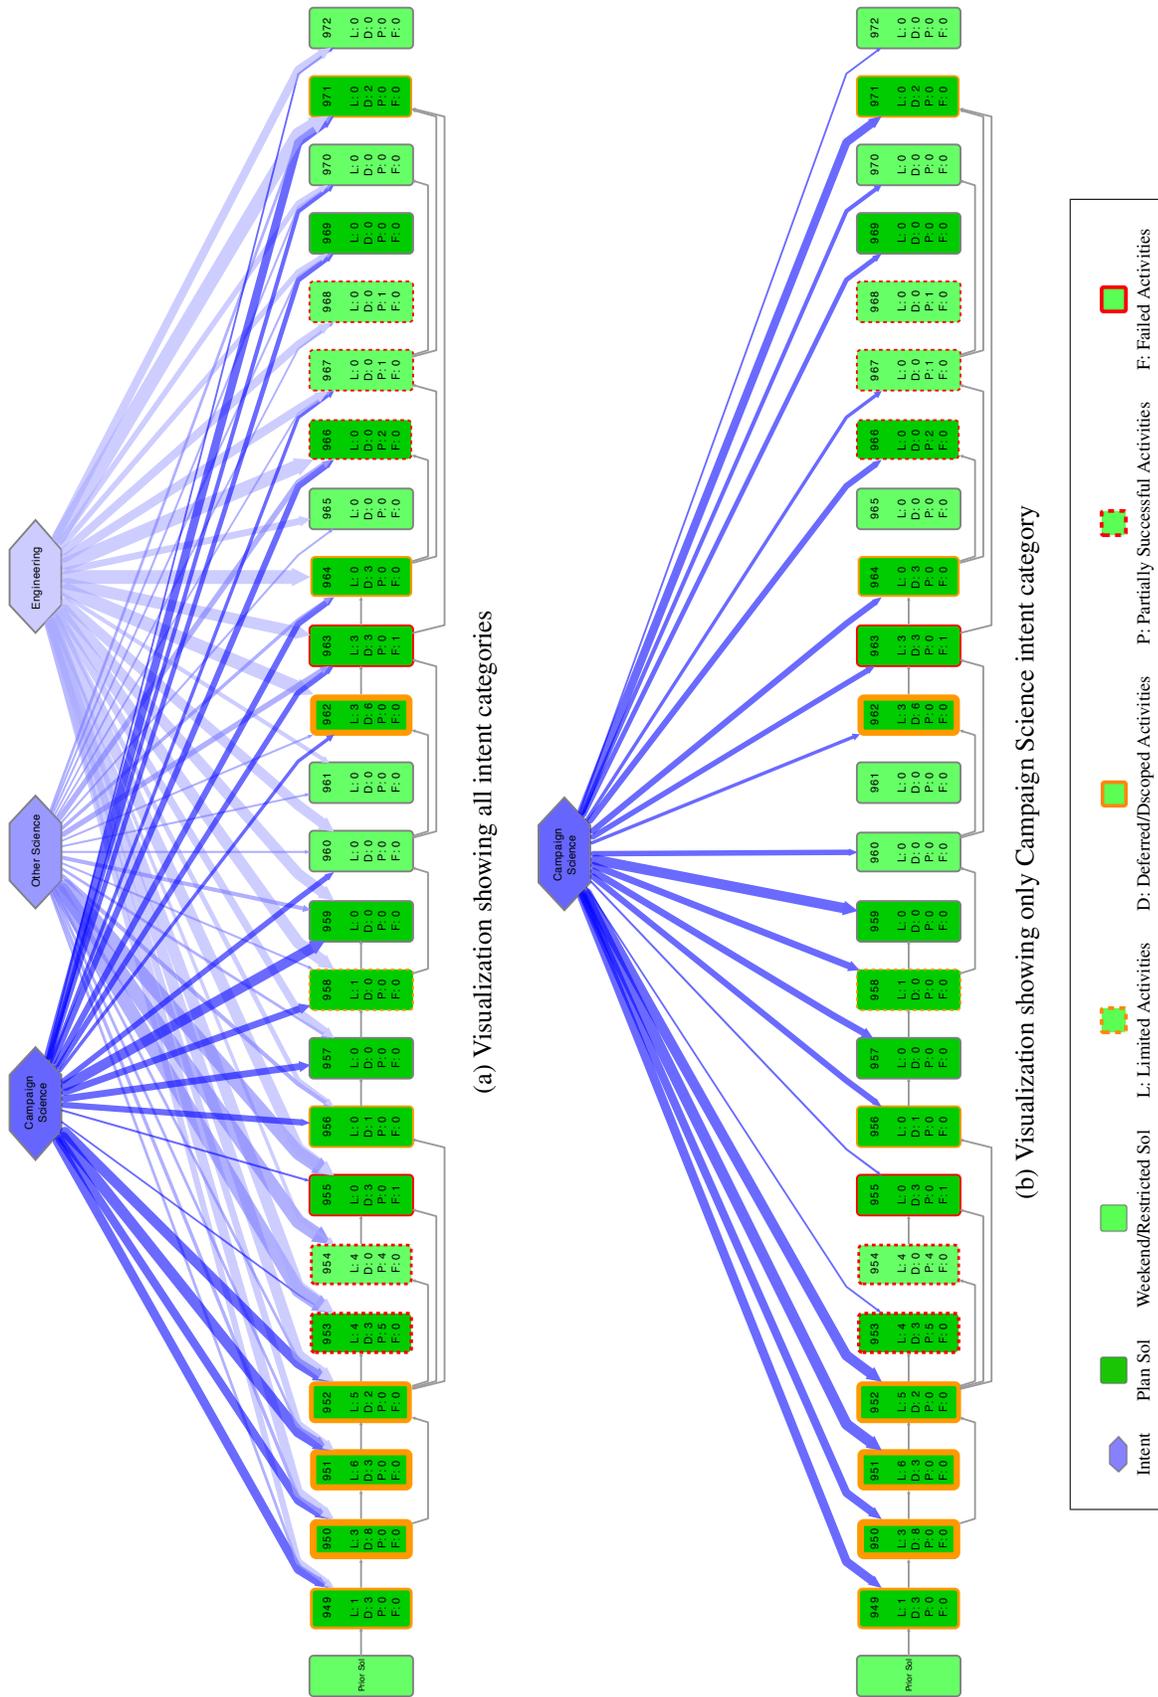


Figure 54: Visualization of the Artist's Drive Campaign.

In particular, Sol 957's drive concluded with a ridge in front of the rover, which limited visibility for further driving, so Sol 958 was employed to drive to the top of the ridge and gain a better view. Sol 963's drive was interrupted by the autonav software identifying a hazardous rock that precluded further progress. Only 17 of the commanded 47 meters were driven. The rock, named "Blackrock", was observed with ChemCam and Mastcam prior to resuming the drive.

- APXS observations of dump pile. Sol 955 was supposed to include a drive, but departing was contingent on getting good APXS data from a drill sample that the rover had deposited on the surface on Sol 954. Once images of the dump pile were downlinked, it was determined that the drill pile was thin enough that the APXS data would be contaminated with observations of the underlying bedrock. Therefore, the drive was deferred so that an additional APXS observation could be made of the bedrock near the dump pile. A ground-in-the-loop cycle was required to analyze the dump pile images and retarget the APXS.
- Scientific discovery. On Sol 963, an interesting feature was discovered that led the team to alter its strategic plan. Instead of heading straight to Logan Pass, they decided to head to a rock lens feature that was dubbed Logan's Run.

5.5 Conclusions

Although the team did not achieve the original goal of getting to Logan Pass, the Artist's Drive campaign achieved the objectives of reaching an alternative Stimson / Murray formation contact location and provided the team with opportunities of imaging the topography along the route. In addition, the brief detour to explore the intriguing outcrop at Logan's Run. The Logan's Run excursion is a good example of how the team identifies new science opportunities in data collected by the rover, and how the team balances these new opportunities with their other objectives. The science team must weigh the potential science gains with the cost (e.g. sols, wheel wear) in accomplishing these new goals.

In terms of productivity, the campaign benefited from beginning near the start of a non-restricted period. This resulted in a relatively large number of unrestricted single-sol plans and, as a result, a high level of campaign activity during the early and middle sols. The mission re-entered restricted sols toward the end of the campaign and a decrease in productivity can be observed. The challenging terrain along Artist's Drive was another significant factor in the campaign's productivity, with several drives limited by available viewshed and a drive that ended early due to an obstacle in the rover's path. The need to re-do the APXS integration on a sample dump pile also impacted productivity as the team had to defer campaign objectives for a sol while the APXS observation was re-acquired. As with Pahrump Hills, our analysis indicated that there is opportunity for increasing productivity during the campaign by making use of shunt energy and otherwise unproductive RCE duration.

6 Marias Pass Campaign

After completing investigation of Logan's Run, the team resumed the trek along Artist's Drive. The next intended destination was Logan Pass, shown in Figure 55 [Jet Propulsion Laboratory Press Release(2015d)]. Orbital imagery indicates the presence of a second formation, called the Stimson formation, that overlies the Murray formation. However, up to this point, the sedimentology and relationship with Murray formation was uncertain [Banham et al.(2016)].

The area called Logan Pass was identified in orbital imagery as a location in which the rover could study the contact between Stimson and Murray formations. The team made multiple attempts to access Logan Pass, but the unfavorable terrain in the area resulted in excessive slip during drives [Jet Propulsion Laboratory Press Release(2015m)].

Due to the difficulty in reaching Logan Pass, the team selected a nearby elevated valley, called Marias Pass (see Figure 55), as an alternative location to investigate the Stimson and Murray contact. Figure 56 show the rover traverse map with the Marias Pass contact between Stimson formation and Murray formation [Jet Propulsion Laboratory Press Release(2015c)].

6.1 Overview of Campaign

The Marias Pass campaign provides some interesting points of similarity and contrast with the Pahrump Hills campaign (Section 4). Both campaigns had the common, high-level objective of exploring a geographical area. During the Pahrump Hills campaign, the objective was to explore the first encounter with Murray formation, the basal layer of

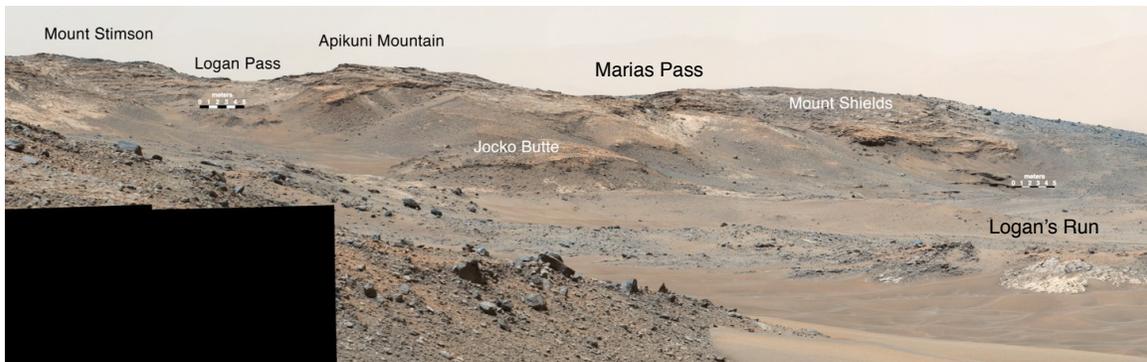


Figure 55: Logan Pass and Marias Pass (NASA/JPL-Caltech/MSSS).

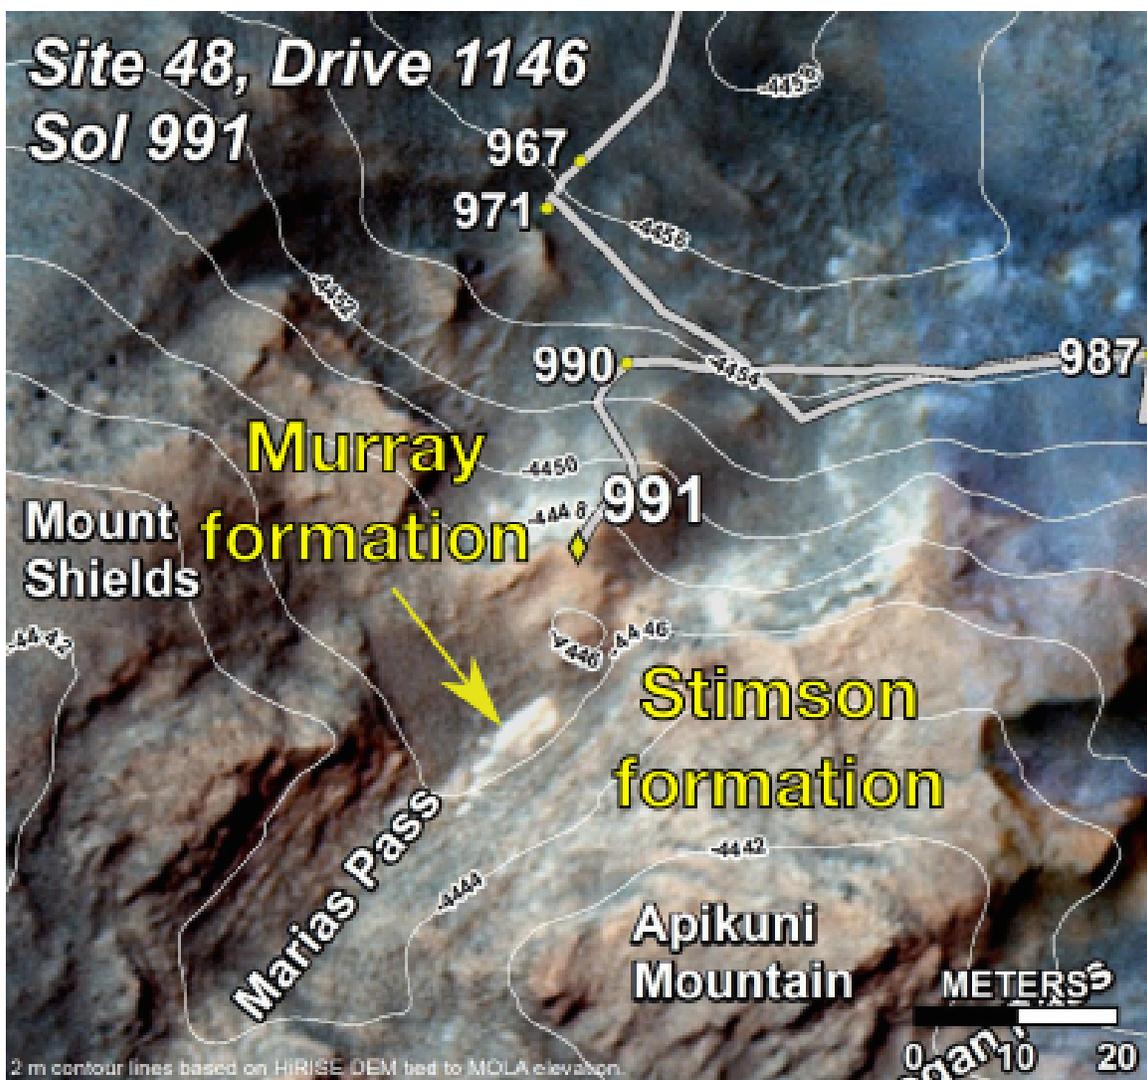


Figure 56: Traverse map through Sol 991 showing contact between Stimson and Murray formation at Marias Pass (NASA/JPL-Caltech/Univ. of Arizona).



Figure 57: Marias Pass, showing the contact between Stimson Formation and Murray formation (NASA/JPL-Caltech/MSSS).

Mount Sharp. For Marias Pass, the high-level objective was to explore the contact between the Murray formation and the overlying Stimson formation.

Once again, the team adopted a linear approach to exploring Marias Pass, which contrasts with the walkabout approach used in Pahrump Hills. In interviews with science personnel, it was stated that, after spending such a long amount of time at the single Pahrump Hills location, there was interest in completing the Marias Pass campaign in a shorter period of time.

Another important difference between the Pahrump Hills Walkabout and Marias Pass campaigns is the difference in geography between these locations. The gradual slope of the Pahrump Hills formation afforded a complete view of the area to be investigated which enabled the development of a relatively detailed strategic drive path. In contrast, Marias Pass is an elevated valley that restricted the rover's view of the contact until it arrived at the site.

The campaigns were also conducted during different Martian seasons. The Pahrump Hills Walkabout campaign was conducted during Spring in the Martian Southern hemisphere. The Marias Pass campaign began in late Summer and continued into the Martian Fall. As the temperatures decrease during Fall and into Winter, more of the rover's energy must be directed toward heating the vehicle. This can result in more challenges for the team as they try to manage the rover's available energy.

The Marias Pass campaign happened to coincide with "Mars Solar Conjunction". This is a period in time that occurs roughly every two years in which communication between Earth and Mars is obstructed by the Sun. Due to interference from the Sun, the team is unable to communicate with the rover for about two weeks during conjunction. In 2015, Mars Solar Conjunction occurred from Sol 1005 (June 4, 2015) through Sol 1026 (June 26, 2015). During this time, the rover executed an extremely limited plan that was pre-loaded before the start of conjunction. We excluded these sols from the Marias Pass case study as the content did not include activity in support of the campaign.

6.1.1 Science Objectives

Figure 57 shows the contact between Stimson formation and Murray formation at Marias Pass in a Mastcam panorama obtained on Sol 992 [Jet Propulsion Laboratory Press Release(2015e)]. The specific science objective of the Marias Pass was to study the contact between Stimson formation and the Murray formation to understand the sedimentology of Stimson formation and relationship with Murray formation. For example, the team wanted to learn if the transition between the environments that led to their formation was gradual, abrupt, or separated by a time gap. In addition, for each formation, the team the facies of the formations, their past environment of the formation and the history of its water-rock interaction.

The campaign would employ the full complement of the rover's science instruments including the use of remote sensing and close contact instruments, as well as the acquisition of a drill sample for ingestion into the rover's sampling instruments.



Figure 58: The hill to climb on Sol 991.

As will be seen in the description of the campaign activity, additional objectives were formed when the team discovered higher than expected levels of silica and hydrogen in ChemCam and DAN observations, respectively. This motivated the team to conduct additional exploration to understand these readings.

In the interest of managing scope for the case study, we chose to restrict our study of the Marias Pass campaign to Sols 991 through 997 and Sols 1027 through 1043. This sol range includes the activity that led to the unexpected silica and hydrogen observations, the initial exploration of Marias Pass and subsequent observations motivated by the unexpected silica and hydrogen findings. The gap in the sol range includes Mars Solar conjunction.

6.1.2 Summary of Activity

Following is a summary of the team's and rover's activities during each tactical planning cycle of the campaign.

Sol 991: Having decided to give up on trying to reach the Stimson / Murray formation contact at Logan Pass the team drove back to an earlier location in the hope that the rover would be able to climb to the alternate contact location at Marias Pass. On Sol 990 the rover drove 52 meters toward Marias Pass where the team would assess if the vehicle could make the climb the hill to the elevated valley of Marias Pass.

The Sol 990 drive went well, placing the rover at the base of the hill that would take it to Marias Pass. Figure 58 shows the view of the hill to be climbed, about a 6 meter change in elevation, as seen from the rover's Navcam at the end of Sol 990. This would be a challenging drive as the vehicle would be transitioning into more rugged terrain with anticipated tilts as high as 23° . This could lead to excessive slippage as the rover tries to ascend the slope. The drive had to be limited to 24 meters because the limited viewshed available to plan the drive did not provide data to plan the drive beyond the top of the hill. As such, an additional drive would be required with the additional imagery obtained at the end of the Sol 991 drive.

The pre-drive science block was used to perform a ChemCam observation on bedrock and nearby angular rocks. Targeted Mastcam observations were used to document the ChemCam targets and to characterize two nearby geological features. The team anticipated that the pre-drive science block might be limited to allow more time for driving and had formed a backup plan with reduced pre-drive science observations. However, due to the

drive being limited by available viewshed, the pre-drive science block was long enough to contain the desired imaging.

Sol 992: The hill climb planned for Sol 991 went well, reaching a max tilt of 20° and slips up to 45%, placing the rover at the entrance to the Marias Pass valley. The team was rewarded with an excellent view of the valley, including the anticipated contact between Stimson formation and Murray formations, as shown in Figure 59.

The team used the post-drive imagery to select the two locations identified in Figure 59 for closer study of the contact. These locations were selected because they appeared to be places where the contact was best exposed with Murray and Stimson formations immediately adjacent to one another without debris covering the interface. They chose to study the first with remote sensing instruments and the second with remote sensing as well as close-contact instruments. As expected, a short drive, aka a “bump”, would be needed to get closer to the locations of interest.

Location 1 was the preferred choice for contact science, but the rippled sandy area in front of the area (better seen following the Sol 992 drive in Figure 61) was judged to be unfavorable for traversability, likely preventing the rover from getting close enough to the contact. In addition, Location 1 was above a small incline. Thus, although Location 2 was more distant, it was selected as a good alternative to Location 1 because it was flat and did not require traversing through sand ripples to reach it.

The pre-drive science block was used to target ChemCam at a piece of light-toned outcrop named Elk and a dark boulder named Bull. Mastcam imagery was also acquired the Marias Pass outcrops. Figure 60 show the Elk target. Also identified is the Lamoose target which will be significant in later sols.

The initial objective for the Sol 992 drive was to drive to the first contact location. Although the team obtained a significantly improved view of Marias Pass with the Sol 991 post-drive imagery, it did not provide sufficient viewshed to plan a route to the desired contact location due to occlusions in the terrain. So, as with Sol 991, the Sol 992 drive was limited due to available viewshed. As a result, the planned drive ended up being a short 6 meter bump to a point where additional imagery could be obtained to complete the drive to the contact.

Sols 993, 994, 995, 996: This was a 4-sol plan that included three sols, to cover the weekend, and a fourth sol to cover the Memorial Day holiday. However, because planning four sols was beyond the scope of the tactical timeline, the team chose to make Sol 996 a “runout” sol. As mentioned in the Pahrump Hills campaign, a runout sol allows the team to generate sequences that cover the sol with minimal effort but results in extremely limited science. In this case, only background REMS and RAD activity would be performed on the runout sol.

The drive from Sol 992 placed the rover close enough to the first contact location (Figure 61) that the science team decided they could more than adequately document the location with Mastcam from the rovers current position. As such, the weekend plan would consist of remote sensing observations of the first contact location before driving on to the second location.

The team took advantage of available onboard storage and the multi-sol plan to perform extensive observations of the contact from this location. In fact, there was an interest in collecting a significant amount of data prior to the upcoming solar conjunction period. During this time, there would be a moratorium on commanding the vehicle, but there would still be opportunities for downlinking data.

The first two sols of the plan consisted of remote sensing observations. The first sol included a large Mastcam mosaic of the contact along with a full 360° Mastcam to take advantage of the rover’s current elevated location. The sol contained a long-range observation of the contact with the ChemCam Remote Microscopic Imager along with a ChemCam observation on a nearby outcrop as part of the continuing characterization of outcrop chemistry. The sol include additional Mastcam observations to document the ChemCam observations, to characterize fractures in the Stimson formation and to further document the Stimson / Murray contact.

The plan also included environmental observations including Mastcam observations of the Deimos eclipse.

The drive was planned on the third sol of the plan, Sol 995. The objective was to approach the second contact location identified in Figure 59. The Rover Planners needed to choose between a more direct route that would cross ripples, which can cause problems due to slippage, or a more conservative, but longer route, that would circumvent the ripples. They ended up choosing the longer route to avoid the ripples, which ended up being a 34 meter drive. Due to the need to avoid the ripples the drive would not reach the desired contact science location. An additional short bump would be required in the Sol 997 plan to complete the drive to the desired location.

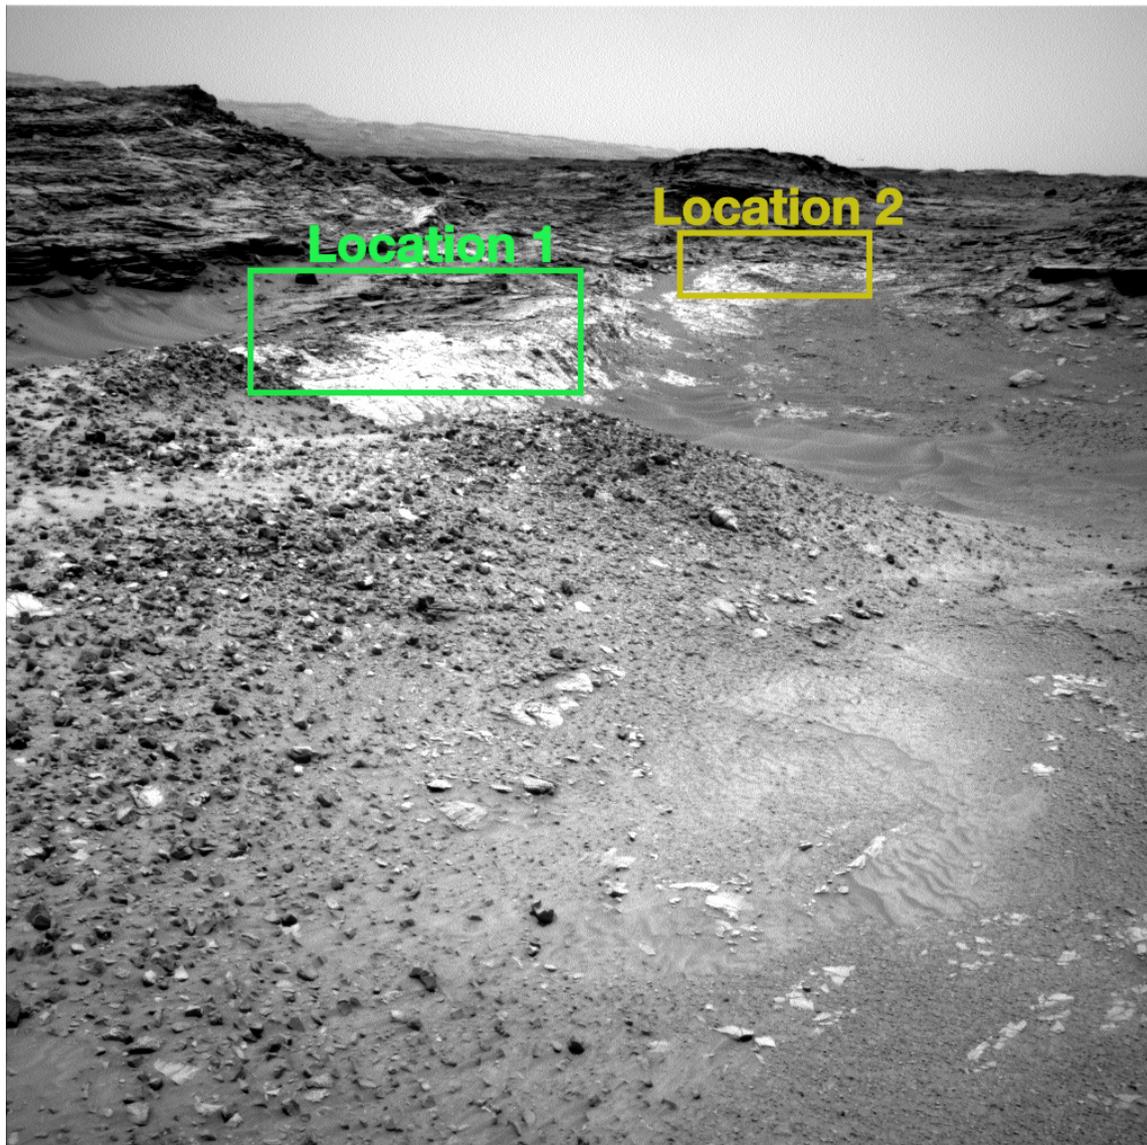


Figure 59: View of the Marias Pass contact following the Sol 991 drive showing the locations selected for closer study.

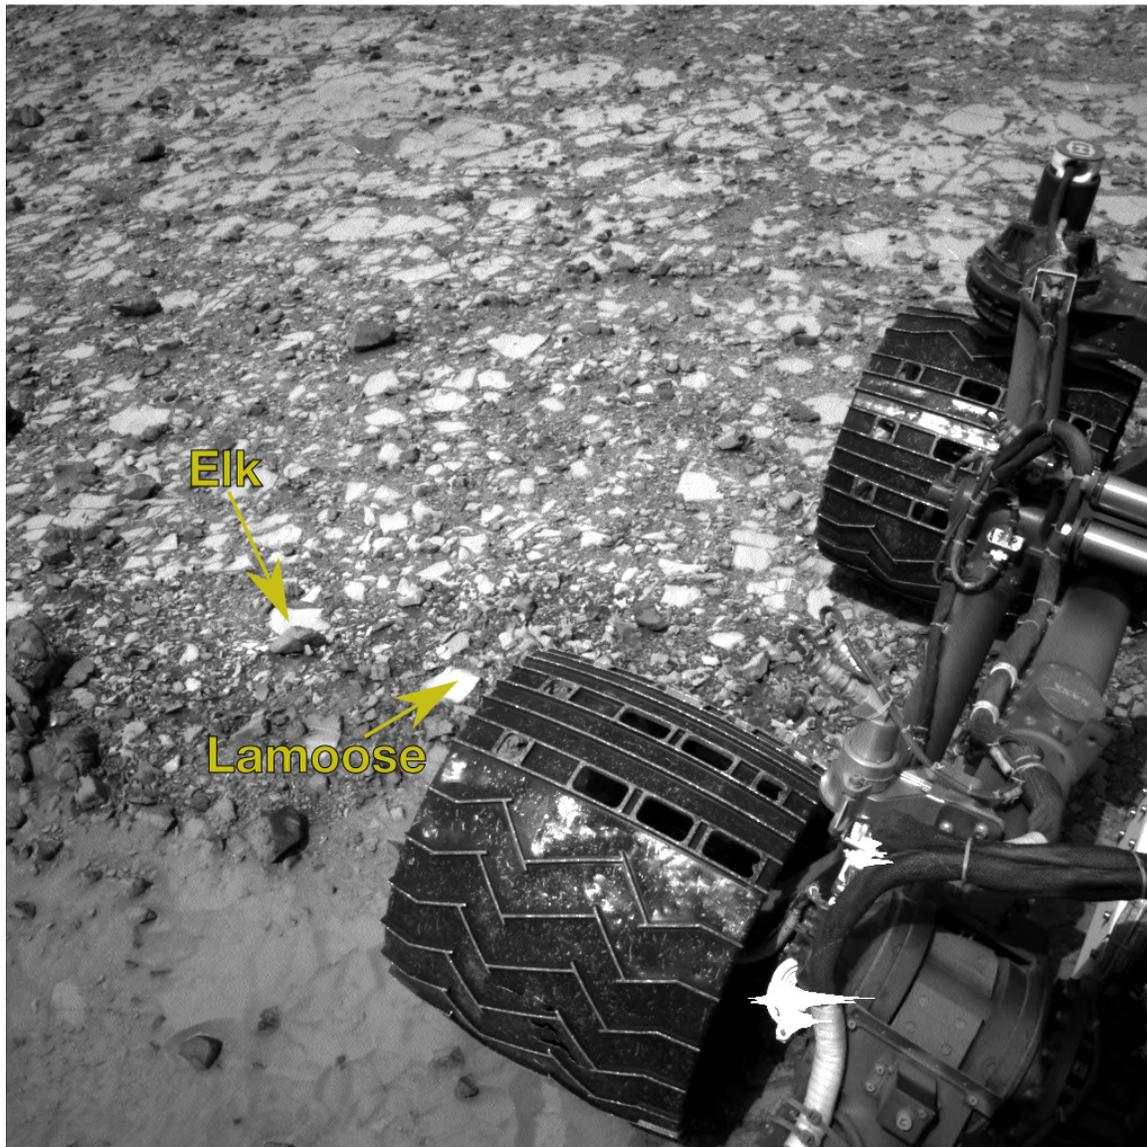


Figure 60: Navcam image taken after the Sol 991 drive, showing the Elk and Lamoose targets.

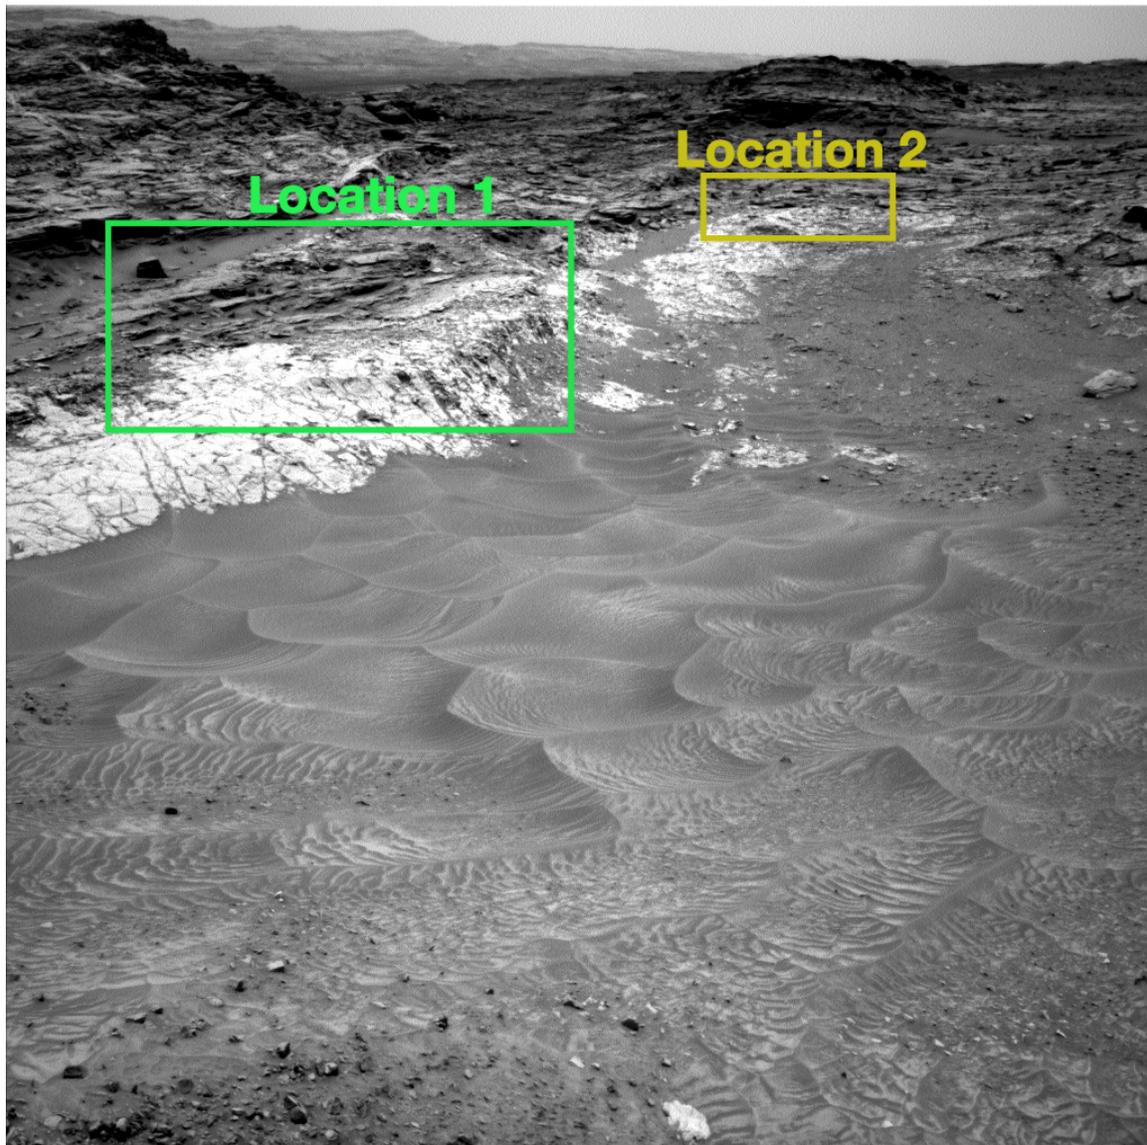


Figure 61: View of the Marias Pass contact following the Sol 992 drive.

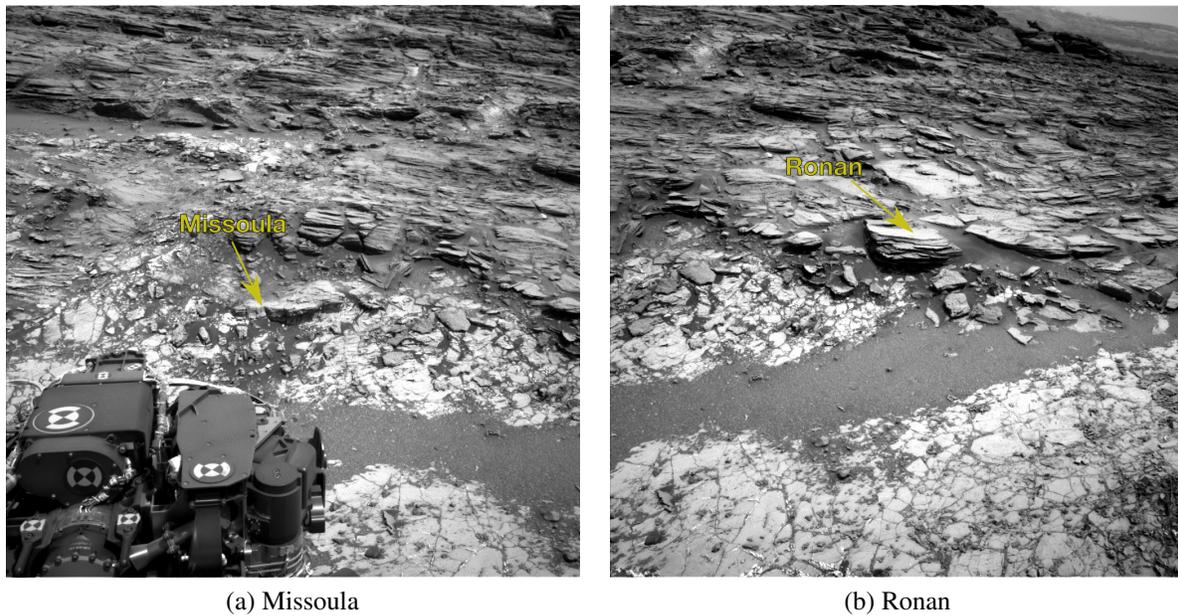


Figure 62: Navcams following the Sol 995 drive, showing the Missoula and Ronan targets.

Sol 997: The drive from Sol 995 went well, leaving the rover close to the area where the team wished to perform contact science. As expected, a short bump would be required to get the rover into the desired position.

The team selected targets within the Murray formation and the Stimson formation for remote sensing observations. ChemCam observations were taken on a Murray formation target located in outcrop close to the rover and on the Missoula target (shown in Figure 62 (a)) which is at the start of the Stimson formation. Missoula was chosen as an important target because it appeared to be a continuous outcrop across the contact, i.e. a single piece of bedrock that had Murray formation at its base and Stimson formation at its top.

The team also picked out a target in the Stimson formation for upcoming contact science observations, including potential use of the DRT. The selected target, named Ronan, is shown in Figure 62 (b). Ronan was of interest because it was a large, nearby block of the Stimson formation, with many internal beds exposed on its sides. A Mastcam mosaic was acquired of the Ronan target prior to the drive.

The bump planned for Sol 997 consisted of a short, straight drive of 2.7 meters with a heading change at the end to be in position for contact science on target Ronan.

Sols 998 through 1026: In order to manage the scope of our case study, we decided to exclude Sols 998 through 1026 from our detailed case study analysis. For completeness we provide a high-level summary of the activity that took place on these sols.

Special attention was required during planning this sol for the team to establish that the rover had not shifted during the long conjunction. Typically the team makes use of images from the Hazard cameras to monitor slippage. However, these were not acquired over conjunction. Instead, the team was able to use results from a periodic attitude measurement activity, based on the rover's inertial measuring unit, to establish that the vehicle had maintained its attitude through conjunction. This enabled the team to go forward with the planned contact science activity.

Sol 998 was largely devoted to studying of target Ronan (Figure 62 (b)). The team acquired ChemCam and Mastcam observations. Next the team acquired MAHLI images of the target prior to brushing with the DRT, then cleaned the target with the DRT and followed up with post-brush MAHLI images. An APXS observation was taken on the brushed Ronan target.

For Sol 999, the team selected a target in the Murray formation, named Wallace, for investigation. As with

Ronan, MAHLI images were acquired of Wallace before and after a DRT brush. An APXS observation was taken on the brushed Wallace target.

Activity was limited from Sols 1000 through 1004 due the approach of solar conjunction. The team became conservative with what activities they would perform with the rover due to the limited opportunities they would have to address any problems prior to the command moratorium during conjunction. Consequently, these sols consisted of remote sensing observations.

Solar conjunction occurred from Sol 1005 through Sol 1026. During this time, the rover executed a long, but highly restricted plan that had been previously loaded. The content of the sols was limited to environmental observations with the RAD, DAN and REMS instruments.

Sols 1027, 1028, 1029: The team returned to planning after the extended solar conjunction hiatus. It was a weekend, so a 3-sol plan was created. The plan included a significant amount of engineering maintenance activity since these activities had not been performed over conjunction.

One of the more important data management activities intended for the plan was to delete onboard data that had been received by Earth. The team had collected a large amount of data in preparation for downlink over the conjunction break. Now that much of the data was received, the team wanted to delete the data that had been sent to make room for new data collection.

Unfortunately, there was a problem with the data management activity in the Sol 1026 plan. A bug in the tools that translate Earth time to the time system used by the spacecraft caused errors in some of the commands to delete data products. Fortunately, this did not result in the accidental deletion of unintended data products. However, given the potential for that type of error, the team took a conservative stance and decided to not send data product delete commands until the problem was better understood and a workaround could be established. This turned out to not directly impact the Sol 1027 plan, but as will be seen, it did have a negative impact on the subsequent plan for Sol 1030.

However, the removal of the data product delete activity had an unanticipated repercussion on subsequent command product development. The team had reserved a large amount of time in the plan for the anticipated data delete sequences. Per standard practice, after the time, the plan would include a cleanup activity to ensure stop the data management sequences if they ran unexpectedly long. When the decision was made to descope the deletes, it opened up a large gap of time in the plan between the now shortened data management activity and the cleanup. The scripts the team uses to put wakeup and shutdown activities in the plan saw the gap and inserted a nap between the data management activity and the cleanup. It was not noticed until later in the day that this resulted in a violation of standard sequencing practice because there was no cleanup between the data management activity and the shutdown. Because the problem was found late in the day, the team removed the extra shutdown, keeping the rover awake until the subsequent cleanup. This incident provides an example of the complexities of command product generation and how issues can propagate in unexpected ways.

On the first sol of the plan, the science team performed remote sensing observations with the intent to compare the scene with pre-conjunction images to see if anything had changed due to winds or potential Mars-quakes. A ChemCam observations were acquired on two Murray formation targets including the Wallace target from Sol 999.

On the second sol, the team used the MAHLI instrument to perform a goniometer observation of the area around the Stimson / Murray contact. This is a process in which the MAHLI is used to image an area from a close standoff at many different angles and enables the team to investigate the light scattering properties of rocks and soils [Johnson et al.(2015)]. APXS observations were acquired on a coarse-grained float rock, called Big Arm, that was previously imaged on Sol 999.

The team encountered additional problems with data management sequences later in the shift. The team needed to perform an unusually large amount of data product retransmits, due to the need to cover the extended conjunction period. The large volume of retransmits uncovered a latent problems with the tools used to generate the sequences including incorrectly generating the retransmit commands and generating a sequence which a size that was larger than the maximum allowed sequence. Fortunately, the command product validations caught the problems. But due to the late time in the shift in which the problems were encountered the team decided to defer the requests to a future plan.

Sol 1030: The team continued the investigation of the Stimson / Murray contact in the Sol 1030 plan. As noted above, the team encountered a problem with a data management sequence on Sol 1026. At the start of Sol 1030 planning, a workaround had not been put in place and, therefore, the team continued the moratorium on deleting onboard data products. With the accumulated data in preparation for conjunction and the additional data acquired in the previous weekend's plan, the team was now at the point that it had to impose constraints on how much data that could collect in the Sol 1030 plan.

The data product restriction resulted in several observations being descoped from the plan. This included a ChemCam target with Mastcam documentation, a systematic geological Mastcam observation and a systematic environmental Navcam observation.

The team picked two vein targets for ChemCam, one target above the contact and one below. Mastcam observations were taken to document the ChemCam targets. The plan also included an observation to observe the shadow of Phobos moving across the Gale Crater rim.

After pre-drive imaging, the team planned a short bump to put the rover in position to perform close up study with arm-mounted instruments on the Missoula target. The rover was on relatively flat ground with no nearby hazards. However, the drive still posed a challenge to ensure the rover is in a good position for use of the arm instruments on Missoula as well as being in a good position to use ChemCam on the target, which would require placing the rover a little further back than otherwise required. In addition, the team would also need to ensure that sufficient visual documentation of the rovers final position is obtained to support the rover stability analysis that is required each time the rover's arm is employed. In this case, the rover planners anticipated that the designed sequence would provide documentation of the terrain on which the front and middle wheels would rest, but not the rear wheels. After further analysis it was concluded that this should be sufficient to perform the stability analysis required for the type of science activity planned for the target. In particular, the team would not be coming into actual contact with the target, but instead "hover" close to the target with sufficient standoff. This drive is a good example of how even a short bump, in this case with a net change in position of about 1 meter, on relatively benign terrain can still be quite challenging to plan. In fact, the tactical Rover Planners enlisted the support of the Strategic Rover Planners to assist in some of the required analysis.

Sol 1031: The hard work of the Rover Planners paid off with a successful bump on Sol 1030 to the Missoula target. A workaround was established for the data management problem that arose during Sol 1030 plan which resulted in lifting the previous plan's data collection constraint. And the engineering team was able to perform the stability analysis necessary to approve use of the arm for the Missoula target activities.

The major activity for the plan was using MAHLI to acquire a "dog's eye" mosaic of the Missoula target. This was a very successful observation and resulted in the mosaic shown in Figure 63 [Jet Propulsion Laboratory Press Release(2015a)]. The mosaic shows mudstone of the Murray formation at the bottom and coarser sandstone of the Stimson Formation at the top.

Later in the day, the team acquired a series of ChemCam observations on several targets across the contact zone to further characterize the contact. One of the ChemCam observations took advantage of a freshly exposed surface on a rock in the Murray formation that was broken apart by the rover's wheels. Mastcam observations were acquired to document the ChemCam observations along with a mosaic to continue the characterization of the contact.

Sol 1032: The plan for Sol 1032 was to wrap up the Marias Pass campaign with contact science on the Missoula target. The team considered staying in the location for additional investigations but instead, during the conjunction break they decided that after finishing up at this location, they would drive back to a previous location, the rover's location during Sol 991 and 992 plans, where some interesting results from the ChemCam and DAN instruments were observed. The ChemCam observation on the target Elk (Figure 60) showed surprisingly high levels of silica. And the DAN instrument detect high levels of hydrogen during a drive in that area. These findings were discussed at a science team meeting during conjunction and the full team agreed to go back to investigate.

But prior to driving away on Sol 1033, Sol 1032 would be spent continuing the Marias Pass investigation near the Missoula target. During our interviews with MSL personnel, it was noted that MSL productivity can benefit greatly when the supractical team is able to pre-plan activities for the tactical team. Sol 1032 is a good

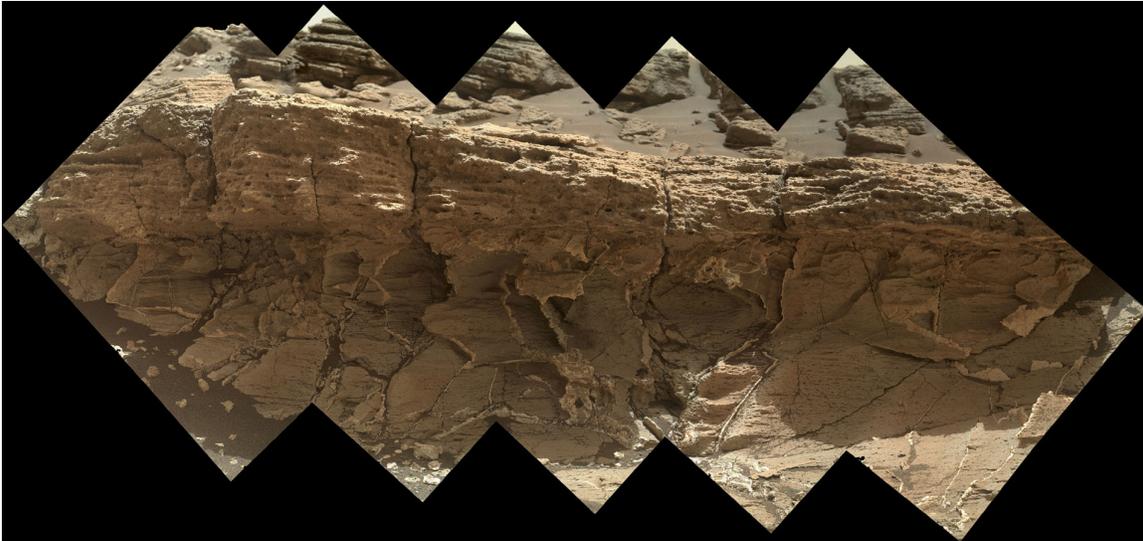


Figure 63: MAHLI mosaic (40cm across) of the Missoula target at the contact between Stimson and Murray formations at Marias Pass (NASA/JPL-Caltech/MSSS).

demonstration of this. While the tactical team was planning the dog's eye view of Missoula on Sol 1031, the supratactical team was planning a MAHLI mosaic on a target called Clark, an area just to the left of the dog's eye mosaic, a second set of MAHLI observations on the rock from Sol 1031 that was broken by the rover's wheels, and finally placing the APXS on a target for overnight integration.

The tactical team expanded on the supratactically provided contact science plan by adding in MAHLI observations on an additional target, called Lumpry.

The arm activities were split into two pieces to allow time between them for a Mastcam video of a Phobos transit.

The plan also included an extensive set of targeted remote sensing observations. A Mastcam observation was re-acquired on a target that had previously been imaged on Sol 1030. It was re-done in the Sol 1032 plan because the target was in shadow during the Sol 1030 observation. Several Mastcam multi-spectral observations were performed on targets as a followup to a pre-conjunction observation. Three ChemCam targets were selected, including a one of the Mastcam multi-spectral targets, a target where the color and texture of a rock change in relation to distance from a vein, and rock with texture representative of other rocks in the area.

There was significant discussion about where to perform the overnight APXS integration. The science team was interested in placing the APXS on a lighter-toned clasts. Although there were many options in the area, none were in reach of the APXS from the rover's current position. Instead, they chose the Lumpry target, on the Missoula-type outcrop, which may be related to the lighter-toned clasts.

Sols 1033, 1034, 1035, 1036: The team developed another 4-sol plan because the team would be taking Friday off as part of the Fourth of July holiday weekend. As with other 4-sol plans, the fourth sol in the plan would be a highly restricted "runout" sol consisting of only REMS and RAD observations.

As mentioned above, the team had decided to drive away from the Stimson / Murray contact during this plan to return to the location the rover was in at the end of the Sol 991 drive. The decision was prompted by analysis of ChemCam and DAN data over the conjunction break. ChemCam results on the target Elk (Figure 60) indicated higher than expected levels of silica. In addition, routine observations made by the DAN instrument during the drive in that area detected high levels of hydrogen.

As such, the team would conclude remote sensing observations at the contact and then drive back to the Sol 991 end-of-drive location. There was interest in doing additional contact science in addition to the drive, but due to complexity of performing both contact science and drive activities, combined with staffing constraints, the team was unable to perform both types of activities, and opted for driving.

Given the extended weekend plan, the team had two full sols of targeted remote sensing. They took advantage of the time to perform some unique observation. A suite of multi-spectral Mastcam and ChemCam observations were taken on a vein target, called Thunderbolt, to investigate the effects of dust cover and time of day on the spectra. Mastcam multi-spectral was acquired before and after the ChemCam observation which exploited the side-effect of ChemCam which has the ability to clear off local dust as part of firing the laser on rocks. Additional Mastcam multi-spectral observations were taken on this same target across the first two sols at different times of the day.

A high priority ChemCam observation was planned on a target to attempt to hit lighter-toned clasts in the Missoula-like outcrop. There had been previous attempts for this type of observation on two different targets in a previous plan, but they appear to have missed the light-tone clast material. There was high interest in this observation as the team wanted to understand a potential relationship between the lighter-tone material in the Murray formation and the Stimson formation. The result was a partial success as two points of the raster hit lighter-toned clast, but on their edges. A ChemCam observation was also performed on the Lumpry target for which an overnight APXS integration was acquired.

Finally, the team made use of the Navcam imagery acquired in the previous plan to design a Mastcam mosaic of the contact. The priority for the mosaic was to capture additional sections of the contact at the base of the outcrop and a sampling of the cross-bedding higher up in the Stimson formation.

The remote sensing plan was filled out with systematic environmental observations to monitor dust and clouds in the atmosphere.

The plan for the drive back to the Sol 991 location was essentially a reverse of the drive the rover took to get to its current position. This allowed the team to take advantage of knowledge they have of the previously traversed terrain. Given the interest in investigating the DAN readings near the Sol 991 location, the 40 meter drive included several DAN observations toward the end of the drive, the last 10 meters between the end-of-drive locations from Sol 992 to Sol 991. The end goal of the drive was to be within reach of the Elk target and other nearby light-toned rocks in preparation for contact science with the arm. This would be somewhat challenging because this location is just over the lip of the 20° hill the rover climbed to access Marias Pass. As such, the rover would be at a pitch of almost -10°, which would pose a challenge for stability for contact science. Given the length of the drive, about 40 meters, the lateral position of the rover could deviate and necessitate an additional bump in the next plan to be in position to reach the targets of interest.

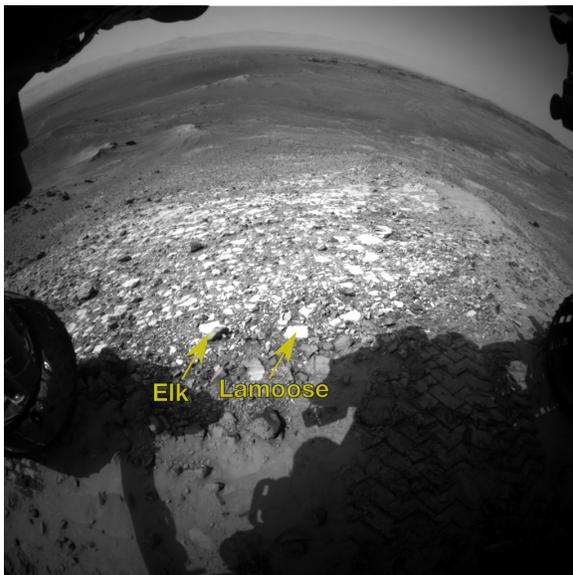
Sols 1037, 1038: Returning from the Fourth of July weekend, the team entered restricted sols, which means they would develop a 2-sol plan and the data from the second sol would not arrive in time for the next planning day.

The team returned to some unfortunate news related to the drive from Sol 1035. The drive ended about 5 meters short of the intended location. The rover uses visual odometry (an onboard technique that estimates motion by comparing consecutive images) to improve its location knowledge as it drives. Terrain conditions, especially lack of sufficient texture, can result in a failure for the visual odometry algorithm to converge. These failures increase the uncertainty in the rover's positional knowledge. On this drive, visual odometry failed as the rover approached the steeper terrain toward the end of the drive. The drive was set to terminate due to the risk of increased uncertainty in the rover's position for this drive. Figure 64 (a) shows the view from the rover's front hazard camera at the end of the Sol 1035 drive. The rover's tracks from its previous drives in this area can be seen along with the lip of the hill can be seen in the distance. The targets of interest are just on the other side of this lip.

Prior to resuming the drive, a pre-drive science block was planned containing three targets for ChemCam observations. One of the targets took advantage of another rock that was broken by the rover's wheels when during previous drives in this area. The broken rock exposed a light green-gray interior. The observation would allow the team to understand the interior chemistry compared to the exterior chemistry of other rocks. A second target was a lighter-toned rock. And a third target was on an area suspected to be bedrock above the Elk target. Collectively, all three observations would enable the team to study chemical variations in the formation. Targeted Mastcam observations were acquired to document these three ChemCam observations. An additional Mastcam mosaic was planned on a region of outcrop that exhibited steeper dipping beds compared to surrounding bedrock. The plan also included Navcam and Mastcam observations in support of environmental monitoring.



(a) Hazcam image at end of Sol 1035 drive



(b) Hazcam image at end of Sol 1037 drive



(c) Hazcam image at end of Sol 1039 drive

Figure 64: Hazcam images following the drives back to the Sol 991 location.

The drive on Sol 1037 was planned for a short 5.5 meter bump to complete the drive back to the Sol 991 location. Additional DAN observations were performed along the traverse. The Rover Planners adjusted the pointing to be used for the visual odometry images to reduce the risk of the convergence failure that prematurely ended the Sol 1035 drive. The objective of the drive was to put a light-toned rock, similar to the Elk target, within reach of the rover, with the possibility of also being within reach of the Elk target.

Sols 1039, 1040: The Sol 1037 drive brought the rover over the lip of the hill with the light-toned rocks in view, as seen in Figure 64 (b). The plan for the day was to have been performing contact science on the light-toned rocks. However, the data following the Sol 1037 drive indicated that the rovers left front wheel was perched on a small scarp at the top of the slope. There would be a chance the rover slipping down-slope if the arm were to be unstowed in this position, which could be a risk to the arm and its instruments. It would be necessary to back the rover up slightly to a more stable position.

A pre-drive, targeted remote sensing block was filled with four ChemCam observations. Three targets were of Elk-like material, including two that was a candidate for an APXS observation once the rover was re-positioned and one that was recently broken by the rover. A fourth target was of dark material that might be from the Stimson formation. Mastcam images were also acquired to document the ChemCam observations.

The plan included Mastcam images of the Sun to look for sunspots. This was a unique opportunity as Mars was currently on the opposite side of the Sun from Earth and the rover's Mastcam was the only available instrument for observing this side of the sun [Jet Propulsion Laboratory Press Release(2015h)].

The drive for Sol 1039 would be a very short bump of 35 centimeters backwards, to place the wheels on more stable ground. The drive avoided some small rocks that would pose potential stability risks for the next plan. The downside of backing up is that it puts some of the desired targets out of reach of the arm. Other positioning options were considered, but in the end it was felt backing up had the best chance of achieving the stability requirements for deploying the arm. The only target anticipated to be within reach after the bump would be too small for use with the DRT.

Sols 1041, 1042, 1043: The bump backwards on Sol 1039 successfully placed the rover in a more stable position. However, because the rover was set back further from the targets of interest, it meant that only one target of interest, Lamoose, was in reach of the APXS. Figure 64 (c) shows the rover's position for the start of Sol 1041.

The first sol of the 3-sol weekend plan included remote sensing and contact science. Targeted remote sensing included ChemCam of a dark rock, Mosquito, which would also be observed with the MAHLI instrument, a light-toned rock, and a passive-only observation on a bright rock. Mastcam observations provided documentation for ChemCam as well as collecting multi-spectral data on Elk, Lamoose and the light-toned ChemCam target in this plan.

Later in the first sol, the arm was used to acquire MAHLI images of Mosquito and Lamoose. The team also used MAHLI to perform a periodic inspection of the REMS instrument. Finally, the APXS was placed on Lamoose for an overnight integration.

Sunspot observations with Mastcam were planned throughout the weekend plan as well as a systematic environmental observation with Navcam.

Ultimately, this location turned out to be unsuitable for drilling into an Elk-like target due to the sloping terrain. Instead, the team planned to drive back toward the Ronan target (Figure 62 (b)) at the contact location, in anticipation of a potential drilling activity. Once again, the rover would be re-tracing past steps enabling the team to make use of information from past drives. DAN and Mastcam observations were be acquired periodically along the drive. Unfortunately, similar to the Sol 1037 drive, the sol 1042 drive also ended early, after about 17 meters, due to visual odometry convergence failure.

6.1.3 Outcome of Campaign

Although we decided to end our case study of the Marias Pass campaign at Sol 1043, the exploration of the Stimson / Murray formation contact continued. The team resumed driving on Sol 1044, following the drive that halted early on Sol 1042. Sol 1046 was largely spent conducting periodic wheel health assessment imaging. The team had not fully achieved their goal of drilling an Elk-like rock due to the challenging slope at the Sol 991 location. However, they

identified a new candidate location, called Lion, on more level ground, though it ended up being a challenge to reach this location.

The team began driving toward Lion on Sol 1049. That drive ended early due to tripping a rover pitch limit. The drive was resumed in the Sol 1051. An additional bump was needed on Sol 1053 to reach the Lion area. The Sol 1053 drive ended early, this time due to a suspension limit being reached and an additional, short bump was needed on Sol 1056. Thus, on Sol 1057, the rover was at the selected drilling location, a target named Buckskin, and began preparing for the drilling activity. The drill sample was acquired on Sol 1060.

This drill sample achieved one of the major objectives of the Marias Pass campaign, which was to collect a drill sample in the Murray formation. The team had initially planned to also acquire a drill sample of the Stimson formation in this area, but instead decided to wait until later for this sample since there would be ample opportunity to drill into Stimson as the rover would be driving on this material for a while, and this gave the team additional time to analyze the results from the Buckskin sample.

Despite the many challenges faced by the team while exploring Marias Pass, it was a highly successful campaign. Data gathered during the campaign enabled the team to characterize the sedimentary facies and sedimentary facies and architectural elements preserved within the Stimson formation, and to determine the depositional history and regional relationships with the underlying Murray formation. These results are important in understanding the later history of the infilling of Gale crater and the potential habitability of the environment [Banham et al.(2016), Watkins et al.(2016)]. The Missoula target and the observations with from Mastcam, ChemCam, and MAHLI provided a rich set of data in support of the development of a depositional history of the Stimson / Murray contact [Newsom et al.(2016)].

The study of high-silica targets lead to the conclusion that the introduction of silica represented one of the most recent water-rock interactions observed in Gale crater [Frydenvang et al.(2016)]. The data would also contribute to the study of the subsequently visited Bridger Basin area [Gasda et al.(2016)].

The Buckskin drill sample contributed the hypothesis of a silicic volcanism process [Morris et al.(2016)] and, combined with previous drill samples helped to form a theory of history of the Murray formation [Rampe et al.(2016)].

6.2 Analysis of Resource Usage

As with the other campaigns, we performed a series of analyses to understand how mission resources were used during the campaign. Please see Section 4.2 for an overview of the plots. Note that for the Marias Pass campaign, the plots have a gap for the solar conjunction period.

Figure 65 shows the allocation of RCE duration for each sol of the campaign. As with the other campaigns, this plot shows how the overall activity in multi-sol plans decreases compared to an equivalent number of single-sol plans. This is not as immediately obvious in this plot for some multi-sol plans as it was for other campaigns. This is due to some sols having a large overall RCE allocation but with significant amount of nonproductive duration (i.e. Idle and DAN Passive Only). For examples, there is extensive Idle duration in Sol 1027. The large Idle duration was due to the problems related to the data management delete sequences described in Section 6.1.2. Sols 1030 and 1040 had a large amount of DAN Passive Only duration. This was due to timing required for short atmospheric observations in the middle of, otherwise, idle time. The team left the vehicle up collecting DAN Passive observations because they had sufficient energy to do so. Note that in the Sol 1040 plan, the team labeled the Idle duration around the atmospheric activity as Margin, which is why Sol 1040 has a larger than usual allocation of Margin duration.

Figure 65 is also similar to the other campaigns in that the allocation to campaign objectives dramatically decreases following drives in multi-sol plans. This can be seen on Sols 1038, 1040 and 1043. Note that this campaign had two cases of four-sol plans, to cover long weekends with holidays. The fourth sol of these four-sol plans had very low activity due to the use of runout to cover the rover's activities on these sols.

This plots shows a few cases where actual execution durations were significantly shorter than predictions. These are due to drive activities not taking as long as predicted. The drive on Sol 991 completed successfully but did not take as long as anticipated. The drives on Sol 1035 and 1042 ended early due to visual odometry convergence failure.

Figure 66 shows how the acquisition of data was allocated during the campaign. As with other campaigns, the larger differences between predicted and actual were with non-critical priority bins. From the sol descriptions above, the most significant sols in terms of data volume constraints were Sols 1027 through 1030, when activities were descope due to concerns regarding onboard data product storage. Figure 66 shows that relatively little data was collected over these sols. The relatively large discrepancy between predicted and actual data acquisition on Sols 1028 and 1030 indicates the team could have descope fewer activities on these sols and still met the predicted data volume

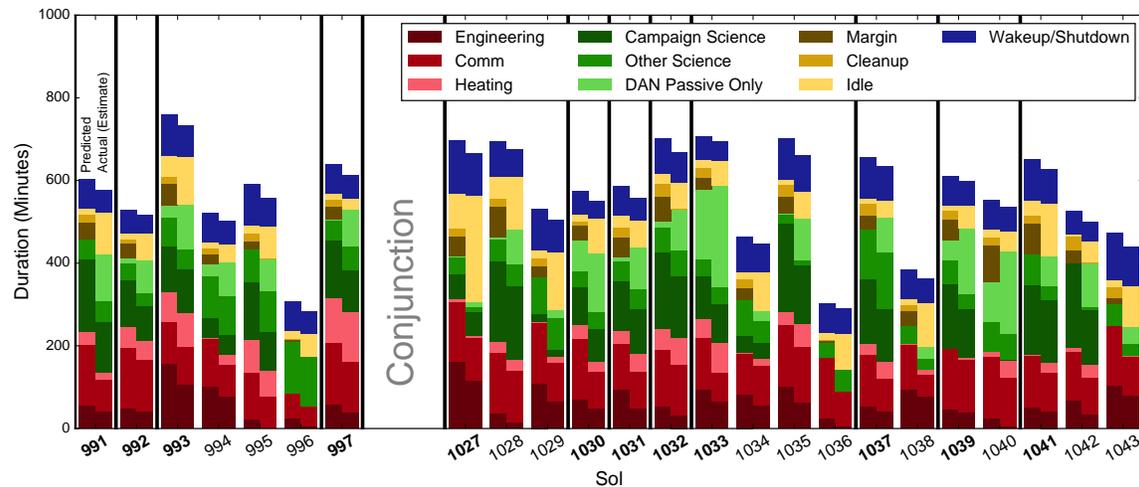


Figure 65: Allocation of RCE duration for Marias Pass campaign.

amount.

Figure 67 shows the predicted and actual performance of the orbiter relay communication windows in terms of number of Mebibits transferred. Overall, actual downlink performance matched predicted performance. The main outlier is the decisional data from the Sol 1027 through Sol 1029 multi-sol plan in which MRO passes highly outperformed predictions. Although there was relatively low decisional volume on Sols 1031 and 1032, due to the type of activity performed on these sols, i.e. contact science, there was little decisional data required for the next plan.

Figure 68 shows the predicted vs. actual warmup heating duration in the plans. The overall amount of warmup heating for this campaign was slightly higher than for Artist's Drive. There was a total of 36 hours of predicted pre-heat and 12 hours of actual pre-heat vs. Artist's Drive which had 30 hours of predicted pre-heat and 9 hours of actual pre-heat. This slight increase in heating is consistent with the seasonal temperatures at Gale Crater as shown in Figure 6. The Marias Pass campaign occurred shortly after the end of Artist's Drive and the temperatures were trending downward.

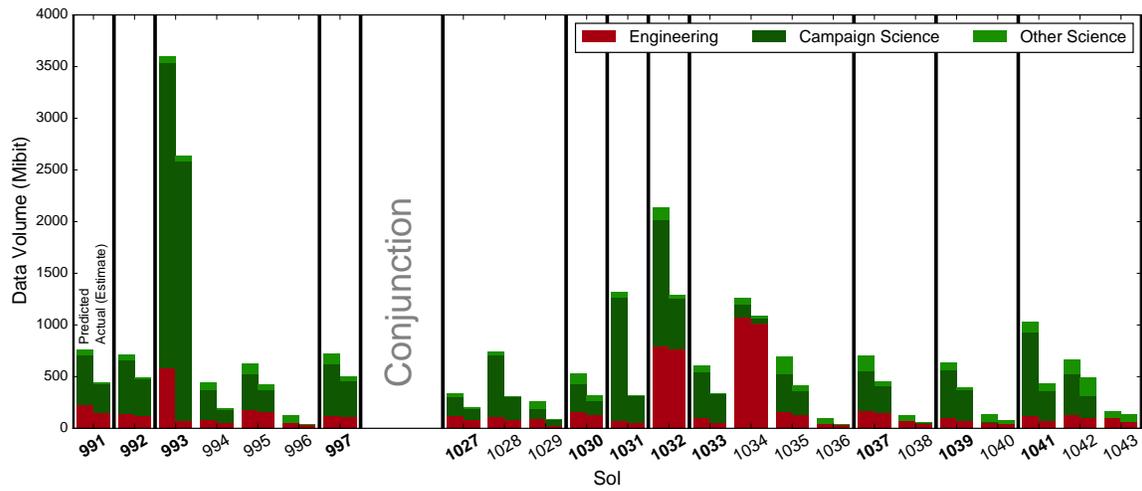
As with the other campaign warmup heating plots, the plot shows that pre-heat duration is frequently over-modeled. The sols with the largest amount of heating, and similarly the largest discrepancies between predict and actual, were in the plans for Sols 997 and 1041. Sol 997 included early morning remote sensing so there was a larger amount of heating for the colder morning temperatures. Sol 1041 ended up including predicted and actual warmup durations for heating activity that began in Sol 1041's plan but spanned the handover into Sol 1042. Because the activity began in Sol 1041, the plot accounts for the duration in the Sol 1041 plan. This explains the larger amount of pre-heat duration for that plan as it includes some heating from Sol 1042.

Figure 69 shows how vehicle energy was allocated during the campaign. As with the other campaigns, the shape of the plot roughly matches the corresponding RCE duration allocation plot from Figure 65. The main deviation is that Sol 1042 had high energy use with relatively low RCE duration. This is because Sol 1042 included a SAM analysis. SAM activities are typically high energy and can be performed while the RCE is off.

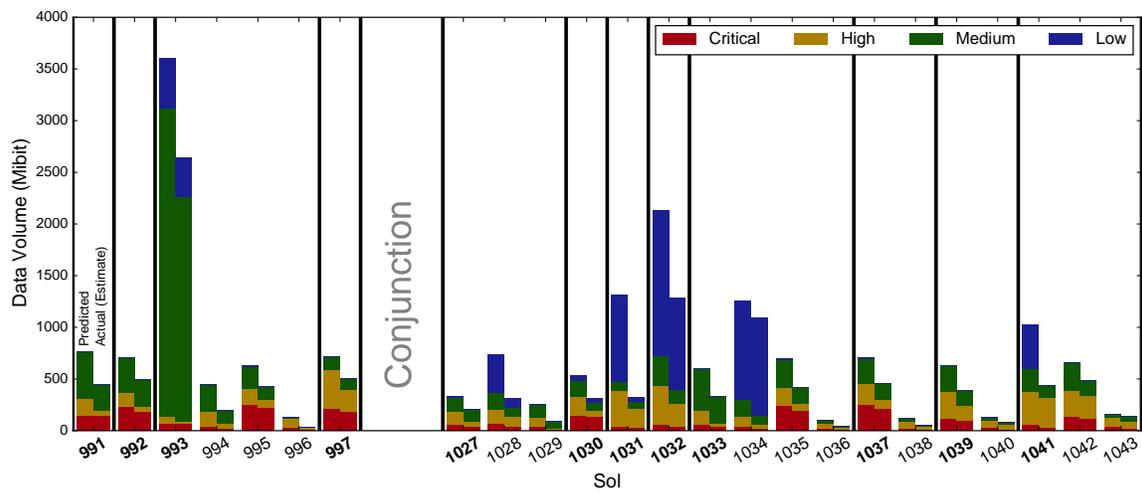
In Figure 70, we look at the predicted and actual state of charge for the rover's battery during the campaign. The plot indicates that energy was generally not a limiter for the campaign activity. For most of the campaign, the minimum state of charge was well above the 40% minimum limit and the handover state of charge values were well above the supratactical guidelines.

The plot does indicate that energy constraints may have become a factor during the multi-sol plan spanning Sol 1041 through Sol 1043. The plan included overnight APXS on Sol 1041 and high energy activities with SAM on Sol 1042 and CheMin on Sol 1043. However, even though the predicted state of charge values were low, the team was able to fit in their intended activity without descope activity.

There was one case during the Marias Pass campaign in which an activity was deferred due to energy levels. A SAM activity considered for the Sol 997 plan was deferred due to concern that if the activity was included, the handover state of charge into Sol 998 would not be sufficient.



(a) Allocation of data volume by intent category



(b) Allocation of data volume by downlink priority

Figure 66: Allocation of data volume for Marias Pass campaign.

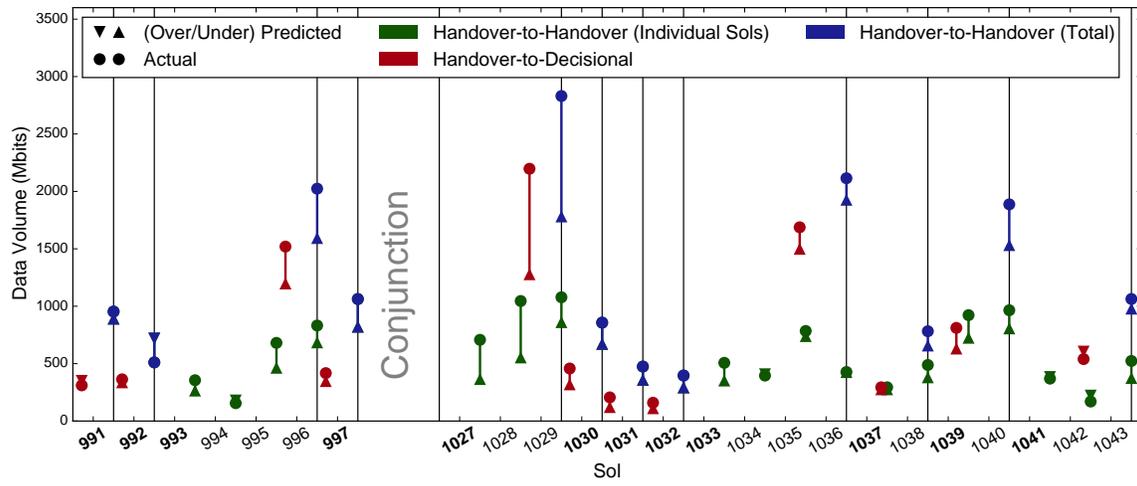


Figure 67: Downlink performance for Marias Pass campaign.

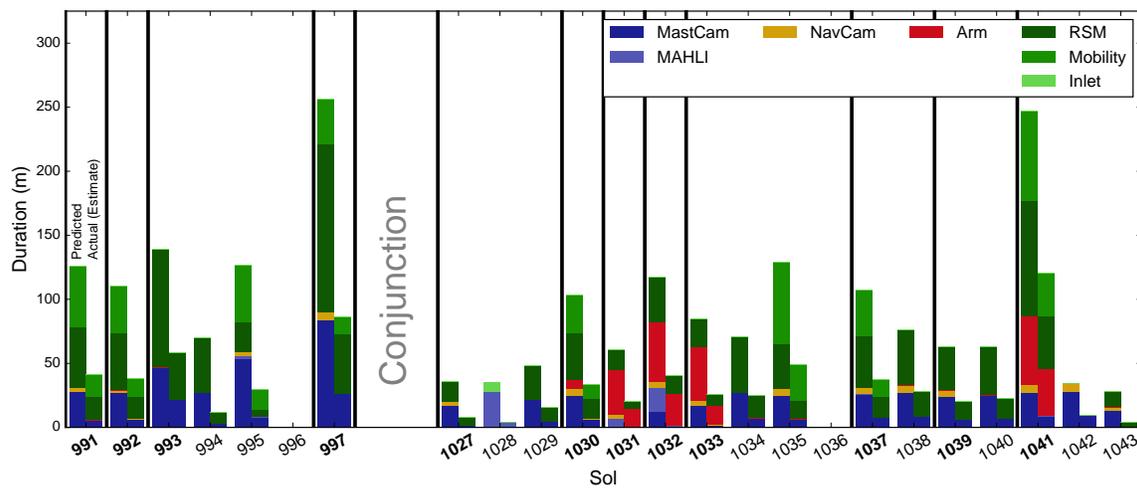


Figure 68: Summary of device warmup activity for Marias Pass campaign.

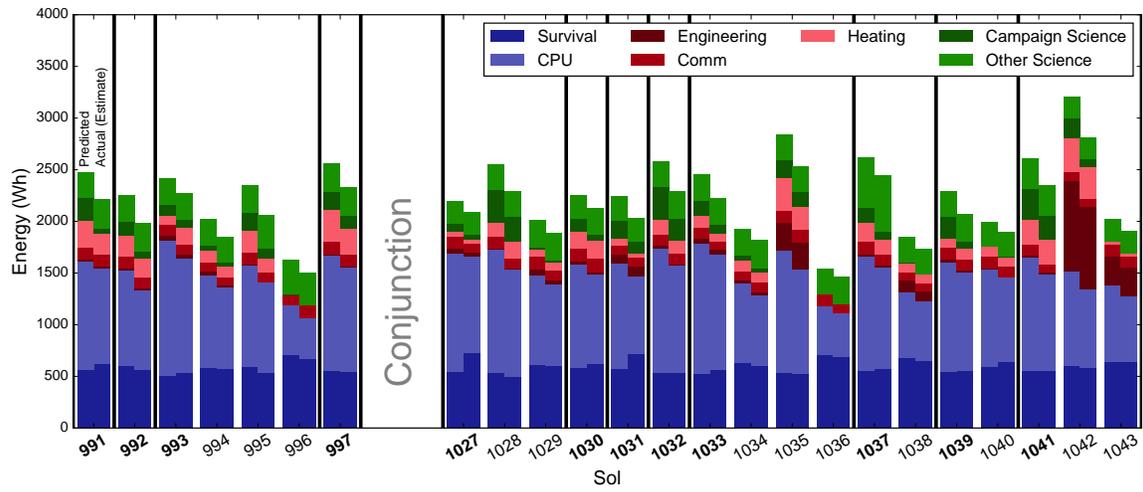


Figure 69: Allocation of energy for Marias Pass campaign.

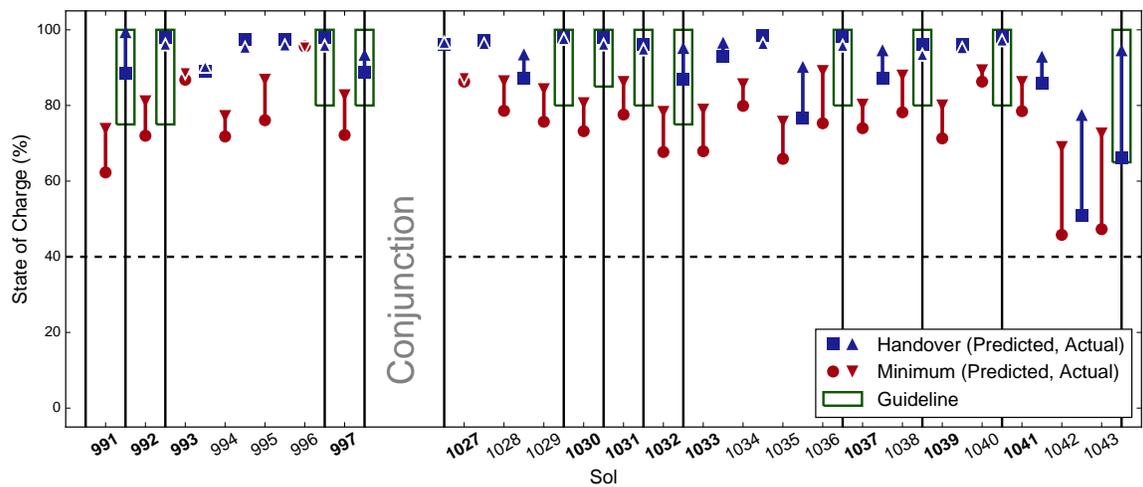


Figure 70: Summary of battery state of charge for Marias Pass campaign.

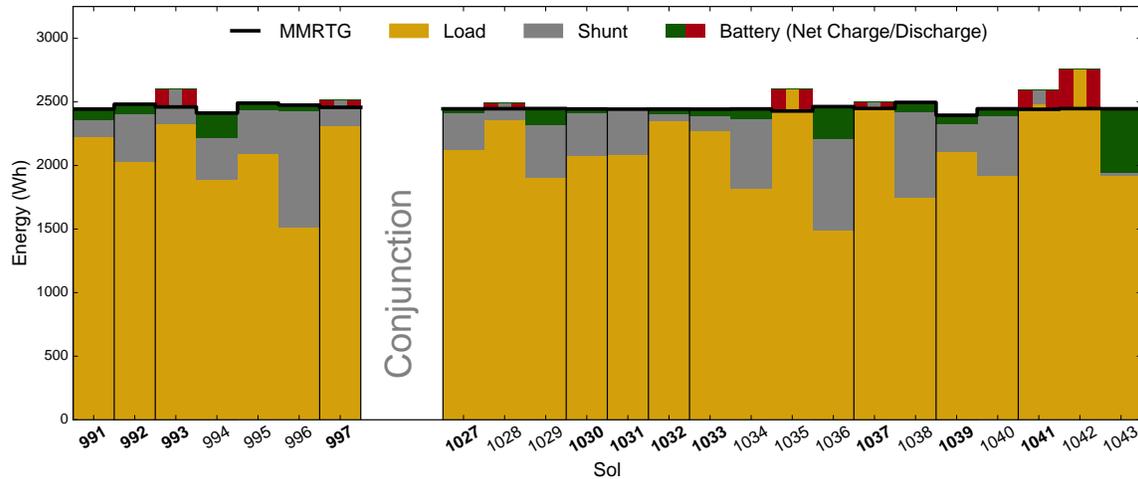


Figure 71: Summary of battery and MMRTG energy for Marias Pass campaign.

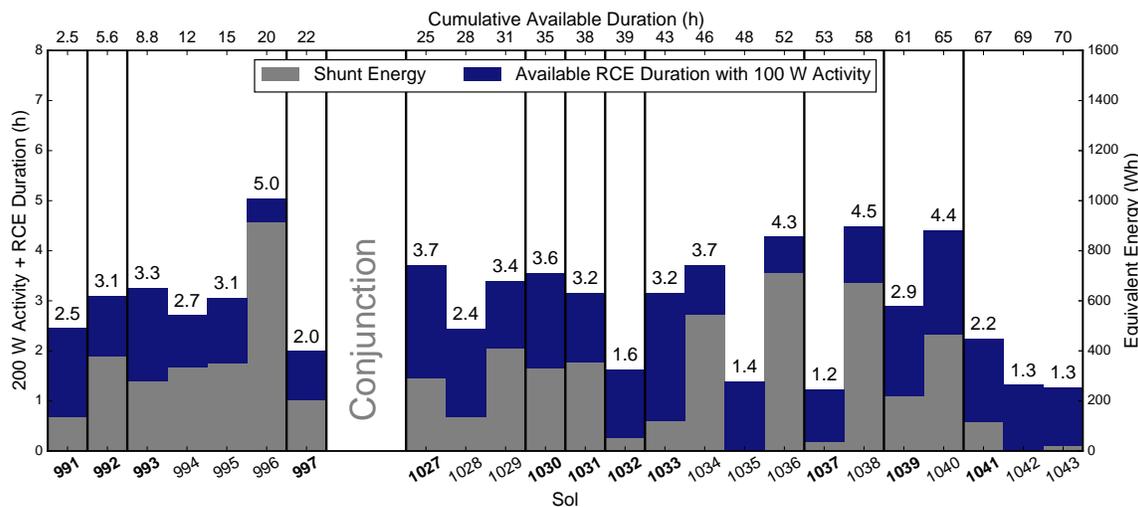


Figure 72: Estimate of extra duration availability for Marias Pass campaign.

As with previous campaigns, Figure 70 shows how actual state of charge values are often higher than predicted values, illustrating the challenge and productivity impact of resource predictions expressed by interview participants in Section 3.3.

Figure 71 shows the how energy from the MMRTG was used during the campaign. The campaign included a few cases with a net drain of the battery. Four sols had relatively higher net drains of the battery. Sol 993 included higher duration of RCE awake time due to a periodic thermal maintenance activity which required the rover to wake up several times throughout the sol. Sols 1035 and 1042 included a CheMin analyses. Sol 1041 included a SAM analysis. Three sols had a slight drain of the battery. Sol 997 included an early morning imaging activity, which required extra heating duration due to the colder morning temperatures. Sol 1028 was a busy sol of contact science with overnight APXS integration and also resulted in a slight drain of the battery. Sol 1037 included a CheMin analysis.

Figure 72 shows the estimated extra duration that would have been available for remote sensing or driving given the amount of unproductive RCE duration from Figure 65 and shunt energy from Figure 71. Over the 24 sols of the campaign, there was a total of 70 hours of estimated extra duration.

Figure 73 shows the amount of unused daylight hours during the Marias Pass campaign. It indicates that there was sufficient unused daylight duration to make use of the extra duration from Figure 72.

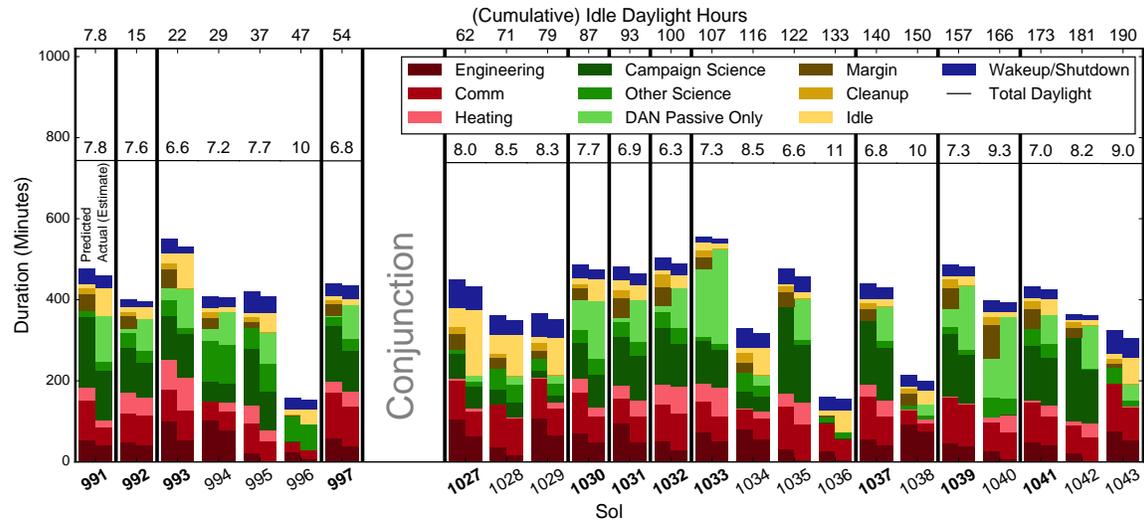


Figure 73: Allocation of RCE duration during daylight hours for Marias Pass campaign.

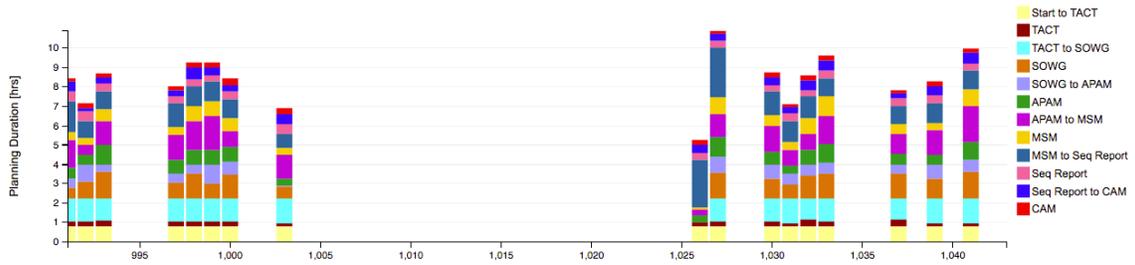
Figure 74 shows the tactical timeline durations for the Marias Pass campaign. During this time of the mission, the nominal tactical duration was 9.5 hours. The plot shows that there were two cases in which the team exceeded the nominal timeline duration. Both of these occurred on multi-sol plans. Sol 1027 was a four-sol plan, using runout for the fourth sol, and Sol 1041 was a three-sol plan. The problems related to data management sequences described in Section 6.1.2 contributed to the long duration of the Sol 1027 tactical shift. The sol 1041 plan included both arm and drive Rover Planner activities. The drive included added complexity to support mid-drive imaging.

6.3 Significant Productivity Factors

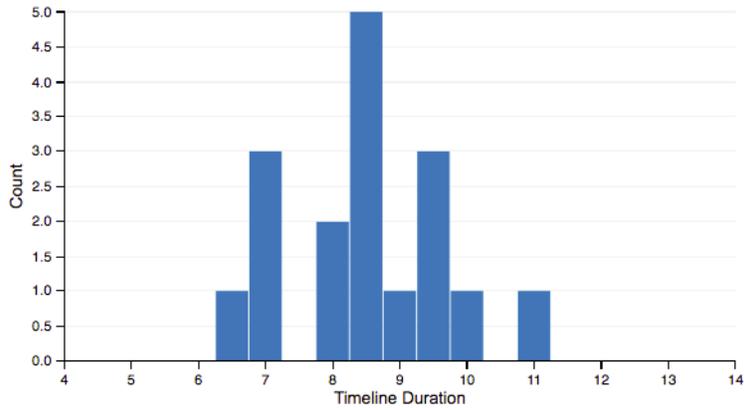
Figure 75 shows the Tom Sawyer Perspectives visualization of the data collected from the Marias Pass campaign. Please refer to Section 4.3 for a description of the diagram nomenclature. As with the other campaign visualizations, Figure 75 helps identify sols in which there was high levels of productivity and those that had lower levels of productivity. The Marias Pass visualization is also interesting in terms of the links between sols, denoting how data from one sol informed the development of activity for a subsequent sol. The diagram shows cases where sols early in the campaign, e.g. Sols 991 and 992, generated data that directly informed activities in sols much later in the campaign, e.g. Sols 1035, 1037, 1039, 1041 and 1042. This is because on those later sols, the team had decided to return to the area the rover had been during Sols 991 and 992. As such, they were able to make use of data collected from those previous sols.

Table 4 provides a breakdown of the sols of the campaign in terms of how they contributed toward accomplishing the campaign objectives. Sols labeled “Campaign” were those that directly contributed to the campaign objectives with remote sensing and/or drives. “Campaign Multi-Sol” sols are those in which significant activity was performed toward the campaign objectives as part of a multi-sol plan, either due to a weekend or restricted planning. The reason for calling these sols out separately is that the presence of the multi-sol plan limited the team’s options for these sols. For example, had there not been a multi-sol plan, the team may have opted for to move up activity that was performed in a subsequent plan (e.g. a drive activity) which would have reduced the overall number of sols required to achieve the campaign objectives. The “Extra Drives” label denotes sols in which unexpected drives were required. “Post-Drive Multi-Sol” sols were those sols in which the team was not able to achieve substantial campaign objectives due to lack of data following a drive during a multi-sol plan. Finally, “Runout” are sols of very low activity that used in cases the team had to create multi-sol plans but the tactical timeline capacity did not allow for sufficient time to develop activities for all sols of the plan.

It should be noted that Table 1 is a coarse breakdown of sols intended to characterize the productivity factors that impacted achieving the campaign objectives. For example, even sols labeled “Extra Drives” still contained activity that contributed toward the campaign goals as the MSL team is adept at exploiting opportunities even when plans do



(a) Tactical durations per planning sol



(b) Tactical duration histogram

Figure 74: Tactical timeline durations for Marias Pass campaign.

Sol Type	Sols	Number of Sols
Campaign	992, 993, 997, 1027, 1030, 1031, 1032, 1033, 1039, 1041	10
Campaign Multi-Sol	994, 1028, 1029, 1034	4
Extra Drives	991, 995, 1035, 1037, 1042	5
Post-Drive Multi-Sol	1038, 1040, 1043	3
Runout	996, 1036	2
Total Sols		24

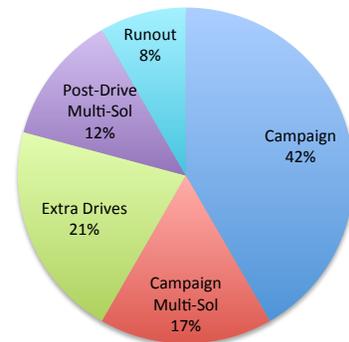
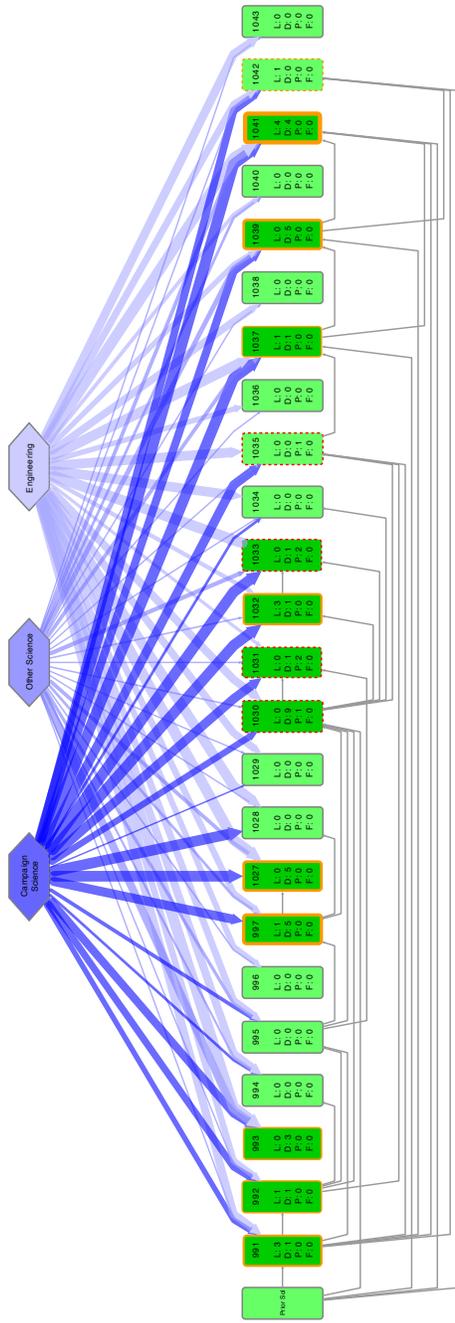
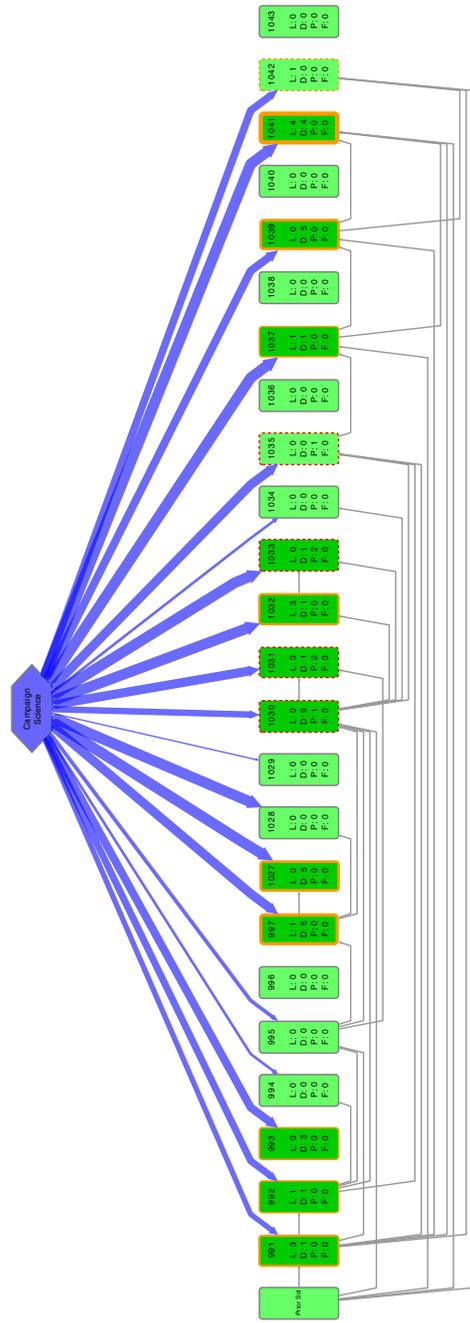


Table 4: Breakdown of sols for the Marias Pass campaign.



(a) Visualization showing all intent categories



(b) Visualization showing only Campaign Science intent category



(c) Legend

Figure 75: Visualization of the Marias Pass Campaign.

not go as expected. However, we felt it was informative to categorize such sols separately as they represent a deviation, and delays, from the intentions the team had for carrying out the campaign.

Table 4 shows that a majority of the sols were in the “Campaign” or “Campaign Multi-Sol” category. This indicates that most of the time the team was able to make direct progress toward achieving the campaign objectives. Though, it would be interesting to know if the team would have altered their activity on the sols labeled “Campaign Multi-Sol” had they planned multiple single-sol plans, rather than multi-sol plans. Multi-sol plans impose different challenges and constraints compared to single-sol plans. A multi-sol plan reduces the amount of tactical time available per planned sol which can limit the overall amount of activity across the multi-sol plan. Further, the lack of ground-in-the-loop cycles between sols places constraints on the types and ordering of activity in the plan. E.g. if the team drives the rover on an early sol, they will be unable to perform targeted remote-sensing and contact science on the subsequent sols of the multi-sol plan.

The Marias Pass campaign had a relatively low number of restricted sols given the total number of sols. This is because Solar conjunction started just as mission was entering restricted sols (Sol 1003). When the team returned from conjunction on Sol 1027, the mission had recently exited the restricted sol period. As such, even though conjunction resulted in a dramatic drop in productivity, it had the advantage of largely occurring on sols during a restricted period. As such, only 4 sols were significantly impacted by restricted sols. These are the three sols labeled “Post-Drive Multi-Sol” and one of the “Runout” sols, Sol 1036, which also happened to be a post-drive restricted sol. We labeled it separately because a “Runout” sol represents significantly reduced productivity compared to a “Post-Drive Multi-Sol” sol.

Table 4 indicates that “Extra Drives” were the largest factor in delaying the team in completing the campaign objectives. This was due to the challenging terrain at Marias Pass. Occlusions in the terrain limited the viewshed available for the Rover Planners to plan routes which resulted in two drives, Sols 991 and 992, being limited from their intended objectives. In the case of Sol 992, despite the limitation, the science team concluded the rover ended up sufficiently close their desired location that we did not count Sol 992 as an “Extra Drive”. Two drives ended early due to visual odometry convergence failure, Sols 1035 and 1042. One drive, Sol 995, was limited due to the complexity of avoiding sand ripples and resulted in an additional drive to reach the intended goal. Finally, the Sol 1037 drive resulted in the rover being placed in a position with insufficient stability and required an additional drive to re-position the vehicle.

Although there were only two “Runout” sols, they are significant because they represent a major drop in productivity for those sols. The original intent of runout is to have the vehicle perform a highly limited and safe set of activities in the event that something prevent the team from uplinking a new set of command products. As such, they are typically only employed in unexpected, off-nominal situations. However, the two runout sols in Table 4 were intentionally planned by the team as in situations when the team needed to create plans that covered a larger range of sols than could be effectively planned in the duration of the tactical timeline. This was due to needing two cover two different holiday weekends (Sol 996 over the Memorial Day weekend and Sol 1036 over the Fourth of July weekend).

Although we did not cover it in our case study, Solar conjunction is an obvious point of low productivity for the mission. Solar conjunction itself spanned 22 sols. In addition, four sols leading up to conjunction were highly limited in their activity to reduce the risk of a problem prior to the command moratorium. Though it only occurs once every two years, such a drop in productivity over the span of 26 sols is highly significant.

Looking at activity within each planning cycle provides additional insight into productivity factors affecting the campaign. Unlike Pahrump Hills, the team did not prepare an extensive strategic plan to guide the Marias Pass exploration. This was at least partly due to the nature of the terrain at Marias Pass. The elevated valley prevented the acquisition of a comprehensive panorama of the area ahead of time. However, the daily activity of the team demonstrates that they were able to quickly react on the tactical timeline to newly arriving data from the vehicle. This included quickly identifying targets for remote sensing and contact science as well as responding to problems such as drives ending early or not achieving sufficient stability.

Although the campaign was intended to employ a linear approach to exploration, as discussed in Section 6.1, rather than a walkabout approach, the campaign ended up including a significant amount of activity which required the rover to re-trace its steps and re-visit previous locations. One of the scientists who participated in the interviews speculated that had a walkabout approach been employed, some overhead in conducting the campaign may have been avoided.

The drive on Sol 1030 is a good example of the complexities involved in planning an apparently simple drive activity. In this case, only a small change in the rover’s position, about 1 meter, was required to get the rover in position for contact science. However, there was a partially conflicting requirement to position the rover such that

ChemCam could be used on the target. The ChemCam has a minimum operating distance which required the rover to be set back a little further from the target than would otherwise be required for contact science. Additionally, the team needed to ensure that sufficiently imagery was acquired to perform stability analysis of the rover's final position. As demonstrated with the Sol 1035 drive, being able to verify the vehicle's stability is essential for enabling contact science. The complexities of this short drive were sufficiently challenging that the tactical Rover Planners solicited the assistance of the supratactical Rover Planners to support the analysis.

The campaign included a few instances in which activities were repeated due to the initial attempt not fully meeting the team's objectives. In addition to the already discussed drive challenges, there were cases in which Mastcam and ChemCam observations were re-done. A Mastcam observation on Sol 1030 ended up being in shadow and was re-taken on Sol 1032. The Sol 1033 plan included a ChemCam observation to attempt to high light-toned clast material that was not sufficiently accomplished with previous targets on a previous sol. In these cases, the team was fortunate to still be in the same area to enable the opportunity to re-acquire the data.

The data management problem that occurred in the Sol 1027 plan combined with the high volume of previously collected onboard data due to the preparation for conjunction, led to significant data volume constraints in the Sol 1030 plan. This ended up reducing the amount of activity the team could perform in that plan including activity that would have contributed toward the Marias Pass campaign. Fortunately, the restriction was short-lived.

6.4 Significant Ground-in-the-Loop Cycles

One of the motivations for conducting these case studies was so that we could understand the reasons for ground-in-the-loop cycles. Toward that end, the following summarizes significant ground-in-the-loop cycles during the Marias Pass campaign.

- Selection of locations for study of the Stimson / Murray contact. The post-drive imagery following the Sol 991 drive provided the team with the first imagery from the rover of the Stimson / Murray formation contact at Marias Pass. The team used this data during the Sol 992 planning to select the two locations, shown in in Figure 59, to be visited to study the contact. Although HiRISE data was available that showed the presence of the contact at Marias Pass, the team required the higher-resolution imagery from the rover to make this selection. The elevation of the valley prevented this type of imagery from being acquired earlier.
- Decision to accept position of rover following Sol 992 drive. The drive planned on Sol 992 was limited by available viewshed and would not reach the original selection for the first contact location in Figure 59. However, based on post-drive data from the drive, the team made the decision on Sol 993 that the rover's current position was more than sufficient to achieve their imaging needs.
- Drive planning. Each drive in the campaign was based on post-drive data collected at the end of the previous drive. The post-drive imagery allowed the team to identify hazards and evaluate routes for reaching the desired end-of-drive locations. Sols in which drives were planned were: Sols 991, 992, 995, 997, 1030, 1035, 1037, 1039, and 1042. The drive on Sol 993 was particularly challenging due to the trade-off between a straight route that would cross sand ripples, with a higher risk of slippage, versus a circuitous route that would take longer and likely require an additional sol of driving to reach the desired goal.
- Stability analysis for contact science. Prior to deploying the arm and performing contact science, the team must use data from the rover's current position to assess the vehicle's stability. If the vehicle were to slip during with the arm deployed, it could result in collision with the terrain with potential damage to the arm and its instruments. This analysis was performed on Sols 998 (not covered in our case study), 1027 (for contact science on Sol 1028), 1031, 1032, 1039 (in which stability analysis concluded contact science could not be performed), and 1041.
- Selection of targets for Mastcam, ChemCam and contact science. Ground-in-the-loop decisions were made for specific targets for Mastcam and ChemCam along with the selection of targets for study with arm-mounted instruments MAHLI and APXS. These types of decisions were made on nearly every sol of the campaign so we refrain from listing them all here, instead pointing out some particularly significant examples. During Sol 997 planning, the team selected the targets at the contact, Missoula and Ronan, that would be studied with arm-mounted instruments. During Sol 1030 planning, the team began using imagery of the Missoula target to plan

the specific contact science observations they would perform such as the planned MAHLI mosaic and potential placements for APXS integrations. This was followed up on Sol 1032 when the actual contact science was planned with refinements to the placement of the APXS instrument. The team also acquired additional Navcam imagery on Sol 1032 to assist in the planning of a subsequent Mastcam mosaic. There were several plans in which the team took advantage of the rover's interaction with the environment to study rocks that had been broken apart by the rover's wheels.

- Unexpected discoveries. One of the most significant ground-in-the-loop decisions was made during the long break for conjunction. Separate members of the team reported on distinct findings, high levels of silica from ChemCam observations and high levels of hydrogen from DAN observations. The team realized they were discussing results from the same location, the Sol 991 location, and made the decision to backtrack to this area after conjunction.

6.5 Conclusions

Although Marias Pass was not the team's first choice for exploring the contact between the Pahrump and Stimson formations, it proved to be a good alternative for studying the contact. Targets such as Missoula contained material from both formations and provided clues to help scientists understand the history of the area. The campaign also uncovered surprising results with the detection of high silica in the Elk target and nearby elevated hydrogen readings from DAN.

The terrain in the area posed a variety of challenges for the campaign. The hill to climb to Marias Pass and obstructions at the top of the hill restricted the viewshed available for drive planning. This resulted in limitations to the drive distances on Sols 991 and 992. A sandy bowl prevented the team from a close approach of the preferred contact location, as seen in Figure 61, due to concerns about slippage. The sandy area also caused delays in reaching the alternate contact location due to the need to circumvent the ripples. There were also two cases of drives ending early due to failures with visual odometry. Based on the sol breakdown in Table 4, the extra drives resulting from these issues had the large negative impact in terms of number of sols to complete the campaign.

Although the campaign had relatively few restricted sols, due to the timing of conjunction which occurred during what would have been a restricted period, there were two holiday weekends during the campaign. This prompted the team to employ low-productivity runout sols to manage the complexity of the resulting 4-sol plans.

Our analysis of estimated extra time available due to shunting and non-productive RCE duration, from Figure 72, combined with the unused daylight hours from Figure 73 shows that there was opportunity to conduct additional remote sensing and driving during the sols of the campaign, which could potentially reduce the total number of sols required to complete the objectives.

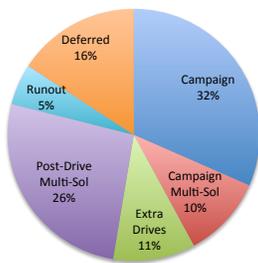
7 Discussion

The selection of campaigns for the case study provides an interesting basis for comparisons. All three were selected as examples of campaigns that emphasized driving and remote sensing. Two of the campaigns, Pahrump Hills Walkabout and Marias pass had a similar objective of characterizing a geological area, while Artist's Drive had the objective of performing a strategic drive with science imagery along the way. Though the Pahrump Hills and Marias Pass campaigns had similar high-level objectives, they were conducted within very different contexts. The terrain at Marias Pass was generally more challenging than at Pahrump Hills. In addition, the team chose to employ very different exploration strategies for these two campaigns, using a walkabout approach for Pahrump Hills and a linear approach, which is more conventional for rovers, for Marias Pass.

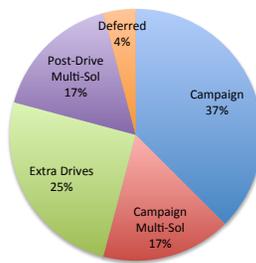
We begin our comparison of the three campaigns by summarizing the sol-by-sol breakdown of campaign activity in Table 5. Despite having different high-level objectives, the sol breakdown for Artist's Drive and Marias Pass appear to be the most similar. This is due to these campaigns having a similar number of restricted plans and both being conducted in similar, challenging terrain conditions.

Comparing the sol breakdown for Pahrump Hills with the other two campaigns shows that restricted sols have a major productivity impact for these types of campaigns. Table 6 shows the number of nominal vs. restricted shifts for each campaign. Pahrump Hills had a total of 9 tactical shifts of which 7 were during restricted periods of the mission. In contrast, Artist's Drive and Marias Pass had more total shifts and fewer restricted shifts than Pahrump Hills.

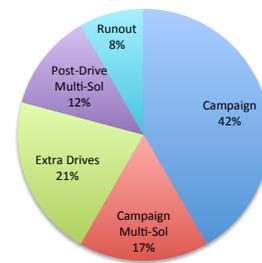
Sol Type	Pahrump Hills	Artist's Drive	Marias Pass
Campaign	6	9	10
Campaign Multi-Sol	2	4	4
Extra Drives	2	6	5
Post-Drive Multi-Sol	5	4	3
Deferred	3	1	0
Runout	1	0	2
Total Sols	19	24	24



Pahrump Hills



Artist's Drive



Marias Pass

Table 5: Breakdown of sols for all campaigns.

Sol Type	Pahrump Hills	Artist's Drive	Marias Pass
Nominal Shifts	2	16	12
Restricted Shifts	7	4	4
Total Shifts	9	20	16

Table 6: Summary of shift types for all campaigns.

Sol Type	Pahrump Hills	Artist's Drive	Marias Pass
Nominal	5	6	3
Viewshed Limited	2	5	2
Drive Fault	0	1	2
Insufficient Stability	0	0	1
Total Traverses	7	12	8

Table 7: Summary of traverses for all campaigns.

The reason for the differences in number of restricted sols between Pahrump Hills and the other campaigns is largely luck of campaign timing. The Pahrump Hills Walkabout campaign happened to begin just as a restricted period was about to start. On the other hand, the Artist's Drive campaign began just after a restricted period had ended. Marias Pass began toward the end of a nominal period but solar conjunction began before the restricted period began. By the time conjunction was over and the team returned to operations, the restricted period had completed. Thus, it is only toward the end of the Marias Pass campaign that another restricted period impacted operations.

Table 6 also highlights the significance of terrain for these types of campaigns. Table 7 summarizes the traverses performed in each campaign. Note that the Artist's Drive traverse marked as "Drive Fault" was also limited by viewshed. Rather than double-count it, we counted it as "Drive Fault" and not "Viewshed Limited". The increased terrain occlusions encountered during the Artist's Drive campaign lead to a larger number of traverses being limited by viewshed than encountered during the Pahrump Hills Campaign. Although the terrain at Marias Pass was more challenging than at Pahrump Hills, it had the same number of traverses limited by viewshed. This is likely because the Marias Pass campaign included returning to previously explored areas, allowing the team to make use of terrain imagery collected on previous sols, as shown in Figure 75. The more challenging terrain of Artist's Drive and Marias Pass resulted in drive faults and rover stability issues in the associated campaigns.

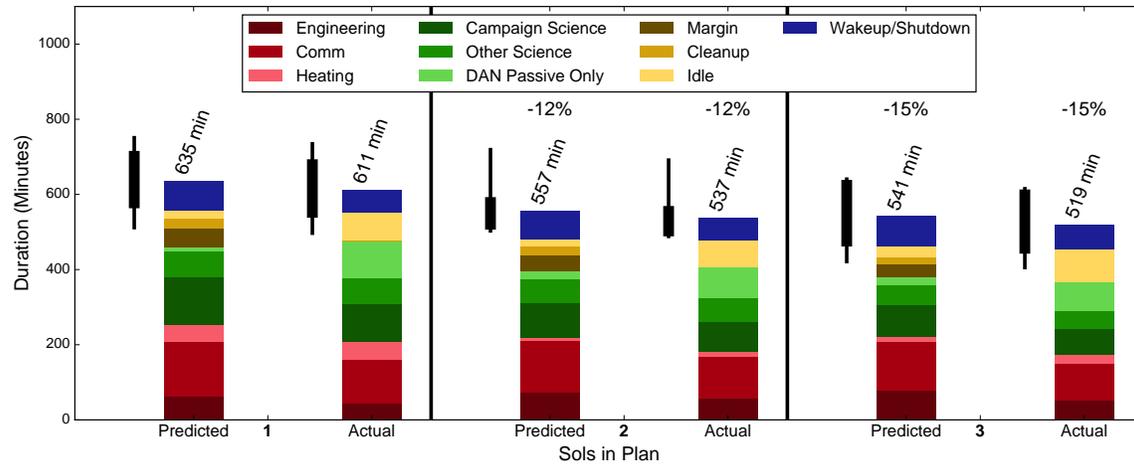
It is interesting to compare the Pahrump Hills and Marias Pass campaigns as they had similar high-level objectives but were conducted with different exploration strategies. Unlike the Pahrump Hills Walkabout, the Marias Pass campaign did not have an extensive strategic plan to direct activity. This was largely due to the geography of the Marias Pass valley. HiRISE imagery provided a high level overview of the region, with sufficient detail to indicate that the area included a promising contact between the Stimson formation and Murray formation, but contained insufficient detail to form a strategic plan for exploring the location. Because the valley was elevated above the Artist's Drive route the rover had been following, it was not possible to obtain the same type of Mastcam panorama that was available for planning the Pahrump Hills Walkabout.

Despite the absence of a detailed strategic plan for the Marias Pass campaign, the team was able to make quick tactical decisions and respond to new data as it arrived such as identifying drive routes and selecting key science targets. This can explain the why Table 5 shows a comparable number of campaign-oriented sols as Artist's Drive. In other words, it seems that the number of restricted shifts and terrain challenges was a bigger factor than the availability of a guiding strategic plan, given the team's ability to react.

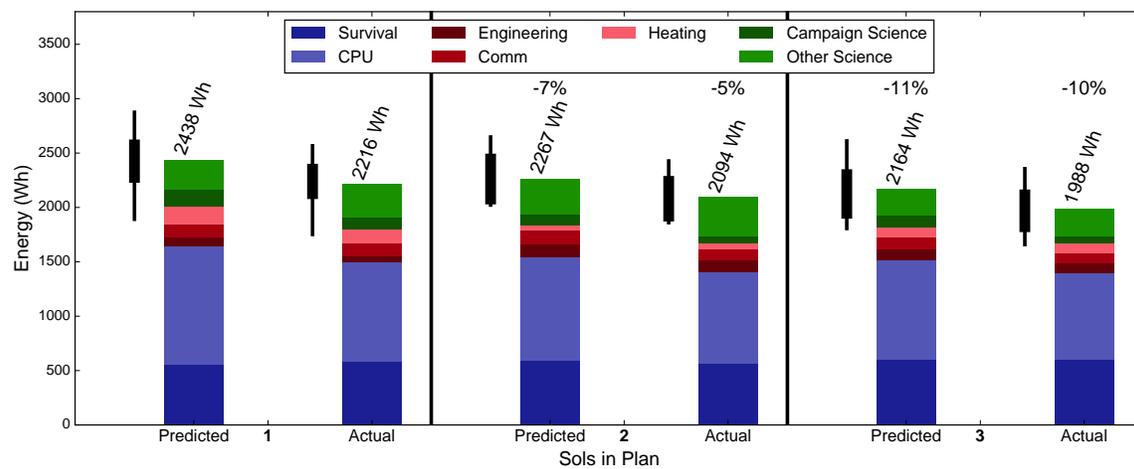
It is also interesting to compare the walkabout approach employed at Pahrump Hills vs. the linear approach used at Marias Pass. Although the team intended to use a linear strategy at Marias Pass, they ended up backtracking to explore data collected near the Sol 992 location. There was additional backtracking in the sols that followed the end of our case study sol range. It was suggested by one of the scientists in our interviews that perhaps the Marias Pass campaign would have been overall more efficient had it employed a walkabout approach.

A full assessment of the benefits of these two exploration strategies is beyond the scope of a single case study. The interested reader is referred to Yingst et al. for additional discussion on this topic [Yingst et al.(2015)]. Their conclusion is that a walkabout approach can take more time to execute, but has the potential for achieving higher quality results. One of the objectives we have with this case study is to leverage what we have learned from these productivity challenges to identify flight and ground approaches that can reduce the overhead of employing a walkabout approach.

The analyses of resource allocations followed a similar pattern for each of the campaigns. First, the analysis showed a general decrease in overall activity across multi-sol plans compared to an equivalent number of single-sol plans. To quantify this reduction in resource allocation, we plotted the per-sol allocation of RCE duration and energy for 1-sol, 2-sol and 3-sol plans. Figure 76 shows the results. The results show that, on average, there is a reduction of 12% RCE duration use and 7% reduction of energy use per sol for 2-sol plans. Usage is further reduced for 3-sol



(a) Per-sol allocation of RCE duration



(a) Per-sol allocation of energy

Figure 76: Per-sol allocation of resource allocations for 1-sol, 2-sol and 3-sol plans.

plans, we a 15% and 11% reductions for RCE duration and energy when compared to single sol plans. This is a reflection of the limited capacity of the tactical timeline to develop and validate command products to make use of vehicle resources.

In addition, the team is limited in the types of activities that can be performed in multi-sol plans due to the lack of ground-in-the-loop cycles between the sols. For example, if the rover drives, it will limit the types of activity that can be performed after the drive, such as targeted remote sensing, until the following ground-in-the-loop cycle.

This latter constraint is also the reason that these campaigns exhibited the pattern of campaign-related activity being significantly reduced on sols following drives in multi-sol plans. In many cases, there was no campaign activity on the subsequent sols. In other cases, there was a small amount of activity due to the use of ChemCam blind targeting.

Energy did not play a significant factor in any of the campaigns. This is partly due to the Martian seasons in which the selected campaigns were performed. While our selected campaigns included three different seasons, they did not include Winter. It would be informative to perform a case study of a campaign conducted during a Mars winter to assess the impact to productivity that energy might have at this time of the Mars year, when more energy is required for heating.

All of the campaigns we studied had unused shunt energy and non-productive RCE duration that could have been used for additional campaign activities. We estimated an additional 72 hours, 62 hours and 70 hours of campaign

activity could have been performed during the Pahrump Hills, Artist's Drive and Marias Pass campaigns, respectively.

As part of the case study, we were also interested in understanding the types of decisions that were made with ground-in-the-loop cycles. Following is a summary of common types of ground-in-the-loop decisions across the campaigns:

Selecting targets for ChemCam, Mastcam and contact science: While distant imagery of terrain provided sufficient information to indicate the value of traversing to an area, the scientists required the higher quality imaging of the area, obtained when the rover arrives at the site, to select specific targets.

Drive planning: Post-drive imagery is also used to provide the data necessary to plan the next traverse, including allowing the scientists to refine their selection of the end-of-drive location and the engineering team to design a route for the rover to follow.

Stability assessment for contact science: Prior to deploying the arm and performing contact science, the team must use data from the rover's current position to assess the vehicle's stability.

Responding to problems in activity executions: It is a complicated mission and plan execution does not always go as expected. Unexpected terrain conditions can cause a drive to end early, resulting in the engineering team assessing the reasons for the problem and re-planning the drive. There are also cases where remote sensing observations do not work as expected. During the campaigns there was a case when imagery was re-acquired due to lighting issues with the first attempt, and cases where ChemCam observations of extremely small features needed to be re-acquired when previous attempts missed.

8 Conclusions

A successfully deployed Martian rover represents an immensely valuable asset. The mission has a strong interest in getting the most out of the vehicle to increase the return on investment. This is further motivated by the fact that the rover's capabilities will inevitably degrade over time. As such there is a strong interest in enhancing the productivity of future surface rover missions.

We conducted the case study of MSL campaigns with the objective to better understand the productivity challenges facing surface missions. The study included interviews with mission scientists and engineers along with detailed study of three science campaigns.

The responses from the interview participants and our analysis of the campaigns showed that there are opportunities for increasing surface mission productivity. In particular, we observed that it is often the case that the vehicle has more available resources than the operations team is able to use. The case study identified a variety of issues that are limiting the productive use of these resources.

Perhaps the largest factor observed in our study of MSL campaigns was due to restricted sols and, more generally, the reliance on ground-in-the-loop to inform a large portion of the rover's activities. This reliance on ground-in-the-loop places constraints on when certain types of activities can be performed. Activities that generate data needed to make decisions for the next shift must be performed prior to the decisional communication pass. Similarly, activities that change the state of the rover in such a way as to invalidate that decisional data (e.g. driving the rover to a new location) cannot be performed after the decisional pass.

The ground-in-the-loop reliance is expected to become an even more significant liability to surface missions as the fleet of aging sun-synchronous orbiters are replaced with non-sun-synchronous orbiters. There are important science motivations for non-sun-synchronous orbiters, such as studying the Recurring Slope Lineae (RSL). However, such an orbit is not well suited for relay operations as it does not provide the same consistent pattern of afternoon passes that current operations relies on to get "end of day" state knowledge of the rover. Instead, there could be many cases in which the communication pass for the plan occurs late in the Martian night or in the morning soon before handover to the next plan. Such cases would result in the ground team having very little time to develop command products in response to the data in time for the next uplink. The result is that the mission would have a much larger number of restricted sols.

The next biggest productivity factor across all campaigns was the need for additional drives to reach objectives due to terrain interactions. There were several cases in which drives were limited due to available viewshed for route planning. In other cases, drives faulted out early due to issues such as visual odometry failure.

The interview participants identified additional significant productivity factors. There was large consensus that effective communication amongst team members was key to productive operations. Several participants agreed that the complexity and overhead of producing command products was a primary factor in explaining why the team often does not use all available vehicle resources. The scientist participants also emphasized the importance and challenge of engagement and situation awareness along with the need for time to analyze downlink results and consider options for activities.

In summary, the following is a list of the significant productivity factors identified in the case study:

Ground-in-the-loop requirements for target selection and effective drive planning: This results in a significant drop in productivity on sols that follow drives during restricted periods of the mission. Even during non-restricted sols, it constrains the timing of activity that can change the state of the vehicle and activity that acquires decisional data to occur prior to the decisional pass.

Capacity of tactical timeline to fill multi-sol plans: Due to the time required to develop and validate command products, the amount of overall activity across a multi-sol plan is generally lower than the amount of activity across a similar number of single sol plans.

Ground-in-the-loop requirements to respond to outcome of activity: We observed several instances where the team decided to re-do an activity, or return to a previous location, after observing the data received from the vehicle. This included re-doing an APXS integration during Artist's Drive, re-acquiring Mastcam and ChemCam observations during Marias Pass, and returning to a previous location in response to interesting instrument readings during Marias Pass. There were also several cases where drives ended early due to unexpected terrain conditions and had to be re-planned.

Use of margin and cleanup duration allocations: Because there is uncertainty in the actual run time of activities, ground operations allocates margin to allow activities to run long without impacting future activities. And to protect against activities exceeding margin, additional time is reserved for cleaning up activities. This results in non-productive time when activities run within predicted durations.

Ability to Exploit Supratactical Work: If the tactical team follows the supratactical plan, they are able to leverage work performed on the supratactical timeline and productivity benefits. If decisional data prompts the tactical team to change course, they may need to start from scratch and are often not able to accomplish as much.

Overhead in Command Product Generation: The effort involved in developing and validating command products can limit the amount of activity that can be accomplished.

Managing Tactical Timeline Complexity: Both the supratactical and tactical teams must judge how much activity can be managed within the tactical timeline. The ability to accurately make this judgment can impact productivity.

Predicting Available Vehicle Resources: Inaccuracies in resource modeling, including activity power and duration requirements, can result in unnecessarily restricting planned activity.

Interpersonal Communication: Effective communication was identified as a significant productivity factor by many of the participants in each of the operations roles we interviewed. This includes communication between the science and engineering teams and among roles within science and engineering teams.

Science Team Engagement: It is important for science team members to be aware of the current and past context of the mission to make informed science decisions. A significant challenge to engagement is that many team members work on the mission part time.

Time to Analyze Data, Make Decisions: The time allocated to the scientists in the tactical timeline to analyze data, make decisions and develop plan fragments for their activities is a significant factor in productivity.

As discussed in Section 2.4, we focused our study on campaigns that emphasized driving and remote sensing. Though two of the campaigns we chose included contact science activity, there would be value in studying campaigns that made an emphasis on contact science as well as sampling campaigns. We believe many of the productivity factors

identified in the current study are relevant for contact science and sampling campaigns, but there are likely to be additional important productivity factors identified in these other campaign types.

Our next objective is to identify changes to flight systems and ground operations practices to overcome these challenges and enable high levels of productivity for future surface missions. The findings from this study will guide the design and development process by helping to define the capabilities required to meet these productivity challenges. We will also leverage examples from the campaigns studied to define scenarios that will be used to focus the development and to evaluate the performance of our work.

Acknowledgments

This research was conducted at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This work was funded by the Jet Propulsion Laboratory Research and Technology Development program.

We are extremely grateful to the MSL team for their support of this project. The generous assistance they provided in enabling access to MSL data, answering our many questions and supporting our interviews was invaluable to this study.

References

- [Banham et al.(2016)] Banham, S., Gupta, S., Rubin, D., Watkins, J., Sumner, D., Grotzinger, J., Lewis, K., Edgett, K., Edgar, L., & Stack, K. (2016). Reconstruction of an ancient eolian dune field at Gale Crater, Mars: Sedimentary analysis of the Stimson formation. In *Proceedings of the 47th Lunar and Planetary Science Conference*.
- [Chattopadhyay et al.(2014)] Chattopadhyay, D., Mishkin, A., Allbaugh, A., Cox, Z. N., Lee, S. W., Tan-Wang, G., & Pyrzak, G. (2014). The Mars Science Laboratory supratactical process. In *Proceedings of the SpaceOps 2014 Conference*, Pasadena, CA.
- [Estlin et al.(2012)] Estlin, T., Bornstein, B., Gaines, D., Anderson, R. C., Thompson, D., Burl, M., Castano, R., & Judd, M. (2012). AEGIS automated targeting for the mer opportunity rover. *ACM Transactions on Intelligent Systems and Technology*, **3**(3).
- [Forni et al.(2015)] Forni, O., Mangold, N., Blaney, D. L., Fisk, M., Wiens, R., Meslin, P.-Y., Nachon, M., Gasnault, O., Maurice, S., Cousin, A., & Moulic, S. L. (2015). ChemCam chemostratigraphy of the Pahrump outcrop, Gale Crater. In *Proceedings of the 46th Lunar and Planetary Science Conference*.
- [Francis et al.(2015)] Francis, R., Estlin, T., Gaines, D., Bornstein, B., Schaffer, S., Verma, V., Anderson, R. C., Burl, M., Chu, S., Castano, R., Thompson, D., Blaney, D., de Flores, L., Doran, G., & Wiens, R. (2015). AEGIS autonomous targeting for the curiosity rovers chemcam instrument. In *Proceedings of the Applied Imagery Pattern Recognition (AIPR) Workshop*, Washington D.C.
- [Frydenvang et al.(2016)] Frydenvang, J., Gasda, P., Hurowitz, J., Grotzinger, J., Wiens, R., Newsom, H., Bridges, J., Gasnault, O., Maurice, S., Fisk, M., Ehlmann, B., Watkins, J., Stein, N., Forni, O., Mangold, N., Cousin, A., Clegg, S., Anderson, R., Payr, V., Rapin, W., Vaniman, D., Morris, R., Blake, D., Gupta, S., Sautter, V., Meslin, P.-Y., Edwards, P., M.Rice, Kinch, K., Milliken, R., Gellert, R., Thompson, L., Clark, B., Edgett, K., Sumner, D., Fraeman, A., Madsen, M., Mitrofanov, I., Jun, I., Calef, F., & Vasavada, A. (2016). Discovery of silica-rich lacustrine and eolian sedimentary rocks in Gale crater, Mars. In *Proceedings of the 47th Lunar and Planetary Science Conference*.
- [Gasda et al.(2016)] Gasda, P. J., Frydenvang, J., Wiens, R. C., Grotzinger, J. P., Watkins, J. A., Stein, N., Edgett, K. S., Newsom, H., Clark, B., Anderson, R., Bridges, N., Clegg, S., & Maurice, S. (2016). Potential link between high-silica diagenetic features in both eolian and lacustrine rock units measured in Gale crater with MSL. In *Proceedings of the 47th Lunar and Planetary Science Conference*.

- [Grotzinger et al.(2015)] Grotzinger, J. P., Gupta, S., Malin, M. C., Rubin, D. M., Schieber, J., Siebach, K., Sumner, D. Y., Stack, K. M., Vasavada, A. R., Arvidson, R. E., III, F. C., Edgar, L., Fischer, W. F., Grant, J. A., Griffes, J., Kah, L. C., Lamb, M. P., Lewis, K. W., Mangold, N., Minitti, M. E., Palucis, M., Rice, M., Williams, R. M. E., Yingst, R. A., Blake, D., Blaney, D., Conrad, P., Crisp, J., Dietrich, W. E., Dromart, G., Edgett, K. S., Ewing, R. C., Gellert, R., Hurowitz, J. A., Kocurek, G., Mahaffy, P., McBride, M. J., McLennan, S. M., Mischna, M., Ming, D., Milliken, R., Newsom, H., Oehler, D., Parker, T. J., Vaniman, D., Wiens, R. C., & Wilson, S. A. (2015). Deposition, exhumation, and paleoclimate of an ancient lake deposit, Gale crater, Mars. *Science*, **350**(6257).
- [Jet Propulsion Laboratory Press Release(2013)] Jet Propulsion Laboratory Press Release (2013). Curiosity team upgrades software, checks wheel wear. <http://www.jpl.nasa.gov/news/news.php?release=2013-374>.
- [Jet Propulsion Laboratory Press Release(2014a)] Jet Propulsion Laboratory Press Release (2014a). Curiosity's route from landing to Pahrump Hills. <http://mars.nasa.gov/msl/multimedia/images/?ImageID=6855>.
- [Jet Propulsion Laboratory Press Release(2014b)] Jet Propulsion Laboratory Press Release (2014b). Geologic layers of Mount Sharp. <http://mars.nasa.gov/msl/multimedia/images/?ImageID=6580>.
- [Jet Propulsion Laboratory Press Release(2014c)] Jet Propulsion Laboratory Press Release (2014c). Slopes of Mount Sharp. <http://mars.nasa.gov/msl/multimedia/images/?ImageID=6573>.
- [Jet Propulsion Laboratory Press Release(2015a)] Jet Propulsion Laboratory Press Release (2015a). Contact zone: 'Missoula'. <http://www.jpl.nasa.gov/spaceimages/details.php?id=pia19829>.
- [Jet Propulsion Laboratory Press Release(2015b)] Jet Propulsion Laboratory Press Release (2015b). Curiosity rover at Pahrump Hills. <http://www.nasa.gov/jpl/mro/pia19114>.
- [Jet Propulsion Laboratory Press Release(2015c)] Jet Propulsion Laboratory Press Release (2015c). Curiosity's traverse map through Sol 991. <http://mars.jpl.nasa.gov/msl/multimedia/images/?ImageID=7235>.
- [Jet Propulsion Laboratory Press Release(2015d)] Jet Propulsion Laboratory Press Release (2015d). Looking toward Curiosity study areas, Spring 2015. <http://mars.jpl.nasa.gov/msl/multimedia/images/?ImageID=7180>.
- [Jet Propulsion Laboratory Press Release(2015e)] Jet Propulsion Laboratory Press Release (2015e). 'Marias Pass,' contact zone of two Martian rock units. <http://www.nasa.gov/image-feature/jpl/pia20174/marias-pass-contact-zone-of-two-martian-rock-units>.
- [Jet Propulsion Laboratory Press Release(2015f)] Jet Propulsion Laboratory Press Release (2015f). Nasa's Curiosity analyzing sample of Martian mountain. <http://mars.nasa.gov/msl/news/whatsnew/index.cfm?FuseAction=ShowNews&NewsID=1777>.
- [Jet Propulsion Laboratory Press Release(2015g)] Jet Propulsion Laboratory Press Release (2015g). NASA's Curiosity eyes prominent mineral veins on Mars. <http://mars.jpl.nasa.gov/msl/news/whatsnew/index.cfm?FuseAction=ShowNews&NewsID=1798>.
- [Jet Propulsion Laboratory Press Release(2015h)] Jet Propulsion Laboratory Press Release (2015h). Nasa's Curiosity Mars rover tracks Sunspots. <http://mars.jpl.nasa.gov/msl/news/whatsnew/index.cfm?FuseAction=ShowNews&NewsID=1842>.
- [Jet Propulsion Laboratory Press Release(2015i)] Jet Propulsion Laboratory Press Release (2015i). NASA's Curiosity rover adjusts route up Martian mountain. <http://www.jpl.nasa.gov/news/news.php?feature=4596>.
- [Jet Propulsion Laboratory Press Release(2015j)] Jet Propulsion Laboratory Press Release (2015j). Quick detour by NASA Mars rover checks ancient valley. <http://mars.jpl.nasa.gov/msl/news/whatsnew/index.cfm?FuseAction=ShowNews&NewsID=1809>.

- [Jet Propulsion Laboratory Press Release(2015k)] Jet Propulsion Laboratory Press Release (2015k). Rover arm delivers rock powder sample. <http://mars.jpl.nasa.gov/msl/news/whatsnew/index.cfm?FuseAction=ShowNews&NewsID=1787>.
- [Jet Propulsion Laboratory Press Release(2015l)] Jet Propulsion Laboratory Press Release (2015l). Testing to diagnose power event in Mars rover. <http://mars.jpl.nasa.gov/msl/news/whatsnew/index.cfm?FuseAction=ShowNews&NewsID=1783>.
- [Jet Propulsion Laboratory Press Release(2015m)] Jet Propulsion Laboratory Press Release (2015m). Unfavorable terrain for crossing near 'Logan Pass'. <https://www.nasa.gov/jpl/msl/pia19662/unfavorable-terrain-for-crossing-near-logan-pass>.
- [Jet Propulsion Laboratory Press Release(2016)] Jet Propulsion Laboratory Press Release (2016). Seasonal cycles in Curiosity's first two Martian years. <http://www.jpl.nasa.gov/spaceimages/details.php?id=pia20600>.
- [Johnson et al.(2015)] Johnson, J., III, J. B., Hayes, A., Deen, R., Godber, A., Arvidson, R., Lemmon, M., Kuhn, S., Carsten, J., & Kennedy, M. (2015). Recent Mastcam and MAHLI visible/near-infrared spectrophotometric observations: Kimberley to Hidden Valley. In *Proceedings of the 46th Lunar and Planetary Science Conference*.
- [Lewis et al.(2016)] Lewis, K. W., Peters, S. F., & Gonter, K. A. (2016). First gravity traverse on the Martian surface from the Curiosity rover. In *Proceedings of the 47th Lunar and Planetary Science Conference*.
- [Mankiw(2014)] Mankiw, N. G. (2014). *Principles of Economics*. Cengage Learning, 7 edition.
- [McBride et al.(2015)] McBride, M. J., Miniti, M. E., Stack, K. M., Yingst, R. A., Edgett, K. S., Goetz, W., Herkenhoff, K. E., Heydari, E., Kah, L. C., Rowland, S. K., Schieber, J., Harker, D., Kennedy, M. R., Krezoski, G. M., Lipkamin, L., Nixon, B., & Beek, J. V. (2015). Mars Hand Lens Imager (MAHLI) observations at the Pahrump Hills Field Site, Gale Crater. In *Proceedings of the 46th Lunar and Planetary Science Conference*.
- [Milliken et al.(2016)] Milliken, R., Hurowitz, J., Bish, D., Grotzinger, J., & Wiens, R. (2016). The chemical and mineralogical stratigraphy of lower Mt. Sharp: Relating rover observations to orbital predictions. In *Proceedings of the 47th Lunar and Planetary Science Conference*.
- [Miniti et al.(2015)] Miniti, M., Beek, J. V., Garvin, J., Goetz, W., Grotzinger, J., Harker, D., Herkenhoff, K., Kah, L., Kennedy, M., Krezoski, G., Lipkamin, L., Rowland, S., Schieber, J., Stack, K., & Yingst, R. (2015). Mapping the Pahrump Hills outcrop using MARDI sidewalk mosaics. In *Proceedings of the 46th Lunar and Planetary Science Conference*.
- [Mishkin et al.(2006)] Mishkin, A. H., Limonadi, D., Laubach, S. L., & Bass, D. S. (2006). Working the Martian night shift: The MER surface operations process. *IEEE Robotics and Automation Magazine*, **13**(2), 46–53.
- [Morris et al.(2016)] Morris, R., Vaniman, D., Blake, D., Gellert, R., Chipera, S., Rampe, E., Ming, D., Morrison, S., Downs, R., Treiman, A., Yen, A., Achilles, C., Bristow, T., Crisp, J., Marais, D. D., Farmer, J., Fendrich, K., Frydenvang, J., Graff, T., Grotzinger, J., Morookian, J., & Schwenzer, S. (2016). High-temperature, perhaps silicic, volcanism on Mars evidenced by tridymite detection in high-SiO₂ sedimentary rock at Gale crater, Mars. In *Proceedings of the 47th Lunar and Planetary Science Conference*.
- [Newsom et al.(2016)] Newsom, H. E., Belgacem, I., Jackson, R., Ha, B., Vaci, Z., Wiens, R. C., Frydenvang, J., Gasda, P., Lanza, N., Clegg, S., Gasnault, O., Maurice, S., Cousin, A., Rapin, W., Forni, O., Banham, S., Gupta, S., Williams, A., Grotzinger, J., Blaney, D., Schroeder, J., Calef, F., Francis, R., Ehlmann, B., Yen, A., Stein, N., Watkins, J., Rubin, D., Bridges, N., Johnson, J., Lewis, K., Payr, V., Mangold, N., Edgett, K., Fey, D., Fisk, M., Gellert, R., Thompson, L., Schmidt, M., Perrett, G., Kah, L., Kronyak, R., Anderson, R., Herkenhoff, K., Bridges, J., Hurowitz, J., Schieber, J., & Heydari, E. (2016). The materials at an unconformity between the Murray and Stimson formations at Marias Pass, Gale crater, Mars. In *Proceedings of the 47th Lunar and Planetary Science Conference*.

- [Rampe et al.(2016)] Rampe, E. B., Ming, D. W., Morris, R. V., Blake, D. F., Bristow, T. F., Chipera, S. J., Vaniman, D. T., Yen, A. S., Grotzinger, J. P., Downs, R. T., Morrison, S. M., Peretyazhko, T., Achilles, C. N., Bish, D. L., Cavanagh, P. D., Craig, P. I., Crisp, J. A., Fairn, A. G., Marais, D. J. D., Farmer, J. D., Fendrich, K. V., Morookian, J. M., & Treiman, A. H. (2016). Diagenesis in the Murray formation, Gale crater, Mars. In *Proceedings of the 47th Lunar and Planetary Science Conference*.
- [Reid & Sanders(2012)] Reid, R. D. & Sanders, N. R. (2012). *Operations Management*. Wiley, 5 edition.
- [Stack et al.(2015)] Stack, K. M., Grotzinger, J. P., Gupta, S., Kah, L. C., Lewis, K. W., McBride, M. J., Minitti, M. E., Rubin, D. M., Schieber, J., Sumner, D. Y., Beek, L. M. T. J. V., Vasavada, A. R., & Yingst, R. A. (2015). Sedimentology and stratigraphy of the Pahrump Hills outcrop, lower Mount Sharp, Gale Crater, Mars. In *Proceedings of the 46th Lunar and Planetary Science Conference*.
- [Thomas(2016)] Thomas, G. (2016). *How to do Your Case Study*. SAGE Publications.
- [Thompson et al.(2015)] Thompson, L. M., Gellert, R., Spray, J. G., & Kah, L. C. (2015). The composition of the basal Murray Formation at Pahrump Hills, Gale Crater, Mars. In *Proceedings of the 46th Lunar and Planetary Science Conference*.
- [Vasavada et al.(2014)] Vasavada, A. R., Grotzinger, J. P., Arvidson, R. E., Calef, F. J., Crisp, J. A., Gupta, S., Hurowitz, J., Mangold, N., Maurice, S., Schmidt, M. E., Wiens, R. C., Williams, R. M. E., & Yingst, R. A. (2014). Overview of the Mars science laboratory mission: Bradbury landing to Yellowknife Bay and beyond. *Journal of Geophysical Research*, **119**(6), 1134–1161.
- [Watkins et al.(2016)] Watkins, J., Grotzinger, J., Stein, N., Banham, S., Gupta, S., Rubin, D., Stack, K., & Edgett, K. (2016). Paleotopography of erosional unconformity, base of Stimson formation, Gale crater, Mars. In *Proceedings of the 47th Lunar and Planetary Science Conference*.
- [Yin(2014)] Yin, R. K. (2014). *Case Study Research: Design and Methods*. SAGE Publications.
- [Yingst et al.(2015)] Yingst, R., Berger, J., Cohen, B., Hynek, B., & Schmidt, M. (2015). A test of two field methods: Determining best practices in reconnoitering sites for habitability potential using a semi-autonomous rover. In *Proceedings of the 46th Lunar and Planetary Science Conference*.