A Case Study of Productivity Challenges in 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 has been a challenge for Mars surface missions. While the rovers have made major discoveries and dramatically increased our understanding of Mars, they often require a great deal of effort from the operations teams and achieving mission objectives can take longer than anticipated. We conducted an in-depth case study of Mars Science Laboratory operations in order to identify the productivity challenges facing surface missions. In this paper, we describe how we performed the case study and analyzed the data. We present and discuss the significant productivity challenges we identified during the study. In addition to informing future surface exploration missions, the study is relevant for a wide range of applications in which operators must interact with a robotic system with limited communication opportunities.

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 19 samples and driven more than 15 kilometers (Grotzinger and others 2015; Vasavada et al. 2014). Curiosity is currently in its second extended mission and continues to make new discoveries as it explores Mount Sharp.

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. These productivity challenges are anticipated to increase as our aging fleet of sun-synchronous orbiters are replaced by non-sun-synchronous orbiters, which do not provide a consistent pattern of "end-of-day" downlink relays.

To better understand the productivity challenges faced by surface missions, we conducted a case study of MSL operations. The case study included structured interviews with science and engineering operations personnel and a detailed analysis of daily activities during three different science campaigns. For each campaign in the study, 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 the team allocated vehicle resources during the campaigns. The interviews with operations personnel were designed to complement the campaign analysis by providing a broader perspective on surface mission productivity factors.

In the following sections we describe the method we used to collect and analyze data from MSL science campaigns and provide a discussion of the significant productivity challenges that we identified during the study. These results are being used to guide the design and development of flight systems and ground practices for future surface exploration missions. In addition, these results are relevant to other applications in which operators interact with robotic systems with limited communication opportunities.

Background on MSL Mission Operations

We begin with a brief overview of some important facets of MSL operations to provide context for the case study. One of the challenges of surface missions is the degree to which operations are impacted 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.

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).

MSL operations augments this tactical process with "strategic" and "supratactical" phases (Chattopadhyay et al. 2014). Strategic planning focuses on developing long-term

---

Copyright © 2017, California Institute of Technology. U.S. Government sponsorship acknowledged.

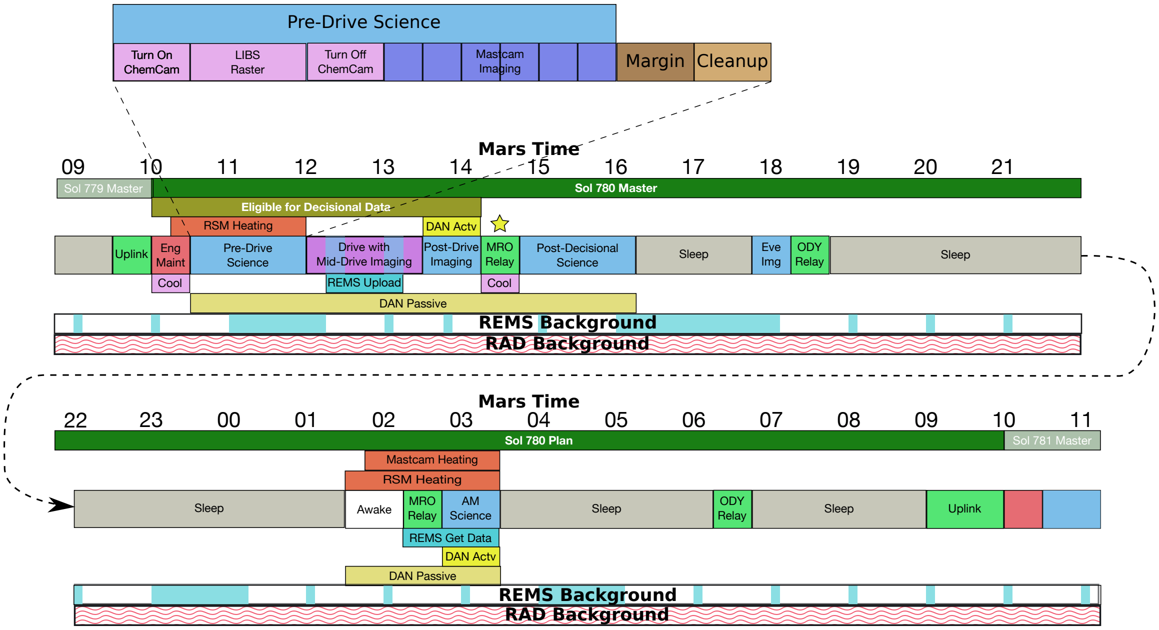

Figure 1: Example sol in the life of the rover.

plans, typically spanning weeks or months, to achieve high-level objectives. Examples of strategic planning include the development of strategies for exploring a large geographical area or a high-level traverse route for reaching a distant objective. 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.

An Example Sol in the Life of the Rover

To provide an idea of how the team operates the rover, Figure 1 illustrates a typical drive sol, derived from an actual Sol 780 command products. Following are some key aspects of the sol.

The plan for each sol begins with an "Uplink" window in which new commands products may be sent to the vehicle from Earth. There are various downlink windows throughout the sol in which the rover uses orbiter relays to send collected data back to Earth. While there are multiple downlink windows, certain downlinks have increased importance based on the time that data in the relay will reach operators. If data from a relay will reach operators by the start of the next tactical planning shift, then they relay is termed "decisional" because data from the relay can be used to make decisions in for the rover's next plan. In Figure 1 the starred "MRO Relay" represents the decisional relay for this sol. It is important to realize that for this plan, only the data collected prior to this pass could be used to inform the next plan.

Another important aspect of Figure 1 is how the team structures the sol into "blocks" of activity. For example, the main portion of the rover's day consists of a Pre-Drive Science block, a Drive with Mid-Drive Imaging bock and a Post-Drive Imaging block. The block structure organizes activity into related groups and allows a "Master" sequence to enforce timing between these major types of activity. The latter has to do with uncertainties in predicting the duration of activity in the plan. Due to environmental conditions such as lighting, scene content and terrain, the time to perform imaging and drive activities varies. The team uses the block structure to ensure that if activity in one block runs longer than expected, it can be cut off to avoid interfering with subsequent activity. To protect against loss of data, the team builds "Margin" into each block, to allow activities to run longer than predicted. To deal with cases where durations exceed allocated margin, the team also sequences "Cleanups" after each block, to ensure that any activity is finished before the start of the next block.

Restricted Sols

The vast majority of the surface mission is conducted with the team restricting operations to daytime hours on Earth. The consequence is that the operations team is often out of sync with the activity of the rover on Mars. Figure 2 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. 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 2. If the vehicle were allowed to make significant changes to its state, in particular driving to a new location, this would significantly limit

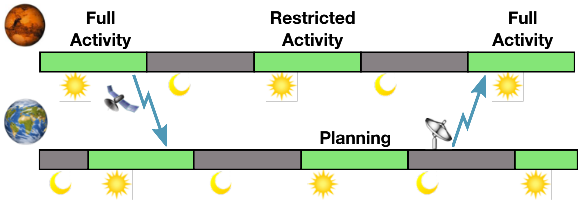

Figure 2: Mars activity vs. Earth planning.

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. Overall, 41% of sols on the MSL mission are restricted sols.

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).

MSL is a complex and ever-changing mission. At the time of our 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.

Selection of Cases

The primary purpose for conducting our case study was to inform our research project in which we will develop technologies and 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.

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 3 shows the campaigns that we have selected for the study.

The objective of the Pahrump Hills campaign was to study the basal layer of Mount Sharp (Figure 3 (a)) (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 on Earth. In a walkabout, the geologist makes an initial pass over the area performing a coarse survey which is used to identify areas for more detailed study on a subsequent pass.

After completing investigations at Pahrump Hills, the team drove the rover toward higher levels of Mount Sharp following a route referred to as Artist’s Drive (Figure 3 (b)). 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 objective of the Marias Pass campaign (Figure 3 (c)) was to explore a contact where the Murray formation (the type of rock from Pahrump Hills) came into contact with an overlying geological formation called the Stimson formation (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.

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.

For the purpose of our case study, we began with the general notion that mission productivity related to 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 expanded on this premise by developing a conceptual model of how objectives are accomplished in a surface mission. This model is shown in Figure 4. 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

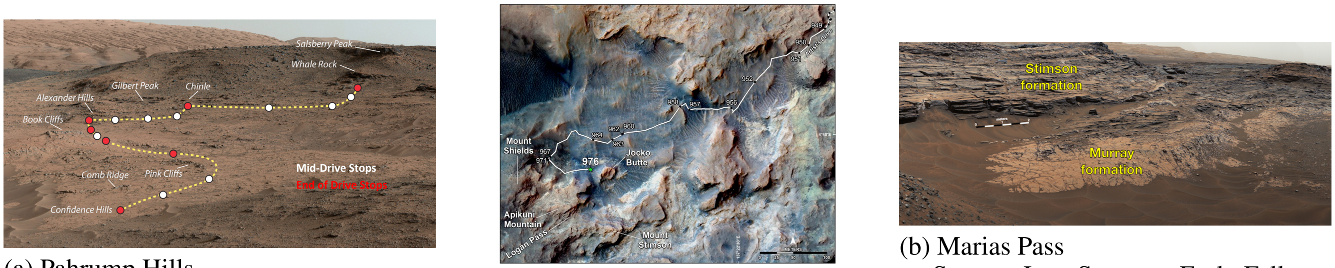

(a) Pahrump Hills
Season: Spring
Sols: 780 - 798 (19 sols)
Drive dist: 152m

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

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

Figure 3: Selected campaigns for case study (NASA/JPL-Caltech/MSSS/Univ. of Arizona).

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 descoping 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.

Data Collection Methodology

Using the conceptual model in Figure 4 as a guide, we developed the data collection schema shown in Figure 5 that includes intent, activities, constraints and data along with relations among these entities. We worked through each sol of the campaigns, sifting through the plans, acquired data and telemetry, and written reports from operations personnel to collect and organize data with respect to this schema. The data gathering process was a combination of manually reading through activity plans and operations reports along with scripts we developed to assist in the collection process. The scripts we developed included utilities to identify links between data products and the activities that used that data and utilities to collect data on predicted and actual vehicle resource allocations. The objective in gathering this data is to identify cases of low and high productivity during each campaign and to help identify the factors that contributed to each.

Caveats

It is important to note some limitations in the case study. First, 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 selected campaigns do not capture all aspects of the mission. In particular, we did not include contact science and 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 planned activity. Operations personnel were generally very good about documenting in 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. We concluded that for our objective of understanding productivity factors 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

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.

Case Study Results

Following is an overview of results from the case study. A comprehensive report of the case study is available in a technical report (Gaines et al. 2016).

Analysis of Campaign Activity by Sol

We begin our comparison of the three campaigns by summarizing the sol-by-sol breakdown of campaign activity in Table 1. Sols labeled "Campaign" were those that directly

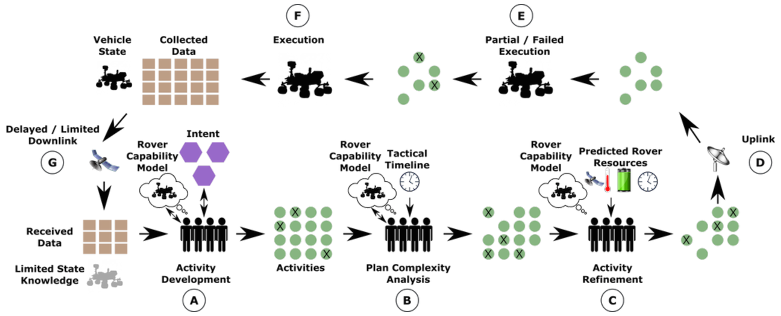

Figure 4: Factors impacting surface mission productivity.

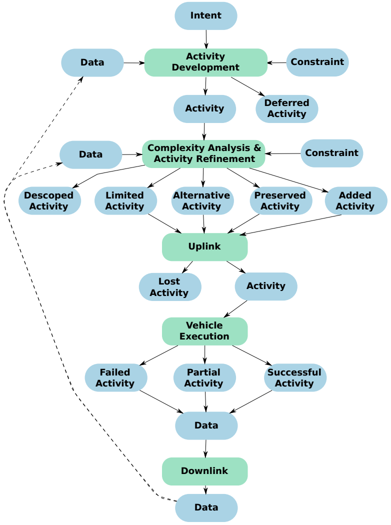

Figure 5: Data collection schema.

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 additional drives were required to reach objectives. This includes cases where a planned drive faulted out early and had to be re-planned on a subsequent sol. “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 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. 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.

The first takeaway from Table 1 is that there are a large percentage of sols in each campaign that are not making significant progress toward campaign objectives. This indicates there is significant potential for increasing mission productivity.

Our second observation is that although Pahrump Hills and Marias Pass shared a common objective, characterization of a geological region, it is Artist’s Drive and Marias Pass campaigns that have a nearly identical campaign activity profile. This indicates, that for the selected campaigns, the overall campaign objectives was not a significant factor in determining productivity. Instead, Table 1 indicates restricted sols and terrain conditions appear to be the dominant factors in the productivity for these campaigns. Although Artist’s Drive and Marias Pass had different campaign objectives, they shared a common pattern of restricted sols and they took place in similar, challenging, terrain conditions. Pahrump Hills had a total of 9 tactical shifts (days the tactical team worked) 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. This difference in number of restricted sols between Pahrump Hills and the other campaigns is largely luck of campaign timing. The Pahrump Hills campaign began just as the mission was transitioning into restricted sols, while both Artist’s Drive and Marias Pass occurred during a period in which most of the sols were nominal.

Table 1 also highlights the significance of terrain for these types of campaigns. Artist’s Drive and Marias Pass both occurred within more complex terrain than Pahrump Hills. The terrain made drive planning more challenging and resulted in more sols in which a drive faulted out and needed to be re-planned the next day.

Analysis of Resource Usage

Another way to evaluate campaign productivity is to look at how vehicle resources were used during the campaigns. We performed a detailed series of analyses on how the team allocated vehicle resources each sol of the campaigns. This included tracking of predicted and actual allocations of flight computer duration, energy and data volume. Figure 6 shows

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

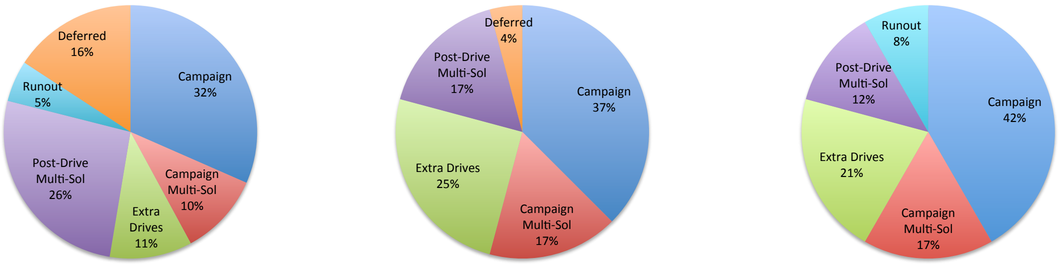

Pahrump Hills **Artist's Drive** **Marias Pass**

Table 1: Breakdown of sols for all campaigns.

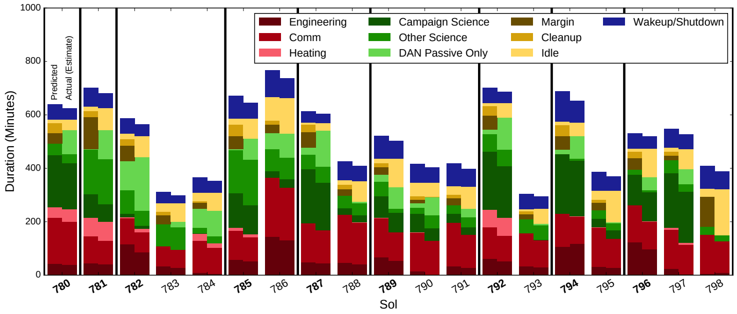

Figure 6: Allocation of flight computer duration for Pahrump Hills Walkabout campaign.

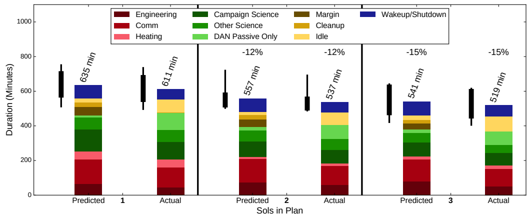

(a) Per-sol allocation of flight computer duration

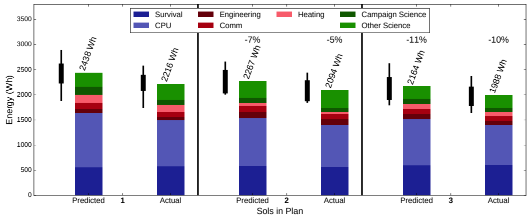

(a) Per-sol allocation of energy

Figure 7: Per-sol allocation of resource allocations for 1-sol, 2-sol and 3-sol plans.

an example using the flight computer duration allocation during the Pahrump Hills campaign. Multi-sol plans are indicated with vertical black lines.

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 in 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. 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.

To quantify this reduction in resource allocation, we plotted the per-sol allocation of flight computer duration and energy for 1-sol, 2-sol and 3-sol plans. Figure 7 shows the results. The results show that, on average, there is a reduction of 12% flight computer duration use and 7% reduction of energy use per sol for 2-sol plans. Usage is further reduced for 3-sol plans, with 15% and 11% reductions for flight computer 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.

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

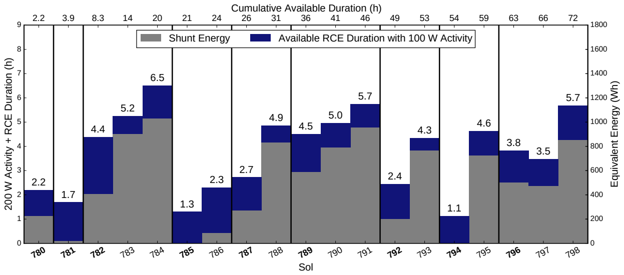

Figure 8: Estimate of extra duration availability for Pahrump Hills Walkabout campaign.

awake time to support an estimated additional 72 hours of campaign-related activity over the span of the 19 sols for the Pahrump Hills campaign, as shown in Figure 8. Similar analysis estimated an additional 62 hours and 69 hours of campaign activity could have been performed during the Artist's Drive and Marias Pass campaigns, respectively.

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.

On a more subjective note, discussions with operations personnel, including scientists and engineers, indicate that the team generally found the Pahrump Hills campaign to be more satisfying than the others. Pahrump Hills is often referred to as the way in which the team would like to conduct exploration campaigns. This is an especially interesting observation given that Table 1 shows that the number of productive campaign sols for Pahrump Hills was less than the other two campaigns in this study.

We believe there are two major factors that lead to the sense of satisfaction with the Pahrump Hills campaign. First, the team put together an excellent strategic plan to guide the exploration of Pahrump Hills. This helped the team to maintain focus during the campaign and make effective use of the relatively limited number of non-restricted sols during this campaign. It also gave the team a clear picture of the consistent progress being made toward completing objectives during the coarse of the campaign.

A second factor is the difference in exploration strategies employed at Pahrump Hills and Marias Pass. The walkabout approach at Pahrump Hills allows the team to build in the application of more cost-intensive resources with lower risk, because they eliminate the "discovery" and "uncertainty" in the first round by identifying the subset of best targets. A walkabout also provides a better match with the science team's internal processing of the data (both machine processing and mental). It gives the team a few weeks or months to digest what they are discovering before leaving.

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.

Drawing from our analysis of campaign activity and interviews with operations personnel, we developed a list of the significant productivity factors we observed in the study. Some of the more significant factors were:

Ground-in-the-loop for target selection and 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 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 the need to re-plan drives that faulted out or require observations that did not have intended results due to lighting conditions or targeting problems.

Predicting Available Vehicle Resources: Inaccuracies in resource modeling, including activity power and duration requirements, can result in unnecessarily restricting planned activity. Because activity durations tend to be conservative, activities did not usually require time allocated for margin and cleanup, resulting in significant idle time during plan execution.

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.

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.

The next biggest productivity factor was the ground-in-the-loop requirements for responding to problems in previous plans. This resulted in extra sols needed to drive to locations and the need to repeat or give up on observations that did not meet expectations.

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. For example, we are focusing on how we can enable the engineering team to effectively interact with the vehicle without knowing the exact state in which the vehicle will be when command products are received. We are also investigating methods to enable the vehicle to detect when activities are not going as planned and robustly respond without the need to wait for ground interaction. We are also leveraging 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. 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

[Chattopadhyay et al. 2014] Chattopadhyay, D.; Mishkin, A.; Allbaugh, A.; Cox, Z. N.; Lee, S. W.; Tan-Wang, G.; and Pyrzak, G. 2014. The Mars Science Laboratory supratactical process. In *Proceedings of the SpaceOps 2014 Conference*.

[Gaines et al. 2016] Gaines, D.; Doran, G.; Justice, H.; Rabideau, G.; Schaffer, S.; Verma, V.; Wagstaff, K.; Vasavada, A.; Huffman, W.; Anderson, R.; Mackey, R.; and Estlin, T. 2016. Productivity challenges for Mars rover operations: A case study of Mars Science Laboratory operations. Technical Report D-97908, Jet Propulsion Laboratory.

[Grotzinger and others 2015] Grotzinger, J. P., et al. 2015. Deposition, exhumation, and paleoclimate of an ancient lake deposit, Gale crater, Mars. *Science* 350(6257).

[Milliken et al. 2016] Milliken, R.; Hurowitz, J.; Bish, D.; Grotzinger, J.; and 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*.

[Mishkin et al. 2006] Mishkin, A. H.; Limonadi, D.; Laubach, S. L.; and Bass, D. S. 2006. Working the Martian night shift: The MER surface operations process. *IEEE Robotics and Automation Magazine* 13(2):46–53.

[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.; and 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*.

[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.; and 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.

[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.; and 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*.