Results from the First Year and a Half of Mars 2020 Robotic Operations

Vandi Verma
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
vandi@jpl.nasa.gov

Mark Maimone
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
mark.maimone@jpl.nasa.gov

Evan Graser
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
evan.graser@jpl.nasa.gov

Arturo Rankin
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
arturo.rankin@jpl.nasa.gov

Kyle Kaplan
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
kyle.w.kaplan@jpl.nasa.gov

Steven Myint
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
steven.myint@jpl.nasa.gov

Justin Huang
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
justin.huang@jpl.nasa.gov

Amanda Chung
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
amanda.chung@jpl.nasa.gov

Kevin Davis
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
kevin.davis@jpl.nasa.gov

Andrei Tumbar
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
andrei.tumbar@jpl.nasa.gov

Iona Tirona
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
iona.tirona@jpl.nasa.gov

Michael Lashore
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
michael.lashore@jpl.nasa.gov

Abstract—**Mars 2020 Robotic Operations is responsible for the development, planning and Mars execution of robotics aspects of the mission. This includes the Perseverance rover's mobility, manipulation, and sampling operations, and the Ingenuity helicopter's flights. As of October 2022 (Martian Solar Day 562, or sol 562), the rover has driven 13,179.5m and collected 15 samples, while the helicopter has logged 32 flights, covering 7281m, with an overall flight time of 3467 seconds. Perseverance and Ingenuity have accomplished several firsts such as coring and caching samples autonomously, and demonstrating powered flight on Mars. Perseverance has also set new planetary rover records such as the longest continuation drive distance (699.9m with no human review), longest single-sol autonomous drive distance (319m), and total autonomously evaluated drive distance (11,594m out of 13,172m total, i.e. 88% of all driving, an order of magnitude more than previous NASA Mars rover missions).**

This paper presents results from the first year and a half of Mars operations to highlight the operations approach that enabled this success, the challenges encountered, and lessons learned. Challenges include an unexpected reboot while driving that led to the discovery of a race condition in the rover flight software, pebbles unexpectedly interfering with key sampling hardware, the Martian winter temporarily grounding the helicopter, and difficult communication situations between the rover and helicopter that arise during periods of limited mobility. This paper also describes the sampling sol path, which is used by the operations team to execute sampling activities in a repeatable manner consistent with the goals of the science team and the capabilities of the sampling hardware. Abrading and sampling performance to date is evaluated, with a particular emphasis on changes to operations as a result of the soft rocks encountered in the Delta region on the western border of Jezero Crater.

978-1-6654-9032-0/23/$31.00

TABLE OF CONTENTS

1. INTRODUCTION

The Mars 2020 Perseverance rover and Ingenuity helicopter landed in Jezero crater on February 18, 2021. The location of primary interest was an ancient delta in Jezero Crater, which scientists believe is one of the best places on Mars to search for potential signs of ancient life. There were hazards for landing in proximity to the delta, which led the onboard Terrain Relative Navigation system to choose a landing spot several kilometers away from it. That decision led to the need to drive 5km back to the delta during a focused Rapid Traverse period, after exploring inside the crater.

The Mars 2020 mission has demonstrated multiple advances in autonomy and mission capabilities over the previous Mars

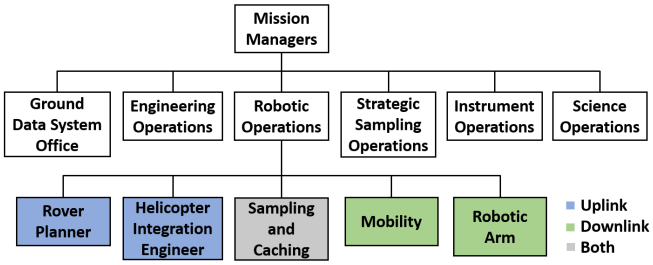

Figure 1. M2020 surface operations organization illustrating Robotic Operations roles

Exploration Rover (MER) and Mars Science Laboratory (MSL) missions. Perseverance drills a core sample from a nearby rock, transfers that sample into an internal Adaptive Cache Assembly, and seals it into a sample tube in a single command cycle; earlier missions required multiple Martian Solar Days (sols) to process samples. Accomplishing the task quickly was necessary to minimize sample contamination, and 15 samples have been collected as of October 2022 (sol 562). Autonomous driving has also been greatly enhanced, due to an updated Thinking While Driving software capability and additional compute power compared to earlier missions. Autonomous drives now move at nearly the same speed as purely pre-planned, human-directed drives. That capability, and the updated planning tools and processes that enable it, have led to unprecedented results. Perseverance now uses autonomous navigation for the majority of its driving; 88% overall, and 94% during the Rapid Traverse period alone. Only 7% of the total drive distance on prior NASA missions was with autonomous navigation. Ingenuity helicopter has demonstrated the first powered flight on another planetary body and has performed scouting flights in support of rover drives.

2. ROBOTIC OPERATIONS ROLE

As illustrated in Figure 1, the M2020 *tactical* surface operations teams that report to Mission Managers (MM) are Ground Data Systems Operations (GDSO), Engineering Operations (EO), Robotic Operations (RO), Instrument Operations (IO), and Science Operations (SO). RO is one of the teams that is new to the M2020 mission. Previous Mars rover missions organized the surface operations team into two general teams, uplink and downlink. The uplink team is responsible for activities related to creating new command sequences to the rover for future execution, and the downlink team is responsible for assessing telemetry downlinked to Earth that resulted from the execution of those commands.

On previous Mars missions like Mars Science Laboratory (MSL), engineers that perform robotic activity planning on the uplink team are assigned the Rover Planner (RP) role within the Integrated Planning and Execution (IPE) uplink team. Engineers that perform assessment of completed robotic activities are assigned to the Mobility/Mechanisms and Sample Acquisition & Sample Processing and Handling (SA-SPaH) downlink roles within the EO team. MSL RP team members earn separate certifications for each role; mobility, robotic arm, and sampling, as well as a related Strategic RP role (to plan a few sols beyond the current uplink horizon).

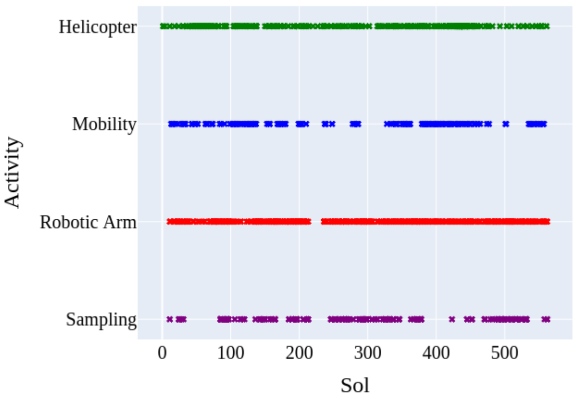

Figure 2. A plot showing interleaved robotic operations activities by sol from sols 0 through 562.

Responsibility for planning Sampling and Caching activities on M2020 was pulled out of the RP role and placed into a new role called Sampling and Caching (SNC) to reflect the increased decoupling of sampling from other activities. The M2020 RP role is responsible for planning Mobility and Robotic Arm (RA) activities and all team members are required to be certified for both. Both the RP and SNC teams also support daily planning of future sols, in Campaign Implementation (i.e., strategic) roles.

When the M2020 organizational structure was created, there was a recognition that the robotic roles work closely together and there was benefit in combining them into a single team. Figure 2 shows interleaved robotic operations activities from sols 0 through 562. Having a single team allows for more agile development of robotics-specific ground software, since its operational impact is limited to the RO team. The roles that were included in the M2020 RO team are RP, Mobility downlink, RA downlink, the new SNC role, and another role new to M2020 called Helicopter Integration Engineer (HIE), which is the rover tactical team that interfaces with the Ingenuity Helicopter Operations (HO) team. Nearly every sol includes hundreds or thousands of commands created by the RO teams, as shown in Figure 3.

The SSO team generally operates on a strategic timeline to develop and release sampling products for SNC team to use during tactical shifts. As the surface operations team prepares to execute a regolith sample collection and the creation of a storage depot of 11 sample tubes on the Martian surface in late 2022, with most of SSO's development effort completed, there is currently an effort to transition the remaining responsibility for sampling strategic development to the SNC team.

Currently, the RO team has approximately 50 members[2], 20% of which are certified in more than one RO role. The RO downlink roles are direct consumers of the RO uplink reports and the RO uplink roles are direct consumers of the RO downlink reports and products. The HIE role staffed both downlink and uplink shifts until approximately sol 80, when its downlink responsibilities were transferred to the HO role and it became an uplink only role. Currently, SNC is the only RO role that staffs both downlink and uplink, with all team members certified to perform both downlink and uplink responsibilities. As the SNC role has broad responsibilities encompassing downlink, uplink, and strategic work, none of

---

[2] A majority of the members work part time in other roles such as flight software development, ground software development, and verification & validation.

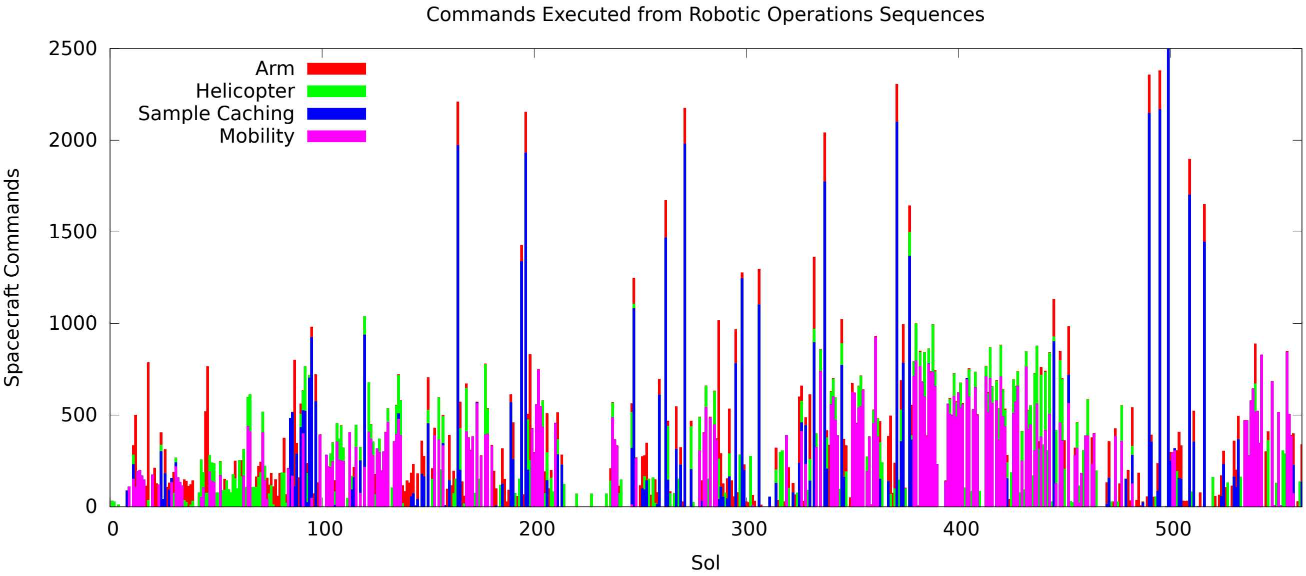

Figure 3. Number of commands executed on the rover each sol by the four RO subsystems.

the current SNC team members are certified for other RO roles.

3. ROVER MOBILITY

The Perseverance mobility system and its initial checkouts already broke records set by other missions during the first 210 sols of operations [1], and that trend has continued into the next year. That initial period covered just the terrain shown in the lower right of the trajectory plotted in Figure 5; the initial landing site, and the southeastern corner of the hourglass-shaped Seitah region. That period of exploration ended with the rover just inside the southern edge of Seitah, near the Crater Floor Campaign location shown in the figure.

The Next Year of Driving

The science team chose to remain inside Seitah for the next four months for the *Crater Floor Campaign*, exploring sand ripples and sampling the new types of terrain there. The mobility team also took advantage of the novel terrain to exercise the Autonomous Mapping capability, to see how well the onboard system could generate 3D maps of the sand ripples using Autonav's terrain understanding software, prior to releasing the full Autonav navigation capability. The crater floor campaign was complete by sol 340, at which point Perseverance began to retrace its route back to a point just east of the landing site. Although we had spent 340 sols driving around and into Seitah, it only required 15 of the next 22 sols to drive 2 km all the way back. Having driven on this terrain before allowed us to confirm our strategic assessment of the suitability of the terrain for Autonav and the full capability was released. This *Return to the Landing Site* drive was our first confirmation that Autonav could be used routinely to cover hundreds of meters in a single sol, as we drove 2,013.97 m in 15 drives from sol 340 through sol 362, averaging 134 m/sol. This period also resulted in our current single-sol longest drive distance record, when we achieved 319.79 m on sol 351, the second sol of a 3-sol drive plan, with all of that distance having been driven autonomously using Autonav.

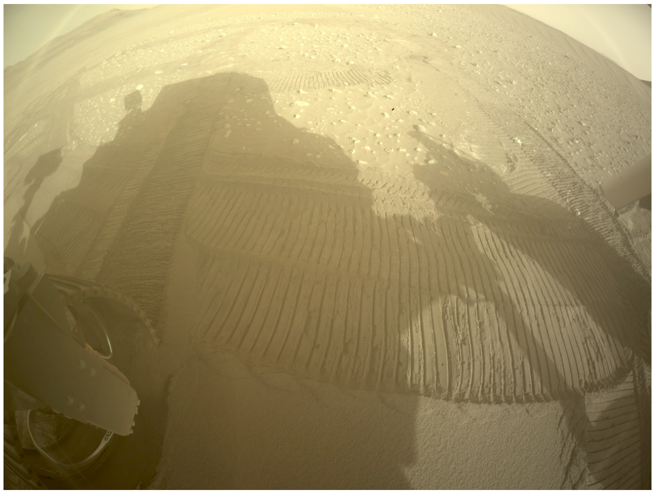

Figure 4. Left Rear Hazcam image from sol 548, after the rover experienced 65.5% slip. The horizontal tracks were created four months earlier, during the sol 428 drive. Note the difference between the sol 428 tracks (smooth, well defined grouser marks) and the sol 548 tracks (unevenly spaced, churning up the regolith).

The next high level driving goal of the mission was the to reach the Delta area to the west of Seitah, on the far side from the Octavia E. Butler landing site as fast as possible. Now that the viability of using Autonav for long-distance driving had been clearly established, we had several choices: drive north along the edges of Seitah for nearly 5 km of mostly benign terrain, attempt to drive 2.4 km due west while crossing the ripples and high risk terrain within Seitah, or drive more than 3 km through moderate risk terrain to the south. There was higher risk of crossing through Seitah and the southern route since at this point in the mission we didn't have statistics on Autonav performance in the dense sandy areas. Drive rate through sand without Autonav was expected to be four times slower than using Autonav along the longer northern route. The Due to the higher confience in the northern route, it was selected.

The period of driving around the northern edge of Seitah to reach the delta became known as the *Rapid Traverse Campaign* [2]. It spanned 5 km of driving, and lasted from sol 379 to sol 409 when we reached the *Three Forks* area just south of the Delta. Several more driving records were set during this part of the mission. It reached its 5 km goal in just 31 sols, 3X faster than any prior mission. Perseverance covered 5,063.4 m using 24 sols for actual driving, resulting in an average progress rate of 210 m/sol; this is apparent in Figure 5 by the spacing of the white dots, which indicate the end of an entire sol's drive and are far more widely spaced along this route. Perseverance covered 528.7 m during a single two-sol drive plan from sol 404 to sol 405, breaking Opportunity's record of 390 m from sol 383 to sol 385 in 2005. And it set a record of driving 699.9 meters over three sols (sol 407 to sol 409) with no human confirmation of each subsequent sol's starting position, using the Opportunistic Extension style of drive plans [2].

The next 100 sols were spent exploring the base of the Delta at Devil's Tanyard and areas to the northwest called Hogwollow Flats. We encountered challenges before successfully acquiring samples in this area as described later in this paper. Based on their analysis of terrain near the Delta, the science team chose to collect another sample back at the westernmost tip of the Enchanted Lake area. It spoke highly of the mobility system and operations team that the science team was confident that the 1 km traverse to that area would go quickly, and the return trip began in earnest on sol 535.

At first that drive went well, the first 770 meters were covered in 8 sols of driving through sol 543, using both manual and Autonav driving. On Sol 548, Perseverance attempted a long 200 meter drive. However, it only drove 20 meters before it was blocked from proceeding due to excessive slip, as measured by the Visual Odometry FSW (Flight Software). RPs had established an upper bound of at most 60% slip being acceptable for the drive, but 65.5% slip was measured while trying to cross westward over a north/south sandy ripple. This was surprising, since the rover had successfully crossed this very same spot on sol 428 four months earlier, and had encountered much less slip (see Figure 4). But one wheel was already at the crest of the hill, so on sol 551 another attempt was made, this time allowing 90% maximum slip. However, the rover FSW supports constraints not only on each individual slip measurement (called a "fast slip" constraint), but also on the average of slips reported over some distance (the "slow slip" constraint). The next attempt on sol 551 increased high individual slip, but also encountered a fault, this time due to high average "slow slip" of 78% over a 1.5m distance. The RP team as a whole had significant experience driving in slip-inducing terrains from earlier rovers and utilized it to safely navigate the vehicle through this first substantial slip event for Perseverance. Since the 551 drive also made forward progress, moving the other front wheel next to the crest of the ripple, the decision was made to continue forward rather than attempting to back away. The third attempt on sol 555 changed all the VO slip parameters based on the data to new high-slip terrain values, and succeeded in crossing the ripple after reaching a peak slip of 78%. The sol 555 drive ultimately ended up not only crossing the ripple, but also driving 237 m total and nearly reaching the Enchanted Lake sampling area, see Figures 28 and 32 for illustrations of the planned drive.

As of sol 562, the rover has reached the *Enchanted Lake site*, and the team has begun preparations to sample it.

Mobility Statistics—M2020 has driven 13,173 m in total as of sol 562. The Mobility system uses the rover's mast-mounted Navcam Engineering cameras for terrain assessment and Visual Odometry (VO) position estimation and slip compensation to assist with automatic navigation across the surface of Mars. Table 1 and Figure 6 give the distances driven using M2020's primary modes of driving. AVOID_ALL or "Autonav" driving, where the vehicle performs terrain assessment onboard and autonomously modifies its route to avoid any obstacles, has been active for 75.6% of the distance driven. 5,063.4 meters were driven during the 31 sol Rapid Traverse campaign conducted from sol 379 through sol 409, an amount that more than doubled the distance driven prior to the campaign. A remarkable 94.8% of the Rapid Traverse drive distance was achieved using autonomous navigation [2].

Mobility Faults

M2020 Mobility

Flight Software monitors and responds to a variety of faults which are capable of stopping or pausing a drive. Dozens of potential faults are checked at 8 Hz while in motion, and several others are checked approximately once per meter traveled. Several faults can be considered nominal under the nominal planning conditions and may be autonomously cleared by the command sequence, allowing a drive to continue. For example, RPs routinely set a maximum cutoff time for each drive, to ensure subsequent science activities and communications links will occur as scheduled. And to maximize the distance traveled, the nominal means of ending every Autonav drive is to reach that time. When that time is reached, the behavior is raise a CUTOFF_TIME fault to skip over the rest of the drive, then clear that fault at the end of the plan. Figure 7 shows all Mobility software faults that have been declared in the mission and specifies which were considered expected. 148 of 169 individual drives have completed successfully during the mission for a success rate of 88%. 15 drives have ended with an unexpected faults, a rate of 9%. The remaining 4% of drives were precluded by issues unrelated to Mobility and were not attempted.

Nominal Drive Faults—Whenever any drive fault occurs, the FSW will raise an error that will prevent future mobility commands from executing. These are some of the nominal ways of managing drive uncertainty using faults:

- CUTOFF_TIME: A cutoff time is set for all drives. If mobility is active when this time is reached, CUTOFF_TIME will halt the drive while leaving sufficient time to clear the fault and perform necessary cleanup activities.
- SAPP_MARGIN: Part of rover position uncertainty management. When the Surface Attitude Position and Pointing (SAPP) software indicates the rover is in need of a sun-based pose update, it will raise this fault and set an error to stop a drive. These can be cleared mid-drive following an explicitly-sequenced attitude update, allowing a drive to resume.
- NO_PATH: Thrown when auto-navigation cannot identify a safe drive path to continue making progress towards its goal. A common, normal, and helpful way for an AVOID_ALL drive to end.
- UNSAFE: Thrown when GUARDED driving determines the drive path would cross unsafe or unknown terrain. Nominal under most circumstances and an acceptable way for a GUARDED drive to end.

Unexpected Drive Faults—These are some of the faults used to safely stop a drive when the vehicle behavior is not as expected when planning:

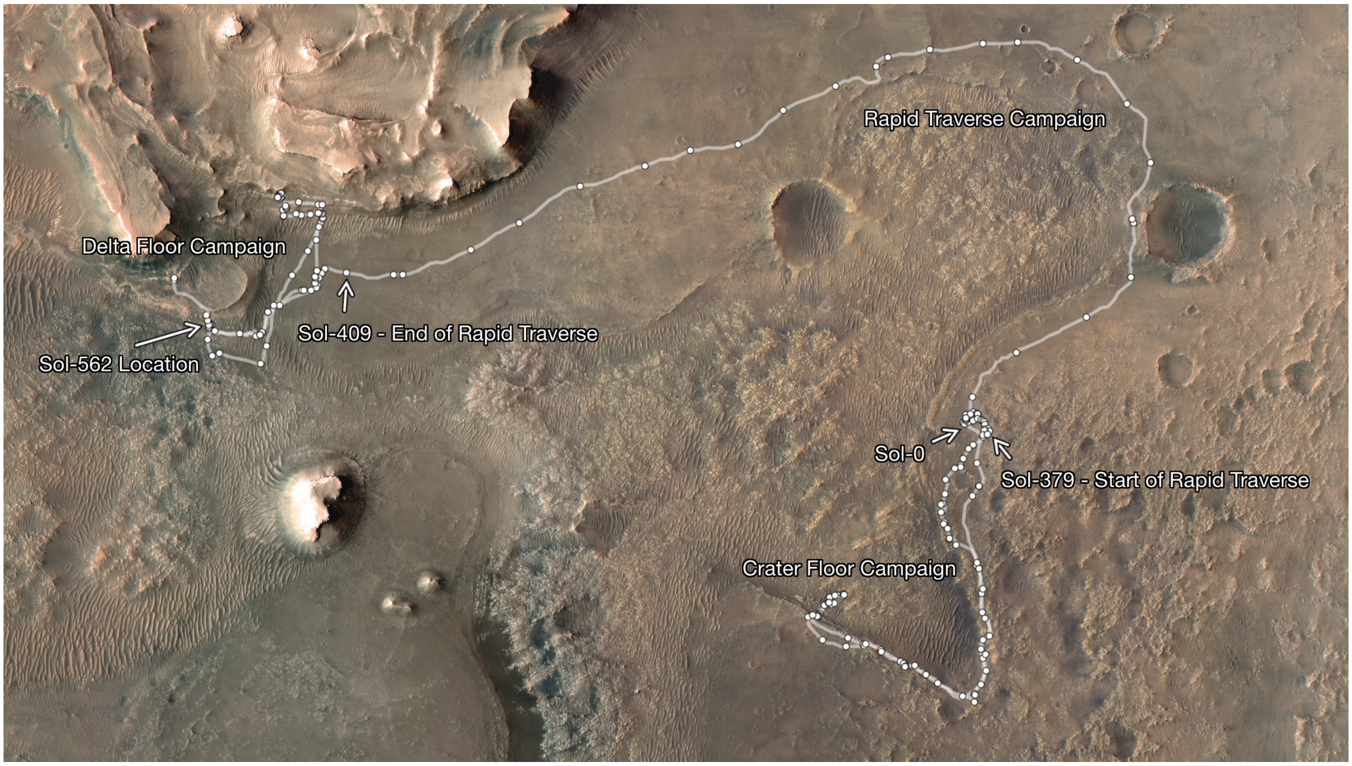

Figure 5. Orbital view of Perseverance's traverse through 562 sols. Most of the image shows the interior of Jezero crater, including the hourglass-shaped region in the middle containing many sand ripples known as Seitah. However the upper left corner in the northwest shows the edges of the Delta outflow, a raised area on the edge of the crater. The lightly transparent white line shows the path of the rover as it drove across the terrain, and the individual solid white dots indicate the final location at the end of a sol. The middle 5 km portion of the drive has much wider spacing than most of the rest of the drive, due to the longer distances covered each sol during the Rapid Traverse campaign.

Table 1. Drive Rate Summary for the first 562 sols, calculated from durations of all mobility commands, including long-range GO_TO and short range ARC and TURN commands.

Drive Mode	Distance [m]	Duration [hr]	Effective Rate [m/hr]	% of Total Distance
AVOID_ALL (AutoNav)	9963.0	101.7	98.0	75.6
GUARDED	126.0	1.06	118.9	1.0
UNGUARDED_Mapping	0.0	0.0	n/a	0.0
UNGUARDED_VO_Mapping	1504.5	14.5	104.1	11.4
UNGUARDED_VO	1269.3	11.7	108.1	9.6
UNGUARDED_Directed	309.8	2.74	113.2	2.4

- Reactive Safety Checks: Faults that occur when rover hardware exceeds a mechanical limit set for the drive. While there are dozens of limits, the following types are adjusted for each drive based on terrain and drive simulation results.
  - SUSP: Rocker-Bogie suspension angle exceeds limit
  - TILT: Rover tilt angle exceeds limit.
  - YAW: Rover yaw is outside its allowed range.
- UNRECOVERABLE: The mobility FSW has received inconsistent commands or results and does not know how to get things back to a consistent state. This was triggered several times in flight by a known FSW bug that always had a low percentage likelihood of tripping during a drive. That bug was corrected in the S7.1 FSW update in Fall 2021.
- VOSLIP_EXCESSIVE: Rover slip detected by Visual Odometry exceeds limits.
- VO_FAILURES: The number of failed Visual Odometry updates exceeds the limit. One instance occurred due to poor stereo performance in sandy terrain with a small number of detectable terrain features.
- MOT: A fault from the Motor control module that occurs while driving. The one instance seen to date was a motion timeout during a high-slip turn on sand.
- FATAL: A FSW assertion while driving. One instance in flight caused by a race condition between the SAPP and DMS (Data Management System) FSW modules.

VOSLIP_EXCESSIVE is the second most common fault to end a drive with 3 instances, all where the rover was traversing sand and experienced slip in excess of its allowed limits. Figure 8 and Figure 9 show slip statistics for all drives, with the 3 high slip faulted drives in excess of nominal limits.

In summary, the drive rates shown in Table 1 show a remarkable improvement over past missions. Every drive mode,

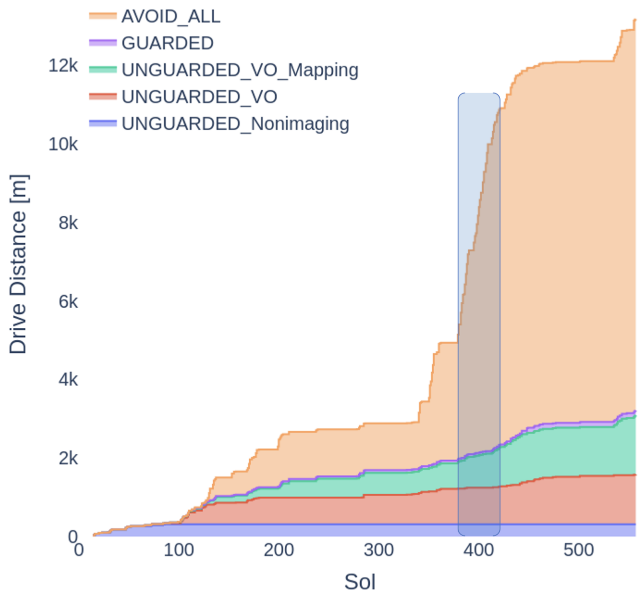

Figure 6. Cumulative odometry by drive mode. Rapid Traverse Campaign highlighted

Table 2. Frequency of Unexpected Drive Faults.

Fault Type	Count	% of Total
UNRECOVERABLE	4	26.7
VOSLIP_EXCESSIVE	3	20.0
SUSP	2	13.3
TILT	1	6.67
YAW	1	6.67
NO_PATH	1	6.67
FATAL	1	6.67
VO_FAILURES	1	6.67
MOT	1	6.67

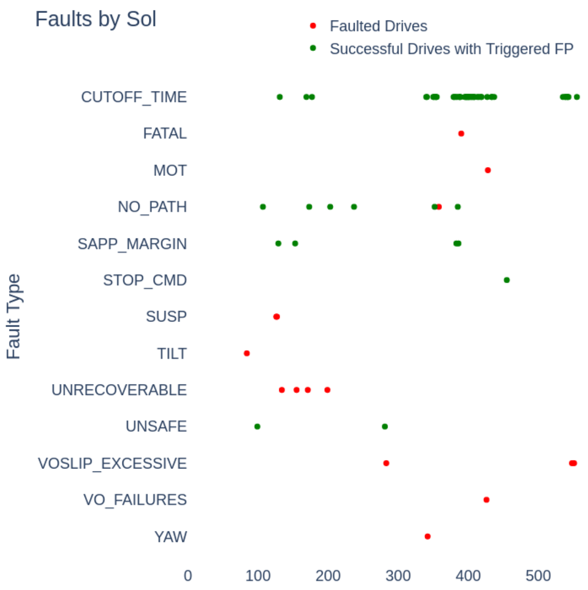

Figure 7. Mobility Faults through 562 sols. Green indicates a fault that was either expected or considered an acceptable way to end a drive. These are often autonomously cleared.

including those using full onboard autonomy, supports drive rates over 100 m/hr. Driving with Visual Odometry is at least 3 times faster than Curiosity (soon expected to be only 2 times faster [3]), and driving with Autonav is 4–10 times faster depending on terrain. And while MER and MSL missions only chose to employ Autonav for less than 7% of all driving, Perseverance has used Autonav and related Guarded and Mapping modes for 88% of all driving as of sol 562 [4].

4. ROBOTIC ARM

The Perseverance Robotic Arm (RA) extends two meters from the front of the rover chassis. The Turret is attached to the end of the RA and includes tools such as the coring drill to abrade the surface and collect sample cores; Gas Dust removal Tool (gDRT) to clear away dust; and the Facility Contact Sensor (FCS) to measure surface locations [1]. The turret also includes science instruments to collect close proximity measurements. The Planetary Instrument for X-ray Lithochemistry (PIXL) is an X-ray fluorescence spectrometer mounted on an articulated hexapod structure, and The Scanning Habitable Environments with Raman and Luminescence for Organics and Chemicals (SHERLOC) instrument performs fluorescence and Raman spectroscopy with an ultraviolet laser and a context imager. The Wide Angle Topographic Sensor for Operations and eNgineering (WATSON) contains a high-resolution camera. The robotic arm can conduct surface proximity observations with Turret-mounted science instruments, collect samples cores, dock with the rover body to transfer them to the sample caching system inside the rover body, self inspect, and perform additional science and engineering activities.

The robotic arm was stowed in front of the rover for landing. The first robotic arm motion occurred on Sol 12, when it was unstowed for the first time on Mars as part of the Surface Operations Changeover (SOX) commissioning phase. The robotic arm is stowed for driving. At each new location the robotic arm is unstowed for engineering and science activities such as proximity science and sample collection. It typically remains unstowed for the duration of placing instruments, conducting proximity observations and sampling operations at each location. There have been 46 robotic arm stow/unstow cycles in the first 562 sols as shown in figure 10.

Basic checkouts were followed by high level checkouts during SOX. Following SOX, First Time Activities (FTAs) were performed in parallel with the science mission, increasing the capability available for science [1]. From Sol 159 through 162 the first abrasion and abraded science campaign was successfully executed, acquiring WATSON, SHERLOC, and PIXL observations targeted within an abrasion patch. Shortly thereafter, the first sampling activity took place on Sol 164. In the first 62 sols the Facility Contact Sensor (FCS) contacted the surface 26 times and the drill was placed on the surface 32 times (Figures 11 and 12). The Corer was docked to the Adaptive Cache Assembly 45 times (Figure 10). The robotic arm placed WATSON 469 times, SHERLOC 23 times, and PIXL 34 times for proximity science observations (Figure 13).

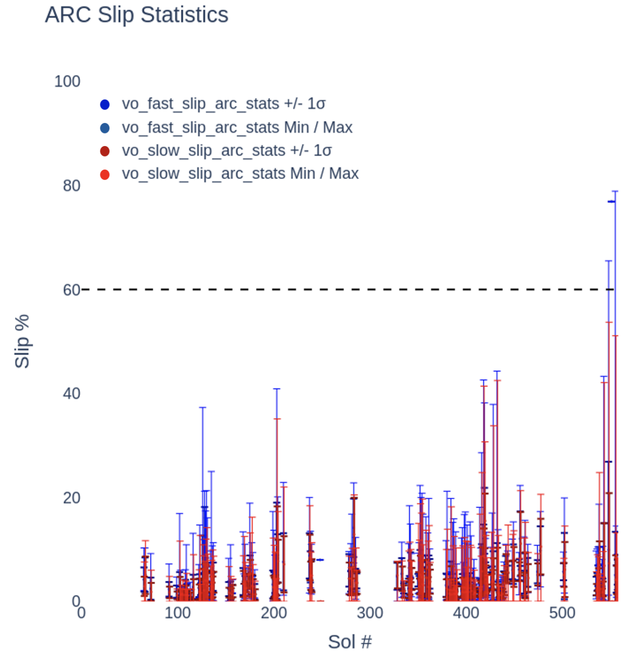

Figure 8. Rover Slip Percentages while Arcing. 60% is the default limit used for most drives.

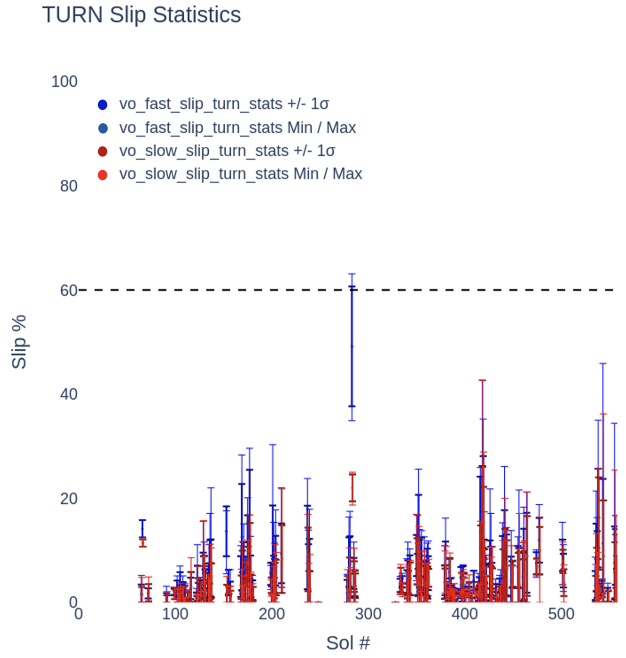

Figure 9. Rover Slip Percentages while Turning-In-Place. 60% is the default limit used for most drives.

5. SAMPLING AND CACHING

The Sampling and Caching Subsystem (SCS) has successfully acquired rock core samples, prepared abraded rock surfaces, and supported proximity science instrument placements. The SCS consists of a coring drill mounted on the end of the RA's turret, referred to as the Corer, the gas Dust Removal Tool (gDRT) also mounted to the turret, and the Adaptive Caching Assembly (ACA) located inside the rover body as seen in Figure 14. The Corer's two stabilizers are preloaded into rocks during abrading and coring, while the Corer feed translates a rotary-percussive drill into contact with the rock. The interface between the Corer and ACA is the Bit Carousel (BC), where 6 coring bits, 2 abrading bits, 1 regolith bit, 1 Witness Tube holder, and 1 Dust Mitigation Tool are stored. The RA docks with the BC to enable Bit Exchange (BX). Together, the ACA and BC store 43 tubes: 38 Sample Tubes and 5 Witness Tubes. Tubes are manipulated and stored within the ACA by the Sample Handling Arm (SHA) and its End Effector (EE) and are transferred to the BC at its lower door. Tubes are installed into a coring bit by the SHA, rotated from the lower to upper BC door, and then transferred to the corer for sample acquisition. As of sol 562, the SCS has collected 12 rock core samples, 1 atmospheric sample, and 2 Witness tube samples which document the SCS's exposure to the Martian elements over time [5].

The SCS is operated by the SNC team, who since landing have assessed and selected potential abrading and coring targets, simulated and validated sampling sequences prior to uplink, and reviewed downlink data from the rover to ensure the health and nominal performance of the sampling system [1]. SNC works closely with the SSO team who test, validate, and deliver the sampling products used by the SNC team to efficiently operate the SCS. Through the first year and half of Mars operations, these teams have worked diligently to ensure the successful operation of the sampling system despite the complexities of Mars surface operations and the occurrence of a few spacecraft anomalies.

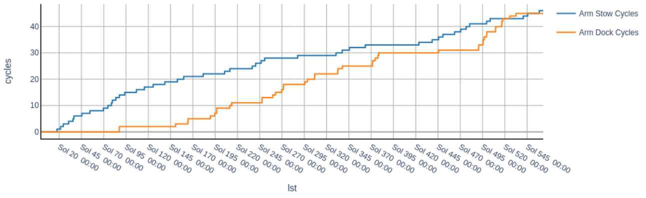

Figure 10. Number of Stow and Dock cycles from sols 0 through 562.

This section outlines the sampling sol path followed by the Mars 2020 team, with particular focus on the inputs to the process provided by the SNC and SSO teams. Abrading and Sampling performance to date is evaluated in the context of the regions of Jezero crater thus far explored by Perseverance. Finally, the response to an anomaly with the Bit Carousel is described, including the lessons learned by the team and resulting changes to operations.

Sampling Sol Path

To provide consistency in the sampling process for both scientists and engineers, the Mars 2020 team follows a "Sampling Sol Path". This sol path defines the order of activities required to nominally collect a single or paired sample from any given rock or patch of rocks. Each acquired sample plan has followed this process.

The sampling process begins with a collaboration between Mars 2020's Science and Sampling teams. Science identifies nearby targets with high value from their perspective and provides these candidate targets to the Sampling team. The sampling team then performs quantitative and qualitative assessment of the candidate targets. At this early point in the target assessment process, SNC is primarily focused on addressing the following questions:

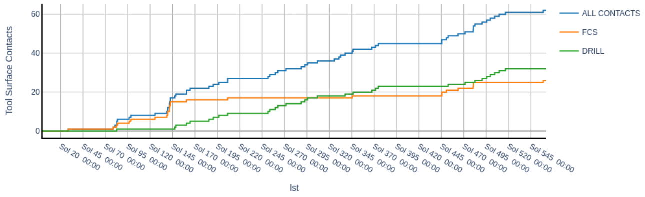

Figure 11. Facility Contact Sensor touches and drill placements through 562 sols.

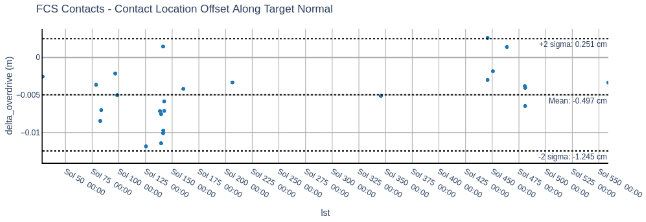

Figure 12. Modelled vs. actual surface target location as found with Facility Contact Sensor.

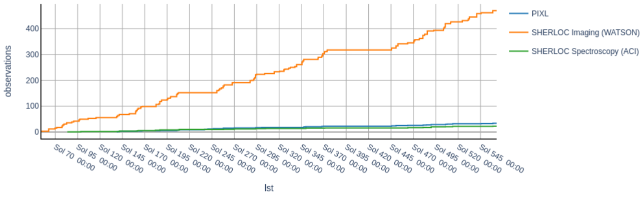

Figure 13. Number of observations taken with each instrument on the robotic arm.

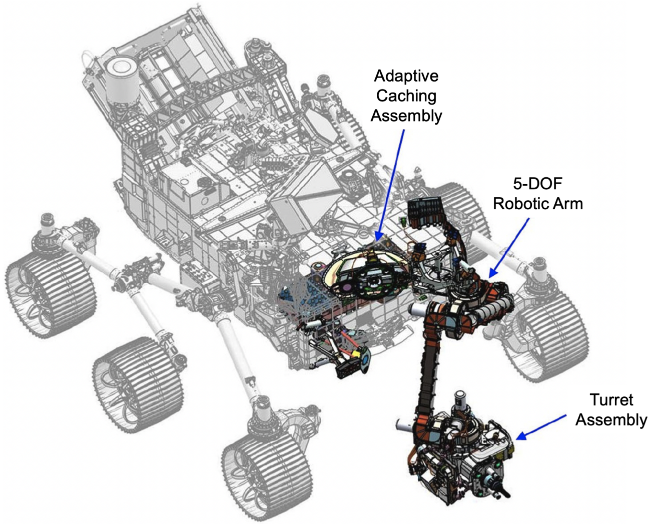

Figure 14. Mars 2020 Rover with SCS components highlighted.

- Is the rock of sufficient size to allow stabilizer placement?
- Is the rock stable and unlikely to shift or fracture when drilled?
- Does the rock have topographical features that would pose a challenge to abrading or coring?

The Rover Planner (RP) team works with the SNC team to ensure the rover can be parked at the desired drilling location and the various Corer and instrument placements are kinematically feasible for the RA. If a rock appears reachable and drillable from a distance, the operations team may begin the sampling sol path with a precision approach or “bump” drive to ensure the selected rock will be in the RA workspace.

Under nominal circumstances, this Sampling Sol Path takes 12 sols to complete, resulting in the creation of 1 new abrasion and the sealing of 2 sample tubes each containing a unique core sample as shown in Figure 15. After bumping to the target, the rover collects detailed workspace imagery in front of the rover, which generates 3-dimensional meshes and other ground-derived products required for surface interactions. A detailed target assessment begins, building on the initial assessment performed at a distance.

Target Assessment—In the sol following the precision approach, RPs and SNCs coordinate to select a set of up to four targets in the workspace to be designated candidate abrading or coring targets. SNCs evaluate each target against a set of Corer Placement Guidelines, which were initially derived from experience placing the Corer on a variety of abrading and coring targets during pre-launch development, and later refined based on lessons learned from flight operations.

Targets are carefully assessed against the guidelines, but evaluating every point in a given mesh for abrading or coring suitability would be an arduous and time-consuming process. To kick start this process and enable the broader science team’s assistance with target selection, goodness maps can narrow down the regions of interest in a particular workspace based on their viability for coring or abrading. Goodness maps are color coded on a red, orange, yellow, green scale and can be applied as an overlay to the workspace imaging taken by the rover’s engineering cameras. Each map is used to visually represent a different aspect of target assessment.

For example, the surface topography goodness map color codes regions of the workspace according to how well the area around a given point adheres to concavity and convexity limits (Figure 16). The area considered in this case is as wide as the distance between the stabilizers. The team uses these maps to quickly narrow down viable abrading or coring targets before moving on to more detailed assessments.

Full target assessment uses RSVP HyperDrive to evaluate target points and surface normals within the workspace mesh. Each target is assessed against a variety of quantitative criteria, including: overall rock size, surface concavity and convexity, surface roughness under the bit and stabilizers, stabilizer distance to the rock’s edge, and target angle relative to gravity. If a target adheres to the limits prescribed for abrading or coring, then it is considered a viable target. In the event that one or more metrics are violated, SNCs may still consider the target viable if the violation in question is not hardware safety critical, the performance risk is well-understood, and the risk is communicated to the science team.

Once a set of up to four targets has been identified by the tactical uplink team, the WATSON camera will take low and high standoff images of each candidate drilling target (at 7 cm and 40 cm distance). The WATSON images serve several purposes. The low standoff WATSON images can be used to reduce the placement uncertainty when abrading or coring by selecting a new target out of the WATSON image. This uses the focus distance from a low standoff WATSON image to perform a relative arm placement, which has a smaller

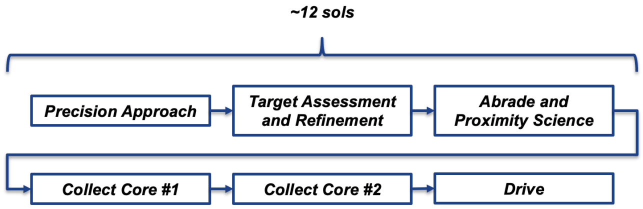

Figure 15. High-level overview of the paired sampling sol path.

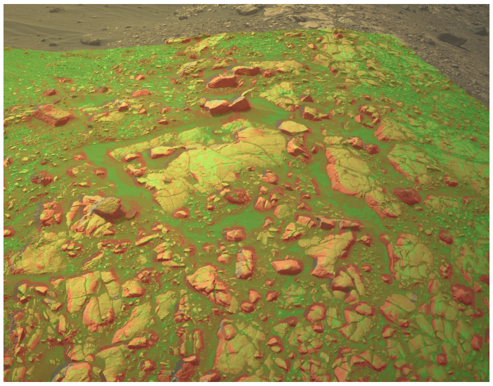

Figure 16. An example of the surface topography goodness map overlaid on a Front Hazcam image at Wildcat Ridge on sol 502.

uncertainty than selecting a target directly out of a Navcam or Hazcam workspace image. This increases the likelihood of success when preloading the stabilizers into the rock prior to coring or abrading and reduces the risk of turret hardware inadvertently coming into contact with the surface. In addition, both the low and high standoff WATSONs are used by SNCs to evaluate the bit and stabilizer placement locations for pebbles, regolith, and cracks which could impede a successful Corer placement. These environmental assessments are more qualitative, and are based on the team's collective experience placing the Corer on a variety of rocks in flight.

Abrading—The Corer is then used to drill an abrasion patch, which removes any surface topography, as well as the weathered outer layer of rock. This allows the Science team to study the underlying rock using a suite of scientific instruments like PIXL and SHERLOC. The team selects a single target out of one of the WATSON images to be the abrading target. The sampling team's primary input to this process is the depth of the abrasion needed to generate a flat surface, which is initially decided by applying margin to the surface roughness under the bit. The RPs then evaluate the placement viability of the instruments on an abrasion of that depth, based on the impact of the placement uncertainty margin on the necessary instrument standoff requirements and turret clearance outside the abrasion area, and provide a second abrasion depth recommendation. If there is a discrepancy, the science team selects the desired abrasion depth, balancing the risk of not creating a flat abraded patch against the need to place the instruments on the patch for analysis.

Sampling—Because an abrasion may disturb the workspace, additional workspace imaging must be acquired and target assessment analysis repeated to ensure the coring targets remain valid prior to sampling. Additionally, WATSON images may be re-taken to further reduce the placement uncertainty. Eventually, two coring targets are identified, and the team proceeds with collecting core samples, sealing them and returning the filled tubes to ACA storage.

Once the second core has been sealed, the team completes any remaining engineering activities and drives away to the next candidate sampling location.

Abrading and Sampling Performance

In addition to driving the target assessment and planning for each sampling campaign, SNC is responsible for evaluating the health and performance of the SCS for each sampling operation. The Mech Data Tools library is used to query, collate, and visualize summary metrics for each abrasion and core, which are then presented quarterly to the larger project [6]. Overall, the SCS has performed to expectations, though the team has improved their knowledge of how to operate this complex system over the first 562 sols of the mission.

To date, Mars 2020 has collected 15 samples. When awaiting data for a new sample, the team is always eager to record the sample height measured by the volume probe inside the ACA (Figure 17). With the exception of the Roubion coring attempt, the SNC team commands coring to a depth of 66 millimeters. However, the length of sample collected will often not match this depth. The rock core may fracture as it moves into the sample tube, leaving unfilled space in the tube, or pieces of the core may fall out after the core breaks away from the rock. If the measured sample height does not meet a parameterized threshold, the ACA will not seal the tube, allowing the operations team to decide if they wish to seal the underfilled tube on a later sol. Most of the samples collected to date have met this required threshold, with the Malay and Coulettes cores falling on the shorter side. In addition, the Pauls sample does not have a volume probe measurement because it was intentionally dumped in response to the Bit Carousel pebble anomaly.

Prodapt Levels— Drilling on Mars 2020 leverages a proprioceptive-adaptive or “prodapt” algorithm. Prodapt targets a prescribed rate of penetration (ROP) by varying the weight on bit (WOB), percussion frequency, and spindle (the rotational actuator) rate. If the controller notices it is drilling too slowly, it will increase the prodapt level, which corresponds to a new triplet of WOB, percussion frequency, and spindle rate. On the other hand, if the rock is on the softer side, the prodapt level will decrease to reduce the ROP. Note that the desired ROP for coring is 10 times greater than the ROP for abrading. Prodapt levels vary on a scale from 0 to 21 for coring and 3 to 21 for abrading. Below level 3, coring will no longer use percussion, referred to as rotary-only coring; abrading will always use some amount of percussion. If the drilling operation hits level 21, a fault will be triggered due to lack of progress in the rock, as was the case with target Atsah visualized in the plot in Figure 18. Prodapt level data are of interest to the SNC team because they are an indicator of how difficult the rock was to drill.

The prodapt algorithm is seeded with an initial prodapt level. For coring, the seeded level depends on data obtained during the start hole operation, which corresponds to the first 5 millimeters of coring. For abrading, the wider diameter of the abrading bit and shallower hole start operation of 2 millimeters means initial data largely depends on surface topography and therefore is not useful for seeding prodapt level. Instead, operators can choose between the default value and a soft setting, which starts the prodapt portion of abrading at a lower prodapt level. Upon arriving to the Delta region on

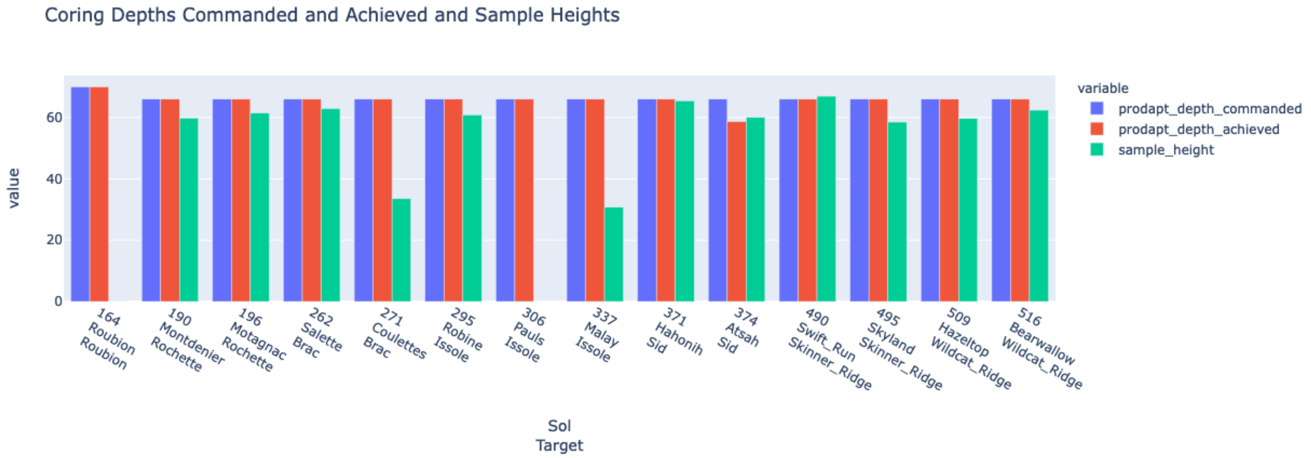

Figure 17. Commanded and achieved coring depths compared to measured sample height for each sample.

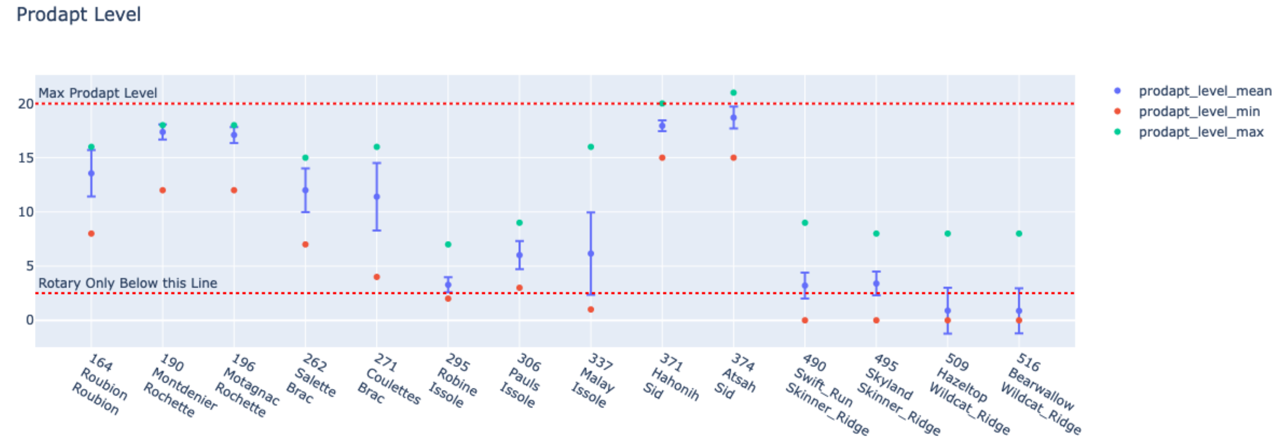

Figure 18. Prodapt levels plotted for each coring attempt.

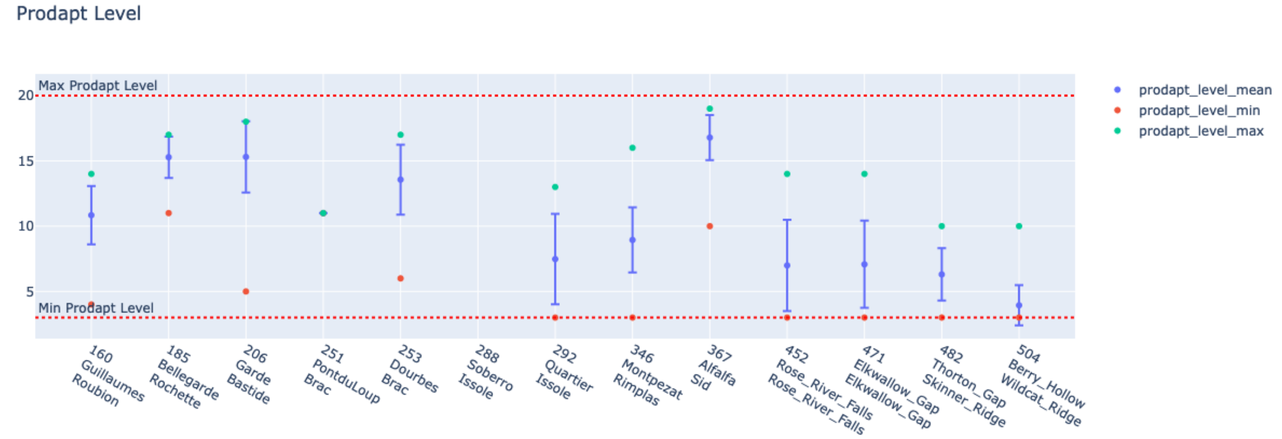

Figure 19. Prodapt levels plotted for each abrading attempt.

the western border of Jezero Crater, beginning with the Rose River Falls abrasion on sol 452, the team encountered rock that would fracture during abrading operations. Fractured abrasions make it difficult or impossible to safely place proximity science instruments on the abraded patch, so the team decided to switch to seeding prodapt with the soft prodapt level. This led to successful Delta region abrasions at Thorton Gap and Berry Hollow (see Figure 20). The softer rocks in the delta were more likely to hold together when less percussive energy was put into the rock. SNC made the decision to seed with the soft level for the rest of the Delta sampling campaign.

Figure 20. Fractured Elkwallow Gap abrasion (left) compared to successful Thorton Gap abrasion (right) seeded with soft prodapt level.

Weight on Bit and Robotic Arm Forces—While drilling, WOB and the force-torque sensor on the RA are monitored to ensure they do not exceed the parameterized fault protection limits. This is for the protection of the hardware, ensuring the forces and moments experienced during surface interactions do not exceed safe limits. During development, the SCS team encountered several instances in which the preloaded stabilizers "walked" across the rock during abrading or coring, tripping these limits or related fault protection. This phenomenon is commonly referred to as stabilizer slip, and has also occurred in flight with the Pont du Loup (Figure 21) and Soberro (Figure 22) abrading targets. In those cases, the stabilizers were either placed in a location with significant topography or placed on a small pebble.

In response to these incidents, SNC has implemented several operational improvements. The target assessment process now requires SNCs to avoid pebbles over a certain size and has limits on the topography, material, and stability of the stabilizer patch. To help with this, SNC also increased the height at which the high standoff WATSON images are taken during target assessment to expose more of the surrounding rock. A stabilizer slip recovery procedure has been developed to help operators verify if a fault was due to stabilizer slip and recover from it quickly. Trending has been established to identify cases in which the WOB and RA forces and moments changed significantly while drilling, allowing the team to respond to potential slip. These practices have helped reduce stabilizer slip faults, the last of which occurred on sol 288.

Bit Carousel Anomaly

On sol 306, the operation to return the sampling bit containing the Pauls sample back to the BC failed during bit insertion. The drill feed was unable to extend to fully insert the sampling bit into its bit holder. The team compared the WOB seen during this motion to the same motion on sol 298, when the coring bit with sample drop off was successful as shown in Figure 23. The amount of force seen by the bit while extending to interface with its bit holder differed significantly.

After some initial recovery actions, the team succeeded in capturing diagnostic images of the BC upper door using the WATSON camera. These WATSON images showed several pebbles around the bit holder, visible in Figure 24, which the team theorized may have fallen out of the bit during the spindle clocking moves that occur just prior to extending the feed to insert the bit into its bit holder for bit dropoff. In order to safely allow movement of the BC in this scenario, the speed of the BC motor was significantly reduced and the resolver-encoder miscompare fault protection was tightened. This protected against hardware damage if a pebble were to cause a jamming effect. The ability to move the BC enabled better imaging and eventually, in combination with rover mobility, helped clear the pebbles from around the bit holder.

Later, the team used the Mastcam-Z camera to record a video of the drill recreating the clocking motions of the bit dropoff activity at an identical gravity orientation. A pebble was expelled from the bit in the video, supporting the theory that such spindle motions can cause pebbles to fall out of the bit.

Following the conclusion of the anomaly and return to sampling operations, the team implemented several corrective actions. A BC Range of Motion (ROM) check was added to the sampling sol path. After driving or coring bit dropoff with sample, SNC must now command a BC ROM check before moving the Bit Carousel at the old, faster speed. This check uses the tightened fault protection at a slower rotor speed to check for jamming at the close clearances inside the Bit Carousel. This reserved additional torque margin so that in a jamming case, there would be further torque applied before hitting the limit. Along with this, the spindle clocking moves are now performed in free space far from the BC, reducing the risk of pebbles falling onto hardware during the sampling bit dropoff operation. With these corrections, the team has so far avoided a recurrence of the anomaly.

6. HELICOPTER

Ingenuity, the Mars helicopter, deployed from Perseverance and successfully completed its initial thirty-day mission of demonstrating powered flight on Mars. Since then, Ingenuity has been flying along Mars 2020's strategic path, keeping in comm range of Perseverance. In total, Ingenuity has completed 32 flights during the first 562 sols of Mars 2020.

Along the way, Ingenuity has taken imagery of the Martian surface. On sol 414, Ingenuity acquired imagery of the rover's EDL (Entry Descent Landing) hardware using the Ingenuity's RTE (Return To Earth) camera as seen in Figure 26.

Ingenuity ran into challenges including thermal, power, and comm. On sol 427, Perseverance was unable to communicate with Ingenuity at its planned wake-up time. Due to lower amounts of energy provided by the winter sun, dust in the atmosphere occluding the solar panels, and lower temperatures, Ingenuity was unable to stay powered throughout the night. Instead, it experienced brownouts, which reset its clock. As a result, it was unable to wake up at the time requested by Perseverance. Instead, it woke up whenever it happened to accumulate enough charge in the morning hours. The Ingenuity team was able to predict when this would occur and updated Perseverance to communicate at that time to reestablish communications. After successfully completing recommissioning activities and successfully completing flight 29, Ingenuity returned to its normal pattern of having regular flights with flights 30 through 32.

Rover orientation at the end of a drive can influence its ability to communicate with Ingenuity. HIEs communicate potential helicopter comm issues to the RPs, who try to accommodate requests for particular final headings, but must first prioritize other orientation constraints. These include avoiding parts of the rover occluding line-of-sight from the High Gain Antenna to Earth, giving science instruments a view of nearby science targets, and enabling arm placement for contact science. Figure 27 illustrates the potential for rover self-occlusion.

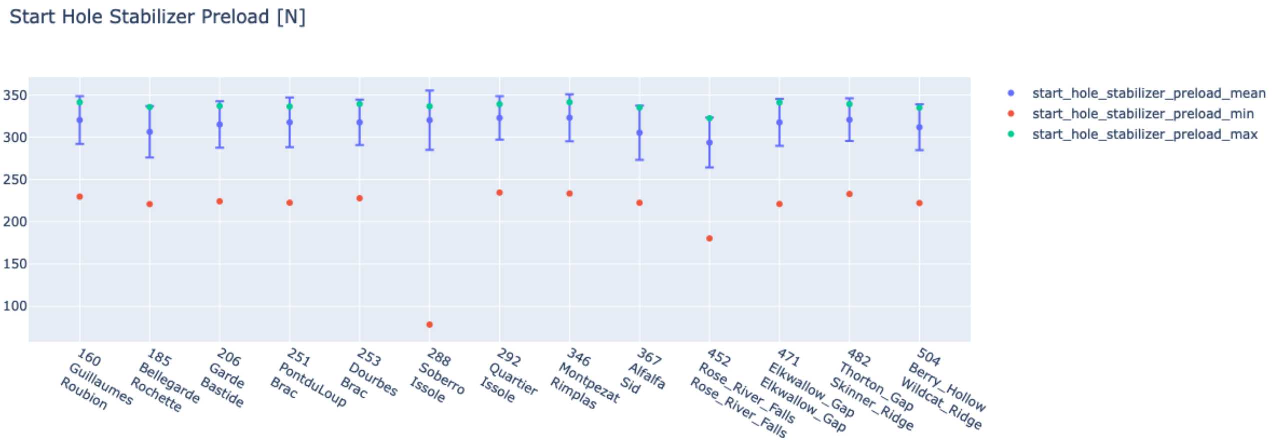

Figure 21. Stabilizer preload during start hole, the first 2 millimeters of abrading.

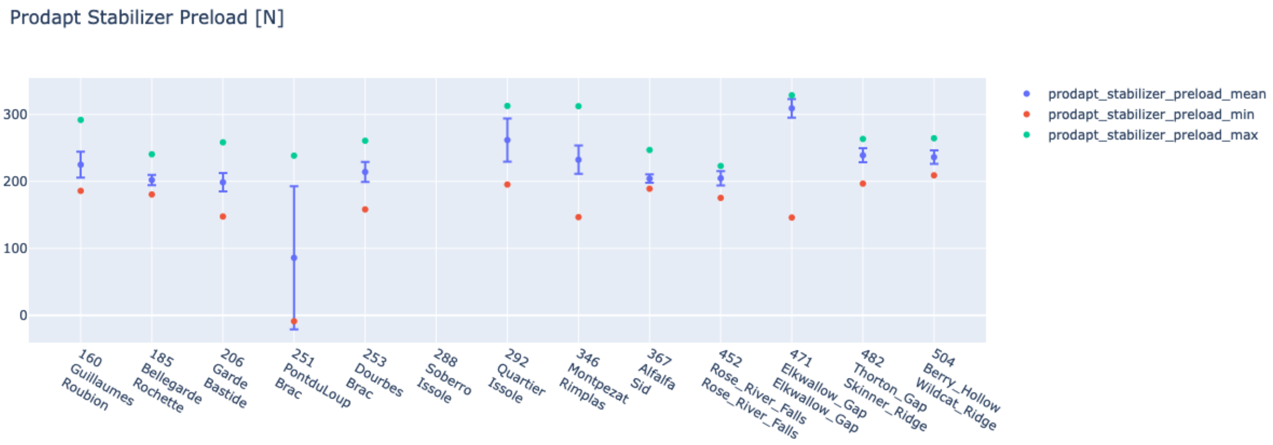

Figure 22. Stabilizer preload during prodapt drilling.

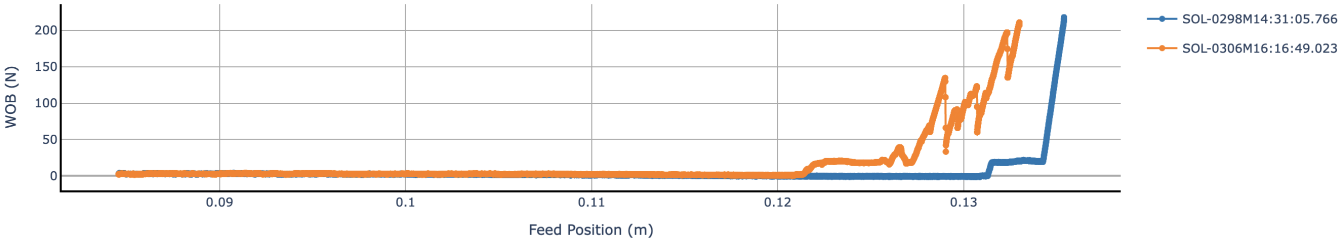

Figure 23. Plot of the feed position against the Weight on Bit for the sol 306 bit dropoff fault and an identical, successful motion on sol 298.

Table 3. Flight Software Transition dates

Sol	Date	Release	Description
00000	Feb 17 2021	C4_2_0	Cruise version
00005	Feb 23 2021	S6_4_0	Initial Surface version
00014	Mar 4 2021	S6_4_1	Primary Surface version
00243	Oct 25 2021	S7_1_0	156 Anomaly fixes, 66 Feature requests, 46 Engineering Change Requests
00468	Jun 13 2022	S7_2_0	62 Anomaly fixes, 36 Feature requests, 30 Engineering Change Requests

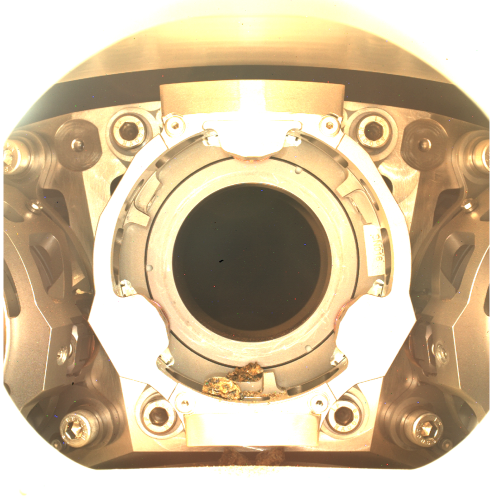

Figure 24. WATSON images showing pebbles visible in and below the Coring Bit 2 bit holder at the Bit Carousel upper door on sol 314.

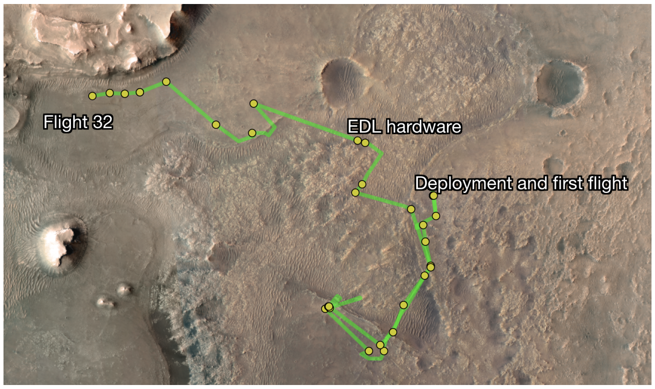

Figure 25. The paths of Ingenuity's flights 1 through 32.

7. ROBOTIC OPERATIONS FLIGHT SOFTWARE UPGRADES

Flight software (FSW) upgrades can improve science return, mitigate known problems or idiosyncrasies, and provide enhanced capabilities. Although the FSW team formally resides in the EO organization, the RO team has significant involvement in FSW updates. Several of the RO staff are themselves current or former FSW team members who own most of the robotics FSW modules. Most of the V&V test procedures run for the mobility and arm FSW capabilities are written and executed on the testbed vehicle by RO members. And some of the benefits of the new capabilities can only be realized by incorporating them into the ground tools that generate plans and the data visualizers that display their results.

The collection of all parameters used and system state maintained by all FSW modules is referred to as the Nonvolatile Parameter Manager (NPM) state. Every planning day, the current NPM state (the combined current values of system parameters and vehicle state) is a starting seed for operations planning; all the values are read into Surface Simulation (SSim) software, which simulates FSW on ground workstations. Knowledge of this data, and the ability to incorporate it into the plan, is critical for robust operations planning.

Figure 26. Image of EDL hardware taken by Ingenuity's RTE camera.

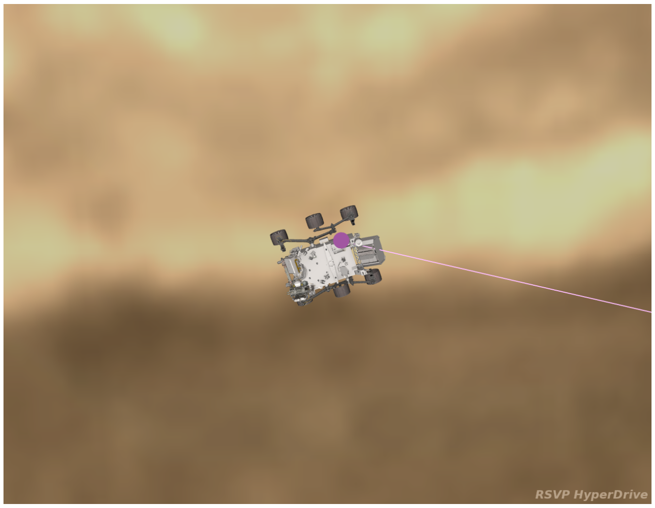

Figure 27. An example of a non-optimal Heli comm situation with occlusions from the Low Gain Antenna and Radioisotope Thermoelectric Generator.

Each new flight software version typically includes a number of changes to this state, with new parameters and new or modified system state variables. So each new FSW release results in a new NPM state that is often incompatible with the previous flight version. This can be problematic for simulation purposes as plans built prior to the flight software transition (and thus containing obsolete NPM) may no longer work with the latest flight software and simulation tools that depend on the flight software. A new tool on Mars 2020 called NPM Direct made it possible to avoid this problem. In additional to its primary goal of decoding NPM, it can also upgrade older NPM information into a format that matches the latest flight software. This feature has helped us in testing our operations tools and flight software by reusing hundreds of previous plans from previous sols in the latest tools that must assume the newest NPM state structure.

Table 3 shows the progression of updates made to the FSW throughout the mission so far.

The first FSW update took place in the first week of operations, as the project transitioned from the FSW version used to travel to Mars and land safely (the Cruise, Entry, Descent and Landing version) to the first Surface version of the FSW. The cruise version did provide some rudimentary surface commands, but not advanced M2020 capabilities. Happily, the transition to full M2020 FSW happened quickly and enabled a quick start to nominal mission checkout operations.

The next versions provided helpful improvements to existing capabilities. Some bugs preventing the use of a feature enabling automated scheduling of command sequences during a drive were eliminated. Bugs discovered by ”software fuzzers” in which extreme values of certain command arguments could result in an unexpected reboot were eliminated. And whole new capabilities were added as well, including updates to the AEGIS science autonomy capability, and future onboard planner and autonomous arm unstow capabilities.

Software Anomaly Diagnosis and Resolution

The Perseverance FSW follows a fault protection model similar to that used on the earlier MER and MSL missions. One aspect is its use of ASSERT statements to ensure that current arguments and system state are within expected ranges. These ASSERTs remain active throughout the entire mission lifetime, including surface operations. Whenever an ASSERT fails due to an unexpected condition, message, parameter or state, the response of the system is to gracefully stop all current operations, shut down the software and reboot into a Safe Mode. ASSERT failures and operating system errors (such as illegal math operations like dividing by zero) will each result in that fault response behavior, typically issuing a terse summary of the issue as a ”FATAL” Event Report.

Thus far only a single drive has been interrupted due to a software FATAL, on sol 390 as shown in Figure 7. That was in the middle of the month-long Rapid Traverse period, and it led to four additional days of non-driving while the team recovered from the event (the first two days were non-planning weekend days on Earth, the actual recovery only took two days). In this case the team quickly narrowed down the cause of the problem and developed a solution. The problem was due a particular code path through the pose-maintaining FSW’s controlling state machine. A message requesting that information about the current pose be saved as an individual Data Product file arrived at an unexpected time; the state machine did not have a response implemented for that message in its current state. That led to the assertion failure and the FATAL. But the team recognized that the errant control path could be avoided completely by simply disabling the writing of that individual product. The same information was being provided elsewhere, and downlink teams confirmed that they would not have any significant degradation in their ability to assess spacecraft state once that change was made. The RO team then updated the commanding macros to ensure that change would be enacted in all future plans, following established processes for validating and reviewing such changes, and driving resumed soon after.

8. ROBOTIC OPERATIONS GROUND TOOLS

One of the major tools we use to operate the rovers is the RSVP (Robot Sequencing and Visualization Program) tool suite. These tools are developed by the Rover Planning Subsystem (RPS) and RO, and are the latest generation of software used to drive the MER rovers [7], upgraded to help the Mars 2020 mission achieve a shorter planning cycle. The suite includes tools that automatically generate rover command sequences, such as ArmSketch and MobSketch.

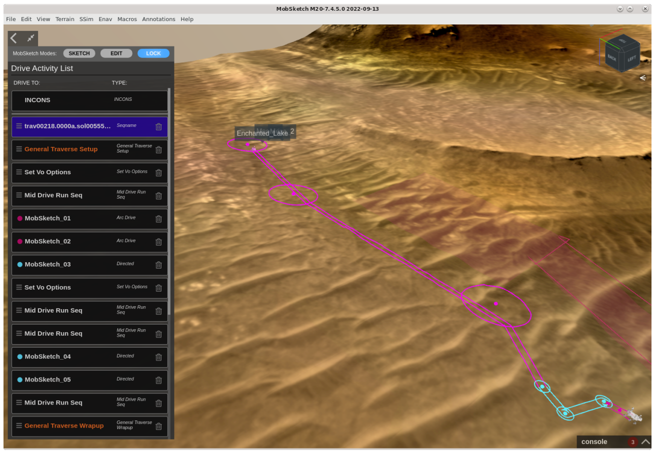

DRIVE TO:	TYPE:
INCONS	INCONS
trav00218.0000a.sol00555...	Seqname
General Traverse Setup	General Traverse Setup
Set Vo Options	Set Vo Options
Mid Drive Run Seq	Mid Drive Run Seq
MobSketch_01	Arc Drive
MobSketch_02	Arc Drive
MobSketch_03	Directed
Set Vo Options	Set Vo Options
Mid Drive Run Seq	Mid Drive Run Seq
Mid Drive Run Seq	Mid Drive Run Seq
MobSketch_04	Directed
MobSketch_05	Directed
Mid Drive Run Seq	Mid Drive Run Seq
General Traverse Wrapup	General Traverse Wrapup

Figure 28. MobSketch with mobility activities for sol 555, displayed with orbital mesh. This is a screen capture from the actual tool, so the font size of Activity List names in the black box appears relatively small here.

Another tool is SSim (Surface Simulation), which uses flight software code to predict the rover’s response to commands in a simulated reconstruction of the Martian environment.

These tools are regularly updated in response to needs found in tactical operations. The RO team tests these tools every few weeks for a new official deployment. In addition to official deployments, there are automatic nightly builds/deployments that allow users to test out new features immediately after RPS developers make them available. This model of development allows for tool updates to be made available to users relatively quickly. This is a contrast to the traditional model of other missions where deployments take several months, which results in operators having to work around tool inefficiencies and bugs for extended periods of time.

MobSketch

MobSketch is a high-level traverse planning tool used by rover planners for planning mobility activities. Designed as a 3D tool which uses Three.js, React and Electron, MobSketch is an improved version of a 2D traverse planning tool currently used on MSL called RSketch, shown in Figure 29. The design and architecture of both MobSketch and its mobility macro infrastructure enables rover planners to change their command sequences throughout the planning day.

Rover planners use MobSketch by loading 3D terrains generated from data downlinked from Mars, as seen in Figure 28. By clicking on the 3D terrains displayed in MobSketch, rover planners are able to create waypoints. Existing waypoints can be easily modified by clicking and dragging to adjust their positions or by editing the drive length and heading in the activity list. Each waypoint forms a drive leg or a mobility activity, which corresponds to an associated macro, and is represented as an entry in the activity list on the left side of MobSketch’s display (see the activity list on the left side of Figure 28). Each mobility activity and its corresponding macros have parameters that can be edited in the activity list in MobSketch by expanding its activity list entry. Figure 30 shows an example of an expanded activity list entry for a single mobility activity.

“No motion” macros (i.e. macros that contains only commands that do not cause the rover to move) and turn macros

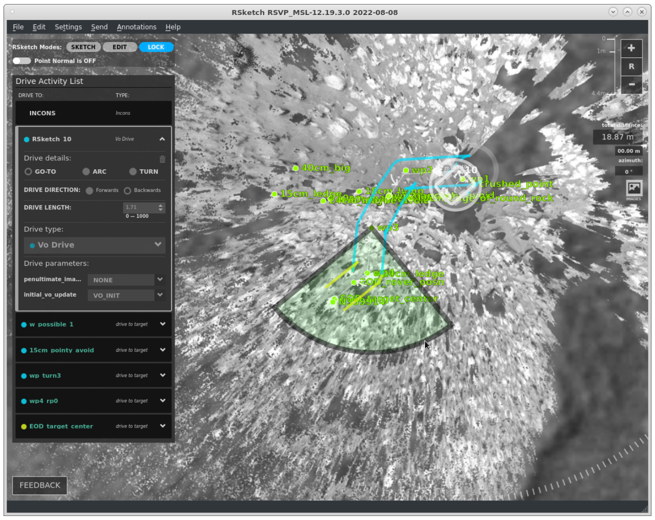

Drive details:
○ GO-TO ● ARC ● TURN
DRIVE DIRECTION: ● Forwards ○ Backwards
DRIVE LENGTH: 1.71 (0 — 1000)
Drive type: Vo Drive
Drive parameters:
penultimate_ima...: NONE
initial_vo_update: VO_INIT

Figure 29. MobSketch predecessor, RSketch for MSL, is a legacy tool shown for illustration purposes only.

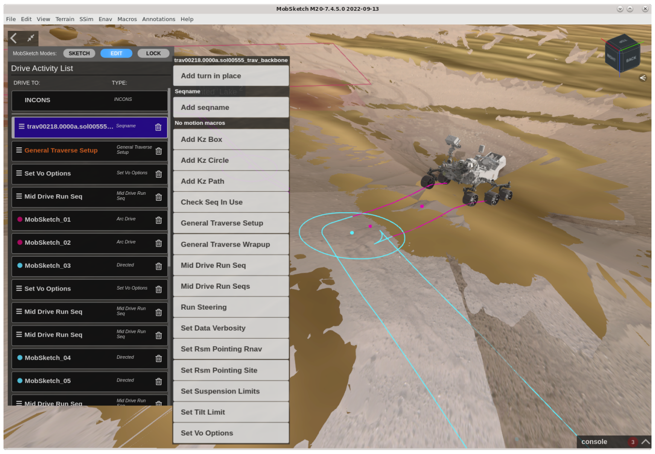

DRIVE TO:	TYPE:
INCONS	INCONS
trav00218.0000a.sol00555...	Seqname
General Traverse Setup	General Traverse Setup
Set Vo Options	Set Vo Options
Mid Drive Run Seq	Mid Drive Run Seq
MobSketch_01	Arc Drive
MobSketch_02	Arc Drive
MobSketch_03	Directed
Set Vo Options	Set Vo Options
Mid Drive Run Seq	Mid Drive Run Seq
Mid Drive Run Seq	Mid Drive Run Seq
MobSketch_04	Directed
MobSketch_05	Directed
Mid Drive Run Seq	Mid Drive Run Seq

trav00218.0000a.sol00555_trav_backbone
- Add turn in place

Seqname
- Add seqname

No motion macros
- Add Kz Box
- Add Kz Circle
- Add Kz Path
- Check Seq In Use
- General Traverse Setup
- General Traverse Wrapup
- Mid Drive Run Seq
- Mid Drive Run Seqs
- Run Steering
- Set Data Verbosity
- Set Rsm Pointing Rnav
- Set Rsm Pointing Site
- Set Suspension Limits
- Set Tilt Limit
- Set Vo Options

Figure 31. Activity list with menu in MobSketch, displayed with both navcam and orbital meshes. This is a screen capture, so the font size of Activity List names in the side boxes appears relatively small here.

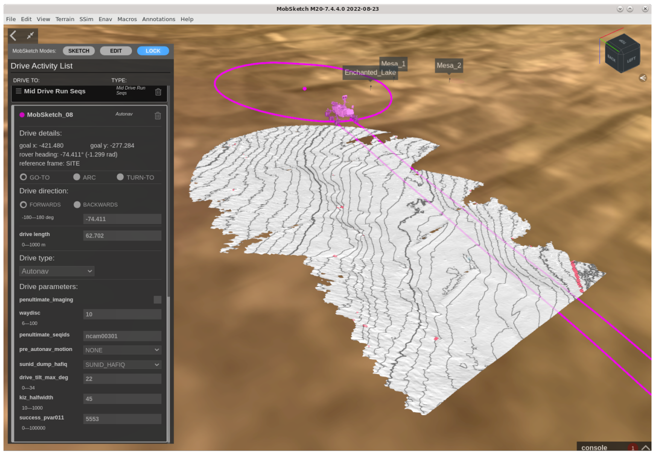

Drive details:
goal x: -421.480 goal y: -277.284
rover heading: -74.411° (-1.299 rad)
reference frame: SITE
○ GO-TO ● ARC ● TURN-TO
Drive direction:
○ FORWARDS ● BACKWARDS
-180—180 deg: -74.411
drive length (0—1000 m): 62.702
Drive type: Autonav
Drive parameters:
penultimate_imaging
waydisc (6—100): 10
penultimate_seqids: ncam00301
pre_autonav_motion: NONE
sunid_dump_hafiq: SUNID_HAFIQ
drive_tilt_max_deg (0—34): 22
kiz_halfwidth (10—1000): 45
success_pvar011 (0—100000): 5553

Figure 30. Activity list and Enhanced Nav heightmap displayed in MobSketch. This is a screen capture, so the font size of Activity List names in the black box appears relatively small.

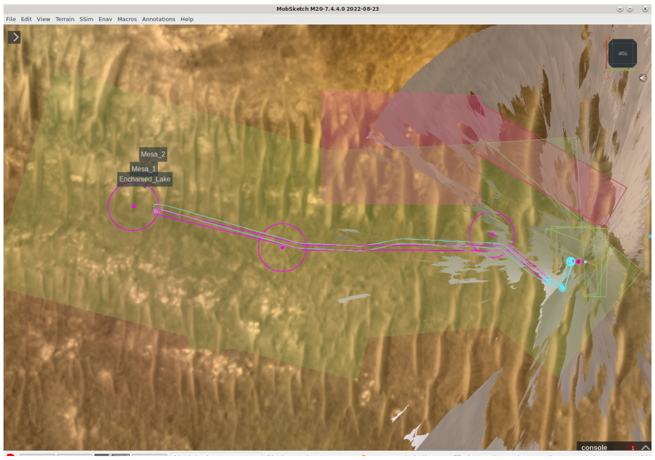

Figure 32. SSim simulation results displayed in MobSketch (in blue). This is a screen capture, so the font size of target names appears relatively small here.

can be added using a menu that appears by right clicking on the activity list, as seen in Figure 31. After updates are made to mobility activities, MobSketch automatically regenerates the full sequence of commands for all activities in the activity list. MobSketch displays simulated drive tracks along with various annotations generated after the full sequence of commands is generated and simulated using SSim. Figure 32 shows an example from Sol 555 of the planned and simulated drive tracks along with annotations such as keep in and keep out zones over Navcam and orbital image meshes in MobSketch. By allowing users to easily modify mobility activities and view the simulation results, MobSketch enables rover planners to rapidly change the mobility sequences, up until locking down all changes for the uplink delivery process. This is in contrast to the standard practice on the MER and MSL missions in which RPs would have to freeze the drive development earlier, to allow time for hand-tuning of the commands in the mobility sequence.

A new feature was added to MobSketch post-landing, the visualization of Enhanced Navigation (ENav) heightmaps, which are terrain heightmaps generated during autonav drives and downlinked as data products, as seen in Figure 30. The visualization of these terrains in MobSketch enables rover planners to better understand what terrain is underneath the rover following Autonav drives (human-directed drives will provide this data by manually adding ”penultimate image pairs”, taken some 5m from the final position). Another feature added was the ability to plan multi-sol drives. This capability was added by updating the associated macros, without requiring significant updates to the MobSketch user interface, thanks to the expressiveness supported by the existing macro architecture. Figure 28 shows the minor UI updates, where the start of each sol in the multi-sol drive is denoted by highlighting the sequence activity in dark blue.

ArmSketch

ArmSketch is a new software tool that generates command sequences for arm activities. It takes as input a list of arm activity objectives (also called the *plan*) from a higher-level planning tool called COCPIT [8], which the mission uses to manage the rover’s activities across multiple subsystems. After importing the plan, ArmSketch users visualize the arm poses involved in each activity and input information specific to that sol's activities. ArmSketch then computes the optimal order of arm activities and computes the commands and arm movements that are necessary in between each pair of activities. Finally, ArmSketch generates the full sequence of commands for the arm activities in the plan.

At a low level, an arm activity corresponds to a *macro*, which is a parameterized command sequence template. Template parameters are exposed as user inputs in the ArmSketch user interface. For example, a boolean argument would appear as a checkbox in ArmSketch, and checking or unchecking the box could cause the macro to include or exclude a block of commands. The plan that ArmSketch imports from COCPIT includes a list of arm activities as well as values for some of the macro parameters. Users can update the macro parameters in ArmSketch after importing the plan, if needed.

Each arm activity has a starting and ending arm pose. These poses are either given as a standoff distance from a target or a joint-space pose. For targets, ArmSketch computes the inverse kinematics solutions and presents for user approval. There are up to two Inverse Kinematics solutions per target. ArmSketch also checks the poses for collisions and only displays reachable and non-colliding poses to the user.

ArmSketch also adds commands between arm activities, which we call in-between moves. In-between moves take the arm from the end state of one activity to the start state of the next activity. Some of these commands simply set the FSW state appropriately (e.g., telling the FSW what the next arm target is). In cases where a straight-line path between the two activities would result in a collision with the terrain, ArmSketch adds in-between moves to avoid collisions. ArmSketch also has logic to obey flight rules in between activities. For example, to minimize the risk of dust settling on the WATSON camera's lens, we do not point the WATSON camera above the horizon while its dust cover is open. If two activities both use WATSON with the dust cover open, and ArmSketch cannot find an in-between path that keeps WATSON pointed down, then ArmSketch will instead add commands to close and later reopen the WATSON dust cover.

ArmSketch also computes the optimal ordering of activities to minimize the plan duration. For example, if there are two activities that both use WATSON, it will place them together to avoid closing and reopening the WATSON cover in between. In general, finding the optimal order of activities given a cost to go between them (the duration of the in-between moves) is the travelling salesman problem, which is NP-hard. We solve this problem exactly using brute force. In practice, the plan imposes constraints on the order of activities. For example, we would not use gDRT to clear an abrasion patch without first completing the abrasion activity. Constraints like this drastically reduce the number of search paths. We also display a progress bar showing the number of paths evaluated. Based on that, users can judge whether to cancel the computation and constrain the problem further.

Figure 33 shows a typical workflow for ArmSketch. As of sol 562, ArmSketch has generated 233 out of 300 arm sequences. RPs did not use ArmSketch for commissioning the arm after landing, nor do they generally use it for fault investigation, fault recovery, and other one-time activities.

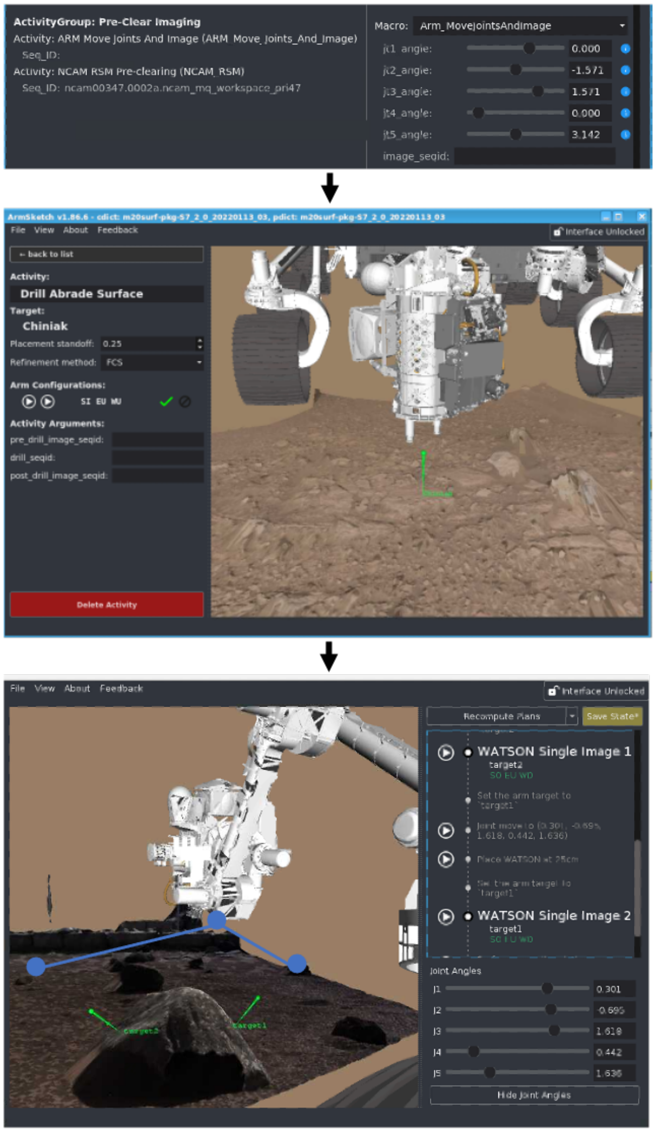
ActivityGroup: Pre-Clear Imaging
Activity: ARM Move Joints And Image (ARM_Move_Joints_And_Image)
Seq_ID:
Activity: NCAM RSM Pre-clearing (NCAM_RSM)
Seq_ID: ncam00347.0002a.ncam_mq_workspace_pri47

Macro: Arm_MoveJointsAndImage
jt1_angle: 0.000
jt2_angle: -1.571
jt3_angle: 1.571
jt4_angle: 0.000
jt5_angle: 3.142
image_seqid:

ArmSketch v1.86.6 - cdict: m20surf-pkg-S7_2_0_20220113_03, pdict: m20surf-pkg-S7_2_0_20220113_03

Activity:
Drill Abrade Surface
Target:
Chiniak
Placement standoff: 0.25
Refinement method: FCS
Arm Configurations:
SI EU WU ✓
Activity Arguments:
pre_drill_image_seqid:
drill_seqid:
post_drill_image_seqid:

WATSON Single Image 1
target2
S0 EU WD
Set the arm target to `target1`
Joint move to (0.301, -0.695, 1.618, 0.442, 1.636)
Place WATSON at 25cm
Set the arm target to `target1`
WATSON Single Image 2
target1
S0 EU WD

Joint Angles
J1 0.301
J2 -0.695
J3 1.618
J4 0.442
J5 1.636

Figure 33. An example workflow using ArmSketch. The first image shows the interface for importing a COCPIT plan. The middle image shows the user interface for editing arguments and visualizing arm poses. The third image shows an example of an ArmSketch-generated plan. The blue dots and lines show the ArmSketch-generated path to avoid colliding with the central rock. This is a screen capture from the ArmSketch tool, so the font size of names in the Activity List found in the black boxes appears relatively small here.

SSim

SSim (Surface Simulation) is a flight-software-in-the-loop simulation tool [9] used to predict behavior of Perseverance given a set of command sequences. Including flight software in the loop is essential given the complex interactions between Perseverance's software, hardware, and the Martian surface. SSim takes in as input the command sequences, initial parameters and persistent onboard state encoded in NPM (Nonvolatile Parameter Manager) records, the initial kinematic configuration, terrain elevation maps, and relevant onboard files. It executes FSW with these simulation inputs and outputs the resulting text message Event Records (EVRs) and kinematic state history. This allows operators to quickly determine whether their command sequences achieve the intent. RO's key goal in developing SSim is to have a simulation tool that quickly and deterministically simulates the behavior of the flight software. These attributes enable quick iteration as we develop command sequences tactically.

SSim is used throughout the planning shift by the RO team in an iterative fashion. Logs from the past year show that on average, RO runs about 60 simulations per planning day. On average, each simulation takes around 20 seconds. Though, a simulation can take a couple minutes if it encompasses three sols of activity, as is often the case for the weekend plans.

SSim is also used to automatically find bugs in flight software. One key component of SSim's extensive automated testing suite is fuzz testing. The fuzzers automatically generate command sequences based on the flight software command dictionary. SSim's fuzzers include a fuzzer that uses the coverage-guided AFL (American Fuzzy Lop) fuzzer and a home-grown Python-based fuzzer that fuzzes the SSim web service, which is a thin web wrapper that makes SSim available to web clients. The fuzzers pipe random command sequences through SSim and detect unexpected or undefined behaviors. Fuzzing has helped find issues in the simulation software itself, but also found tens of fatal bugs in flight software. Finding and fixing these flight software bugs early meant we avoided hitting them during Mars operations.

NPM Direct

NPM Direct is a tool we developed to extract the onboard state of the Mars 2020 rover from its flight software's Non-volatile Parameter Manager opaque data stores. The extracted state is a key input to SSim, which is required to initialize flight software's state. NPM Direct's output is also used by various other tools including RPyCheck, the static analysis tool used to validate Rover Planner command sequences. In operations, the data from NPM is dumped from Perseverance and transmitted as a few hundred kilobytes of compressed data. This data is uncompressed and then decoded by NPM Direct to reveal the underlying several hundred thousand variables. These variables' types are in the form of integers, floating point numbers, strings, and enumerations.

NPM Direct is an innovation from our previous mission, (MSL) Mars Science Laboratory, where a combination of handwritten and autocoded scripts had to be written to infer the NPM state based on EVRs, Channelized Engineering Health and Accountability scalars (EHA), and binary data products. NPM Direct avoids the need for laboriously writing these scripts, which can be error prone. Instead, NPM Direct directly extracts the NPM state from the NPM module itself. The challenge is that the NPM dump is an opaque memory dump and the computer architecture of the rover on Mars is PowerPC whereas SSim runs on x86. NPM Direct automatically analyzes flight software to document the memory layout of each memory dump as a string similar to Python's struct format. This NPM Direct format string is able to describe the true byte of any structure or array and allows a mapping of raw bytes between different platforms. NPM Direct will decode the raw NPM records downlinked from the rover and process them into a JSON file that no longer has platform dependency. NPM Direct can then re-encode this JSON file for various architectures including SSim's x86 platform.

Since the initial successful roll out of NPM Direct at the beginning of the Mars 2020 mission, a new release has been made. This new release reimplements the backend to use the LLVM compiler's "Pass" API similar to AddressSanitizer [10] to extract NPM records from flight software more efficiently (in space and time by orders of magnitude) and to allow NPM Direct to extract all of the NPM records without requiring any special cases. By using the LLVM compiler, NPM Direct is able to provide layouts for any platform supported by LLVM. This version also supports upgrading old flight software's NPM for use by newer flight software by way of a scripted transform. This ability to automatically re-express NPM from an older version of FSW has allowed Mars 2020 to test new flight software against the existing plans originally used on Mars with older FSW.

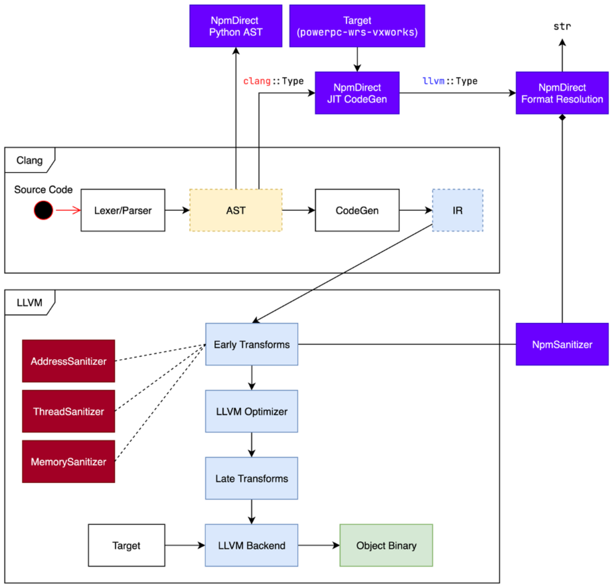

Figure 34. A data flow diagram showing NPM Direct using LLVM to automatically infer NPM dump memory layouts.

NPM Direct will compile SSim with an LLVM compiler pass called NpmSanitizer. This pass is able to detect loads/writes of NPM records as well as additional metadata such as the name of the symbol being written to disk and a string describing its true layout in memory. NPM Direct will then combine information received and cached at runtime about the layout of the record as well as offline information from the source code such as structure field names, to build its platform layout. Figure 34 shows where the pass will be hooked into the rest of the compiler pipeline. Once the instrumented SSim is run and each record is described relative to flight software structures, NPM Direct is able to use its Just-in-Time (JIT) compiler to describe the resultant memory layout on any target platform. NPM Direct's JIT Compiler is an adaptation of LLVM's official CodeGen stage shown in Figure 34 which takes a parsed abstract syntax tree (AST) that describes the source code and lowers it into LLVM's intermediary representation (IR) to apply transformations later down the pipeline. NPM Direct's JIT CodeGen will instead be able to lower a single structure to IR and then query this structure's memory layout on the target platform. For example, a user could query the resultant memory layout of powerpc-wrs-vxworks for any C structure in our flight software as well as its layout on i386-pc-linux-gnu. These two platforms are used onboard the rover and to build SSim respectively. By knowing where records are loaded into these C structures, a job performed at runtime by the instrumented SSim binary, we can use this functionality to decode and encode binary data into/from these two platforms providing a compatibility layout between downlinked data and simulation.

The success of NPM Direct on M2020 resulted in MSL requesting a backport onto their mission for a select use case.

NPM Direct will allow MSL to test its upcoming R13 flight software against the thousands of existing R12-based plans previously executed on Mars [11].

9. CONCLUSIONS

Here are some lessons learned related to the creation of the RO organization from the first 1.5 years of surface operations,

Mission success created un-forecasted RO staffing issues

Ingenuity was originally intended to only complete a 30-sol technology demonstration, after which three members of the HIE team (who were already certified for RP) would increase their RP availability and one HIE team member would start RP training. But due to its tremendous success, Ingenuity's mission was extended, contributing to a shortage of RP staff. The HIEs had planned to move to other roles or projects, and all but one was able to delay their transition. An HIE training program was also developed to train a second generation of HIEs. As we move toward an extended mission with reduced staffing, HIE and HO teams have begun cross-training with the possibility of merging into a single team.

RO team structure has greatly facilitated tool development

The Robot Sequencing and Visualization Program (RSVP) suite of tools, developed by the Rover Planning Subsystem (RPS) team, is exclusively used by all of the RO roles during tactical shifts. One month after landing, an agile engineering point-release test process was developed to quickly respond to RSVP needs identified by the RO tactical team. After a development period where RPS fixes reported bugs and implements new high-priority features, the RO roles of RP, SNC, and HIE all participate in regression testing of RSVP candidate point releases, typically over a two-week period. 12% of RO team members are also part-time RPS developers, facilitating close coordination between the teams. Since landing, there have been 29 RSVP engineering point releases.

Training RPs from within RO is ideal but a staffing challenge

RP certification training currently takes approximately one year to complete at a half time commitment. As greatly employed in the MER mission and to a lesser extent in MSL, the Mobility and RA downlink roles are a useful pipeline for the RP role, since candidates from those teams start with a working knowledge of Perseverance's robotic capabilities and RSVP. In the first RP training class after landing, all six trainees came from within RO teams; three from Mobility, two from RA, and one from HIE. Staffing the RP role from within RO is ideal, however, it creates the staffing challenge of sufficiently over-staffing the RO downlink teams in advance of an RP training class to maintain sufficient certified staff once an RP training class starts.

The RO organization structure has facilitated robotic testing

When RO team members are not scheduled for a tactical shift, their remaining M2020 time is filled with strategic work, such as residual testbed Verification and Validation (V&V) or FTA testing. It is common for testbed shifts to be staffed by multiple RO roles. For example, a testbed test of sample caching products would involve representatives from both RP and SNC. To coordinate testbed training and test shifts for the RO roles, an RO testbed lead was selected. The RO testbed lead and testers meet weekly to review activity reports from the previous week, upcoming test shifts, staffing, and preparation for shifts.

Collocating RO tactical uplink roles in same room is ideal

Given SNC sequences are called from an RP arm backbone sequence, and both the RP and SNC teams assess the viability of proximity targets, it is ideal for those planning roles to be collocated in the same room during tactical and campaign implementation (CI) shifts to collaborate. Similarly, there is benefit to HIE cross-pollination and discussions with RP and SNC when Ingenuity flights are being planned during a tactical shift. Due to the COVID-19 pandemic, at the beginning of surface operations, there was a restriction that allowed no more than four masked RO team members in the RO uplink planning room at one time. When none of the RO uplink roles are released, there are six RO roles supporting tactical shifts on Monday-Thursday; three RPs, two SNC, and one HIE. If desired by the Science team, on Fridays two additional RPs are staffed to support the ability to implement complicated arm and drive activities in a single weekend plan. During pandemic restrictions, some roles had to participate in uplink planning remotely from another room in the M2020 Operations building or from home. Once pandemic restrictions were lifted, four additional RSVP workstations were set up in the RO Uplink Planning room to support as many as eight RO team members during a single uplink shift.

Starting strategic route planning early has paid dividends

The Strategic Route Planning (SRP) team performs long-range route planning to targets of interest for the Science team using an orbital mesh. Prior to landing, the SRP team was formed from a subset of the RP team, an SRP procedure was generated, reviewed and approved, and two Operational Interface Agreements (OIA) between the Science and SRP teams on proposed strategic routes and strategic route evaluation were approved. The SRP team works closely with the Science team to provide routes and the predicted number of sols it will take to reach each candidate drive target. The responsiveness of the SRP team to Science requests and the accuracy of the predicted number of drives has greatly improved compared to previous missions. The SRP team's work helps guide years-ahead extended mission planning, weeks-ahead selection of future routes toward the next scientific area of interest, and also provides helpful context during weekly traverse reports given to project and science leadership, and daily planning.

RO organization facilitates leadership communication

In previous missions, robotic operations discussions with project and science leadership occurred through many roles, including an RP during a tactical shift, the RP lead, the Tactical Downlink Lead (TDL), the IPE lead, or the EO lead. One advantage of having a separate RO organization is the RO Team Chief can function as the primary point of contact with project and science leadership in management meetings, working groups, and quiet hours. The RO team chief acts as a conduit, collecting and processing critical robotic operations information from the RO role leads and concisely reporting it to leadership. In addition, the RO Team Chief passes down information and requests from the project and science leadership to the RO role leads and team members.

Crossover of pre-landing RO personnel enabled continuity

Prior to launch, phase leads led testing campaigns on engineering model testbeds and on Perseverance during Assembly, Test, and Launch Operations (ATLO). Full RO team staffing began approximately six months prior to landing and the phase leads responsible for testing robotic functionality became members of RO. During the lead-up to landing, there were surface operations Thread Tests (TT) and Operational

Readiness Tests (ORT), as well as strategic product development and planning. Crossover of key RO personnel between pre-landing activities through post-landing surface operations has enabled continuity of strategic development, the release of new capabilities, and retention of key domain knowledge.

Previous RP experience reduced surface operation training

One requirement new to M2020 when staffing RPs was to only seek candidates that were already certified RPs from a previous mission (MSL, MER). That enabled M2020 RPs to be certified very quickly, within a few weeks through an abbreviated RP training process. The shorter training period increased RP time available to support tasks such as creating macros, reusable sequences, procedures, and OIAs prior to landing. This turned out to be very helpful. As expected, schedule slips during earlier phases of the mission led to less time available for operations readiness and slower transitions of some RO staff. During the first year of surface operations, a more comprehensive RP training program was developed, and in February 2022 the first post-landing RP training class began the one-year training program with 6 trainees.

Restructuring RP domain for M2020 was extremely beneficial

While developing the RO organization, several changes were made to the established MSL RP domain. Tactical commanding for the sampling system was removed from the RP domain and incorporated into the domain of the new SNC role. Also, strategic product and capability development for the sampling system was removed from the RP domain and included in the SSO domain. These changes have been extremely beneficial to maintaining the rapid cadence of sampling in surface operations.

ACKNOWLEDGMENTS

The research described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration(80NM0018D0004). The authors would like to thank the Mars 2020 Program for supporting this research.

REFERENCES

[1] V. Verma, F. Hartman, A. Rankin *et al.*, "First 210 solar days of Mars 2020 Perseverance Robotic Operations – Mobility, Robotic Arm, Sampling, and Helicopter," in *2022 IEEE Aerospace Conference (AERO)*. IEEE, 2022, pp. 01–20.

[2] A. Rankin, T. Sesto, P. Hwang, H. Justice, M. Maimone, V. Verma, and E. Graser, "Perseverance Rapid Traverse Campaign," in *IEEE Aerospace Conference*, 2023.

[3] M. Maimone, N. Patel, A. Sabel, A. Holloway, and A. Rankin, "Visual Odometry Thinking While Driving for the Curiosity Mars Rover's Three-Year Test Campaign: Impact of Evolving Constraints on Verification and Validation," in *IEEE Aerospace Conference*, 2022.

[4] A. Rankin, M. Maimone, J. Biesiadecki, N. Patel, D. Levine, and O. Toupet, "Driving curiosity: Mars rover mobility trends during the first seven years," *Journal of Field Robotics*, Jan. 2021.

[5] R. C. Moeller, L. Jandura, K. Rosette *et al.*, "The Sampling and Caching Subsystem (SCS) for the Scientific Exploration of Jezero Crater by the Mars 2020 Perseverance Rover," *Space Science Reviews*, vol. 217, no. 5, 2021. [Online]. Available: https://doi.org/10.1007/s11214-020-00783-7

[6] K. Kaplan, K. Davis, D. Klein, J. Schachter, and B. Wolsieffer, "A Modular Framework for Integrating and Visualizing Telemetry for Mars 2020 Rover Mechanism Operations," in *2022 IEEE Aerospace Conference (AERO)*. IEEE, 2022, pp. 01–15.

[7] J. Wright, F. Hartman, B. Cooper, S. Maxwell, J. Yen, and J. Morrison, "Driving on mars with rsvp," *IEEE Robotics & Automation Magazine*, vol. 13, no. 2, pp. 37–45, 2006.

[8] I. Deliz, A. Connell, C. Joswig, J. J. Marquez, and B. Kanefsky, "COCPIT: Collaborative activity planning software for Mars Perseverance rover," in *2022 IEEE Aerospace Conference (AERO)*, 2022, pp. 01–13.

[9] V. Verma and C. Leger, "SSim: NASA Mars rover robotics flight software simulation," *2019 IEEE Aerospace Conference*, pp. 1–11, 2019.

[10] K. Serebryany, D. Bruening, A. Potapenko, and D. Vyukov, "AddressSanitizer: A fast address sanity checker," in *USENIX ATC 2012*, 2012.

[11] A. Holloway, J. Denison, N. Patel, M. Maimone, and A. Rankin, "Six Years and 184 Tickets: The Vast Scope of the Mars Science Laboratory's Ultimate Flight Software Release," in *IEEE Aerospace Conference*, 2023.

BIOGRAPHY

Vandi Verma *is the Deputy Section Manager for Mobility and Robotics Systems at NASA Jet Propulsion Laboratory, and the Chief Engineer of Robotic Operations for Mars 2020. She holds a Ph.D. in Robotics from Carnegie Mellon University and specializes in autonomous robots, and robotic operations. Robotics capabilities she has worked on are in regular use on the Perseverance, and Curiosity rovers, and in human spaceflight projects. She has worked on the Mars Exploration Rovers, Curiosity rover, Perseverance rover, Ingenuity helicopter Technology Demonstration, Europa Clipper Autonomy Prototype, Europa Lander, and autonomous research robots in the Arctic, Antarctica and Atacama.*

Mark Maimone *is a Robotic Systems Engineer in the Robotic Mobility group at the Jet Propulsion Laboratory. Mark designed and implemented the MER and MSL GESTALT self-driving surface navigation Flight Software; was MSL Deputy Lead Rover Planner, Lead Mobility Rover Planner and Flight Software Lead; automated MER and MSL downlink tools; and is now Mars 2020 Robotic Operations Deputy Team Chief, and member of the Rover Planner and Rover FSW development teams. He holds a Ph.D. in Computer Science from Carnegie Mellon.*

**Evan Graser** is a Systems Engineer in the Engineering Operations for Surface Missions group at JPL. He is a Mobility downlink engineer for Mars 2020 and is the Mobility & Mechanisms subsystem team lead for MSL. He is also a member and lead trainer of the MSL Engineering Operations Systems team. Evan received is MS in Aerospace Engineering from the University of Colorado, Boulder.

**Arturo Rankin** received his Ph.D. in Mechanical Engineering at the University of Florida in 1997 and has worked at the Jet Propulsion Laboratory since then. He is currently a Robotic Systems Engineer in the Robotic Systems Staff group and the Mars 2020 Robotic Operations Team Chief. Other Mars rover roles he has held include Mars 2020 Robotic Operations Deputy Team Chief, MSL Mobility/Mechanisms Lead and FSW Lead, and MER Mobility/Robotic arm downlink analyst.

**Kyle Kaplan** received a B.S. in Aerospace Engineering from the University of Maryland in 2018 and an M.S. in Astronautical Engineering from the University of Southern California in 2021. Kyle joined JPL in 2018, serving as a systems engineer on the Mars Science Laboratory project with focuses in sampling and fault protection. On the Mars 2020 project, he has supported sampling operations development and systems engineering. He is now the lead of the Sampling and Caching operations team.

**Steven Myint** is a Group Lead in the Robotics section of the Jet Propulsion Laboratory. His recent work for Mars 2020 includes developing flight software, applying modern software fuzzing techniques to finding problems in safety-critical flight software, leading the development of SSim (Surface Simulation), the primary simulation tool for Mars 2020 and Mars Sample Return operations, and leading the development of RSVP (Robot Sequencing and Visualization Program). In Mars 2020 Robotic Operations, his roles include Rover Planner, Helicopter Integration Engineer, and Mobility Downlink Engineer.

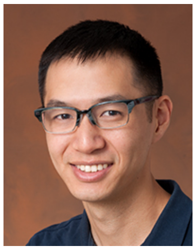

**Justin Huang** joined the Jet Propulsion Laboratory as a Robotics Systems Engineer in 2018. He develops software tools for rover planners, particularly focusing on robotic arms. He also works as Mars 2020 testbed operator and robotic arm downlink engineer, and he is helping to build next-generation software for the Mars Sample Return testbed. Justin developed the collision model for the Mars 2020 rover and the COLDArm project. He received a B.S. in Computer Science from UC San Diego and a Ph.D. in Computer Science from the University of Washington.

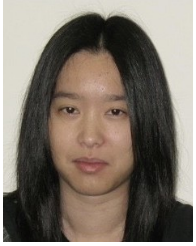

**Amanda Chung** is a Data Visualization Developer in the Science Data Visualization group at the Jet Propulsion Laboratory. She works as a software developer on RSVP (Robot Sequencing and Visualization Program) for the Mars 2020 and Mars Science Laboratory missions, with a focus on the development of mobility tools. She received a B.A. in Asian Humanities from the University of California, Los Angeles and an M.S. in Computer Science from Georgia Institute of Technology.

**Kevin Davis** received a B.S. (2010) and M.S. (2015) in Aerospace Engineering from the University of Maryland. Kevin joined JPL in 2015, working primarily in operations and systems engineering roles for the Mars Science Laboratory and Mars 2020 projects, with a focus on the robotic arm subsystem on both projects. They currently serve as the deputy lead for the Mars 2020 Robotic Operations Downlink Team, and also serve as a Rover Planner for the Mars Science Laboratory and Mars 2020 rovers.

**Andrei Tumbar** is a Robotics Engineer at the Jet Propulsion Laboratory. He works on simulation and operations tools for missions including the Mars 2020 Perseverance Rover, Ingenuity Helicopter, MSL, and COLDArm. In addition to simulation, he designed the compiler infrastructure around the NPM Direct tool as well as built a multi-mission platform for simulating FPrime based flight software. He is currently leading a project to efficiently simulate the Ingenuity helicopter for use in operations.

**Iona Tirona** is a Mechatronics Engineer in the Spacecraft Mechanical Engineering section of the Jet Propulsion Laboratory. She worked on the Mars 2020 Sampling and Caching subsystem from early development through surface operations, and led the design, fabrication, and testing of the Mars 2020 abrading bits. She is currently the deputy lead of the Sampling and Caching operations team. Iona received a B.S. and M.Eng. in Mechanical Engineering from Cornell University.

**Michael Lashore** is a Robotics Engineer currently serving as the Bit Exchange Technical Authority for the Sampling and Caching (SNC) Surface Operations team. Before his stent in surface operations he worked on various sample and caching subsystem (SCS) development efforts, some of which include, the Qualification Model Dirty Testbed campaign for SCS and Corer-level verification and validation (V&V) during the flight assembly integration and test (AI&T) phase. He received his B.S. and M.S in Mechanical Engineering from California Polytechnic State University San Luis Obispo.