

Robotic Operations During Perseverance’s First Extended Mission

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

Ellen Thiel
 Jet Propulsion Laboratory,
 California Institute of Technology
 4800 Oak Grove Dr.
 Pasadena, CA 91109
 ellen.r.thiel@jpl.nasa.gov

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

Nadya Balabanska
 Jet Propulsion Laboratory,
 California Institute of Technology
 4800 Oak Grove Dr.
 Pasadena, CA 91109
 nadya.balabanska@jpl.nasa.gov

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

Noah Rothenberger
 Jet Propulsion Laboratory,
 California Institute of Technology
 4800 Oak Grove Dr.
 Pasadena, CA 91109
 noah.z.rothenberger@jpl.nasa.gov

Ethan Schaler
 Jet Propulsion Laboratory,
 California Institute of Technology
 4800 Oak Grove Dr.
 Pasadena, CA 91109
 ethan.w.schaler@jpl.nasa.gov

Stephen Kuhn
 Jet Propulsion Laboratory,
 California Institute of Technology
 4800 Oak Grove Dr.
 Pasadena, CA 91109
 stephen.kuhn@jpl.nasa.gov

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

Joseph Carsten
 Jet Propulsion Laboratory,
 California Institute of Technology
 4800 Oak Grove Dr.
 Pasadena, CA 91109
 joseph.carsten@jpl.nasa.gov

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

Harel Dor
 Jet Propulsion Laboratory,
 California Institute of Technology
 4800 Oak Grove Dr.
 Pasadena, CA 91109
 harel.dor@jpl.nasa.gov

Abstract—The Mars 2020 Robotic Operations team 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. The rover has driven 30 km and collected 28 samples as of sol 1266 (September 11, 2024). This paper presents the results from the mission after successfully establishing the Three Forks Sample Depot on January 28, 2023, where it deposited 10 of these samples on the surface of Mars. We discuss the climb onto the delta, the upper fan exploration, and exploration of the margin unit as the rover works its way toward the Jezero crater rim. We discuss the autonomy and wheel motor challenges encountered navigating the terrain and strategies for addressing them. We describe the successful rover flight software update to enable Global Localization to enable even longer drive distances and the operational challenges encountered during deployment. We discuss the energy savings from onboard planning to allow more time for driving and sampling. We present updates made to our contact science strategy that have resulted in reducing the number of sols needed to perform an abrasion, and updates to sampling during this extended mission. We discuss recalibrating the Robotic Arm during this period, which is a first for this mission. We also discuss analysis that was performed to ensure that the gas dust removal tool will have sufficient gas and the drill percuss mechanism has sufficient life for the extended life of the mission. Successfully driving through this challenging terrain has enabled the rover to reach a milestone in Mars planetary science exploration: “While previous rovers have driven in fluvial terrains, even small channels draining local relief, this is the first time we have visited the floor of a large Mars river”.



Figure 1. Route of the Perseverance rover from sols 0-1281. Each white dot indicates the stopping point at the end of a Martian sol. The light blue dot on the left shows the rover’s location on sol 1281 (26 September, 2024).

TABLE OF CONTENTS

1. INTRODUCTION2

2. MOBILITY3

3. MOBILITY CHALLENGES3

4. GLOBAL LOCALIZATION8

5. ABRASION SOL PATH OPTIMIZATION9

6. SAMPLING CHALLENGES 10

7. SENSOR RECALIBRATION 13

8. STAFFING CHALLENGES 14

9. INGENUITY OPERATIONS 16

10. SIMPLE PLANNER 17

11. LESSONS LEARNED 18

BIOGRAPHY 20

Table 1. Accumulated mobility odometry and performance as of sol 1265. AVOID_ALL performs fully autonomous drives with Visual Odometry (VO), hazard detection (Terrain Mapping) and avoidance. UNGUARDED drives ignore any onboard map and do not predictively evaluate terrain safety, but may employ any combination of VO and Terrain Mapping; the bottom four rows break out those components from the UNGUARDED summary line above. “Nonimaging” indicates neither mode is active, “VO.Mapping” means that both are active. GUARDED uses VO and hazard detection, but simply stops with a fault rather than avoid a hazard in its path. AVOID_KOZ does no terrain understanding though will avoid any human-specified Keepout zones, but has not been used yet during this mission. This table details the accumulated values at the rightmost edge of Figure 2.

Mode	Distance	Duration	Effective Rate	Total Odometry	Max Distance	Max Distance
	meters	hours	meters/hour	Percent	Sol Number	meters
AVOID_ALL	22332.00	243.18	91.83	75.08	753	331.00
UNGUARDED (overall)	7274.86	87.93	82.74	24.46	1247	78.07
GUARDED	180.40	1.78	101.52	0.61	448	36.49
AVOID_KOZ	0.00	0.00	0.00	0.00	-	-
UNGUARDED_Nonimaging	315.75	2.90	108.75	1.06	47	39.99
UNGUARDED_VO	2405.19	24.31	98.95	8.09	123	73.61
UNGUARDED_Mapping	0.00	0.03	0.00	0.00	-	-
UNGUARDED_VO_Mapping	4553.92	45.67	99.70	15.31	1247	78.07

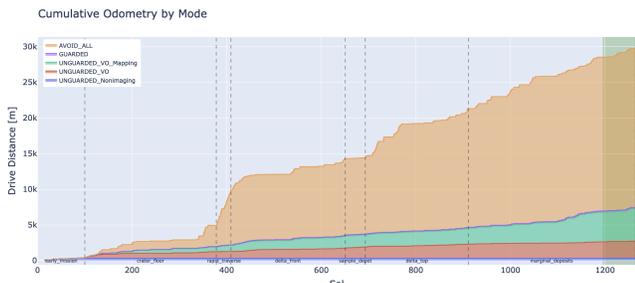


Figure 2. Cumulative plot of Perseverance Mobility Odometry, labelled by drive mode. This is a visualization of the data underlying Table 1.



Figure 3. Selfie taken by the Perseverance rover on Sol 1219 in front of Cheyava Falls.

1. INTRODUCTION

Perseverance has performed extremely well since landing in Mars’ Jezero Crater on 18 February, 2021. It has set new records for autonomous mobility by a planetary rover [1]. In the extended mission we have reduced the time needed to perform abrasions on the surface, further increasing the science return. Observations made by instruments in a clean abrasion return considerably better observations than from a dusty natural surface. The team has incorporated new software technologies onboard and in ground operations, and transitioned Ingenuity helicopter operations to a new phase as part of our extended mission. Commands from the Robotic Operations (RO) team execute nearly every day (see Figure 4).

We have also encountered new challenges along the way. New terrain types led to multiple and ongoing efforts to reconfigure the rover, changes to mission lifetime requirements have led to new constraints on Mars 2020 continued operations, and the steep terrain we are beginning to experience as we begin to explore the Crater Rim region led us to strategically reassess processes and flight and ground software capabilities. We also had to recover from staffing challenges that arose due to layoffs in February 2024.

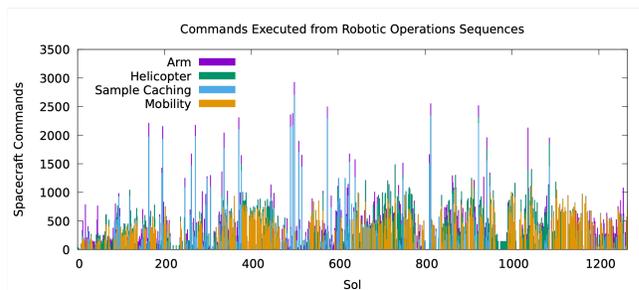


Figure 4. Sol by sol plot of the number of commands executed for each RO activity type. Simple Planner has been active for every RO activity as of Sol 934.

This paper describes how the robotic operations team has addressed these challenges, enabling the rover to navigate challenging terrain to access locations of high science value and make close up science observations with the robotic arm and collect samples. Figure 3 shows an example. The Cheyava Falls observations provided some indications it may have hosted microbial life billions of years ago [2]. Multiple scans of Cheyava Falls by the rover’s SHERLOC (Scanning

Habitable Environments with Raman & Luminescence for Organics & Chemicals) instrument indicate it contains organic compounds. A sample was collected at this location which is on the northern edge of Neretva Vallis, an ancient river valley measuring a quarter-mile (400 meters) wide that was carved by water rushing into Jezero Crater long ago.

2. MOBILITY

Overall our mobility system has performed very well. Table 1 describes the overall performance in each drive mode, and Figure 2 shows the accumulation of odometry in each mode over time. Available drive modes are:

AVOID_ALL The primary Autonomous Driving mode, the rover uses Visual Odometry to maintain position knowledge, and employs Terrain Mapping (stereo vision processing and traversability analysis) to understand the geometric properties of the surrounding terrain, and drive around any perceived hazards.

UNGUARDED Driving without using any predictive on-board knowledge of the terrain geometry.

GUARDED Driving while using onboard terrain knowledge to proactively stop before driving into a predicted terrain hazard, in contrast to AVOID_ALL which can maneuver around predicted hazards.

AVOID_KOZ Driving without using onboard terrain knowledge, but autonomously maneuvering around any human-specified Keep Out Zones.

We have demonstrated unprecedented levels of autonomous driving. A dedicated Vision Compute Element (a second RAD750 computer with Virtex 5 FPGA) on the rover enables high resolution Stereo Vision and Visual Odometry image processing to complete in milliseconds, thus enabling onboard autonomy analysis to occur in parallel with drive motion. Since the effective rate of autonomous driving is comparable to non-autonomous modes, the operations team often chooses to drive autonomously. As a result, **Perseverance is the first Mars Rover to have driven over 90% of its total odometry autonomously** (in an AVOID_ALL, Guarded, or Unguarded Mapping mode) [1].

Traverse Highlights

Perseverance has been climbing its way out of the Jezero Crater Floor area since establishing the Three Forks Sample Depot Cache on Sol 690 [3], [4]. We began driving northwest up onto the Delta fan, experiencing our highest tilt to date of 25.4 degrees at the base of the fan on sol 700 (see Figure 6). We then drove 4.4 km to the Emerald Lake area just west of Belva crater very quickly, averaging 170 m/sol (sigma = 87.9 m/sol) in just 26 drives over 70 sols.

The area west of Emerald Lake was predicted to be more challenging for Autonav, with a greater density of boulders than we had seen before (rocks visible in the HI-RiSE orbital views with at best 25cm resolution). But it ultimately was relatively simple for Autonav to work through.

The route got more challenging for mobility once we entered the Tuxedo Park area about 0.5 km west of Emerald Lake on sol 896. Multiple steering motor faults slowed our progress for several days, as described in Section 3 below.

Our drives along the Margin Unit (toward the Bright Angel target) went well until we entered more complex terrain on sol 1096 (see Figure 9). At that point we found that the

conservative settings of our Autonav system were unable to discover viable paths through that most challenging terrain seen to date, as explained in Section 3. After two months of manually guiding Perseverance through the complex terrain, we made our way into the river valley and once again were able to drive autonomously. We spent a month visiting different sites within the riverbed, reaching the Bright Angel area on sol 1182 after 1.1 km of travel using 9 drives averaging 124.8 m/sol (sigma 58 m/sol).

We departed Bright Angel on sol 1219, heading south toward the beginning of Jezero Crater Rim. In order to reach the crater rim we would have to climb slopes consistently above 20 degrees for the first time. This presented some initial challenges (see the *High Tilt* section below), but we were ultimately successful in achieving full autonomous driving even here, and look forward to continued exploration of high tilt areas in our future extended missions.

3. MOBILITY CHALLENGES

Perseverance encountered multiple mobility challenges since completing the Three Forks sample depot on sol 690 [4], [5]. And although Perseverance's mobility mechanisms were originally required to survive up to 20 km of travel, recent changes to the plans for overall Mars Sample Return [6] and recognition of our having already driven 30 km on Mars have led the Mars 2020 mission to reevaluate our operational strategies in response to these challenges, to help us drive even farther and better support potential rendezvous for sample transfer.

In this section we describe some of the more concerning faults summarized in Figure 5.

Steering Faults

After driving more than 20 km during its first two and a half years of operation, Perseverance's steering motors experienced their first faults on sols 896, 897 and 899. These faults occurred on different motors and had multiple signatures, which led us to conclude that they were occurring due to terrain interactions, not internal issues. Perseverance often commands large cumulative steering changes during Turn in Place motions (around 90 cumulative degrees per corner wheel on average), and each of these faults occurred while steering into or out of a Turn in Place configuration.

The first two events were STALL faults (indicated as MOTOR faults in Figure 5), during which the wheels were pushing up against rocks and were not able to achieve their commanded turn amount. And the third was a CMAX_SLOW fault, indicating that the current being used to drive an individual wheel had exceeded a threshold of 2.5 Amps for more than 2 seconds. Visual inspection of the terrain showed no motion of adjacent rocks during the stall events, but there was some resettling motion during the CMAX_SLOW fault (see Figure 7). Discussions with the science team led to the conclusion that during the first two faults we had encountered rocks behaving like "icebergs"; a small fraction of the rock being visible above ground, but most of the rock being firmly held in place underneath the surface. That led us to change our short term driving strategy to minimize use of Turns in Place and only command shorter distances while in that terrain, review and update motor fault protection parameters (described in part in the *Hard-brakes* and *Delayed Brake Engagement* sections below), and also to create long-term goals to modify the flight software (FSW) to better enable

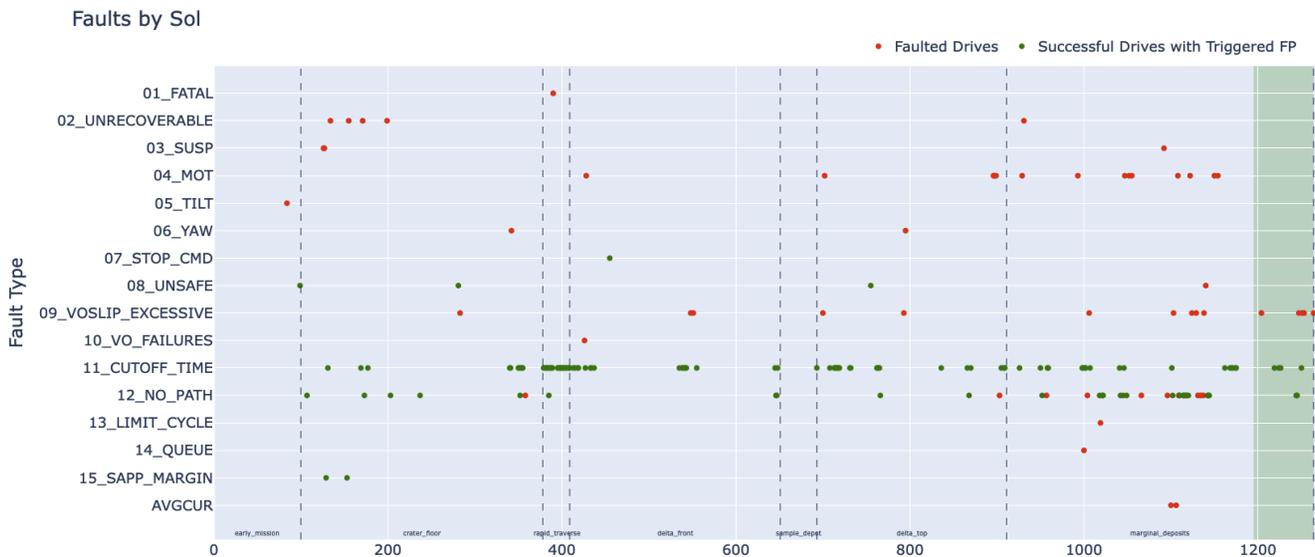


Figure 5. Sol by sol plot of Perseverance Mobility Faults. Here “fault” means any asynchronous interruption of a mobility command. Some faults are nominal, e.g., every Autonav drive is expected to end with a CUTOFF_TIME fault because that means it has used all available time to drive as far as possible, until driving was cut off by the time limit. Similarly, most NO_PATH faults are considered nominal since the rover successfully avoided driving over a previously unknown obstacle.

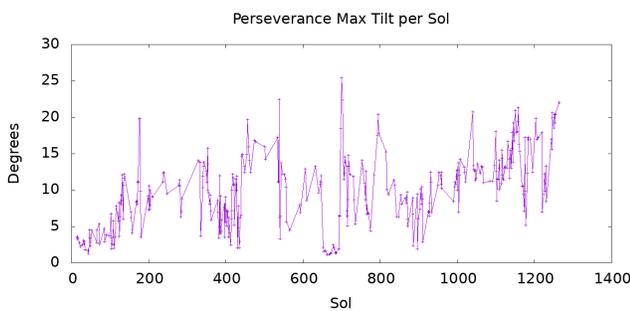


Figure 6. Maximum tilt experienced by Perseverance on each drive sol.

future driving using fewer or no Turns in Place.

Hard-brakes

The increased frequency of stalls due to terrain revealed an undesired stall response in the mobility brake system. The stalls on Sols 897, 929, 1047, and 1052 all resulted in hard-brake events, in which the brakes are closed before the motor speed has decreased below a safe threshold. This can cause significant damage to the brakes and the motor via detrition and debris generation. Because of this, the mobility motors are limited to five hard-brake events each over the mission lifetime. After the Sol 1052 hard-brake, the M2020 Steer_RR motor had seen three of the five allowed hard-brakes in quick succession over the course of only 156 sols. This frequency of hard-brakes has not been observed on MSL, which has only seen three total hard-brakes on mobility motors over the course of 2531 sols, with only two of those being on the same motor. The Perseverance Rover needs to maintain its ability to traverse large distances well into the future, as it may play a pivotal role in a future Mars Sample Return mission. The increased frequency and number of hard-brakes endangered

Perseverance’s ability to do that. Because of this risk, driving was halted while the underlying cause was investigated and a solution was worked to prevent future hard-brakes.

The underlying cause was determined to be the grouped brake configuration in the M2020 Rover Motor Control Avionics box, in which a brake driver channel controls two brakes at once. In other words, when one brake is opened or closed, its grouped brake will also open or close, and vice versa. This configuration was a change from MSL and was necessary because the number of brake driver channels was halved from MSL to M2020, but the number of actual brakes only decreased from 54 to 41. Therefore, in order to operate all brakes, some motor brakes had to be “grouped” to share a driver channel with another motor, and all of the mobility motors were grouped with other mobility motors. Unfortunately, this grouping created a new issue for the motor control stall response behavior. When a stall occurred, the stalled motor’s brake could not close until its grouped non-stalled motor had completely stopped. While the non-stalled motor was stopping, the stalled motor was only being held in place via dynamic braking, which is a passive method that uses back EMF to slow the motor down. In the case of many steer stalls, the back EMF would not be enough to hold the stalled motor in place, so the motor would start backdriving to unwind the tension built up during the stall. When the brakes finally closed, the stalled motor would still be backdriving at a speed above the safe threshold, and, therefore, it would experience a hard-brake event. If the stalled motor was not grouped, its brake would have been able to close immediately, preventing any backdriving or hard-brakes.

The problem was fixed by changing motor control FSW parameters to prevent the FSW from using dynamic braking on stalled motors. Nine total parameters were changed with the primary purpose of causing a CMAX_SLOW fault to trip before a stall fault is tripped. This prevents hard-brakes because CMAX_SLOW faults use the ramp-down stop mode instead of dynamic braking. In ramp-down, the motor is actively

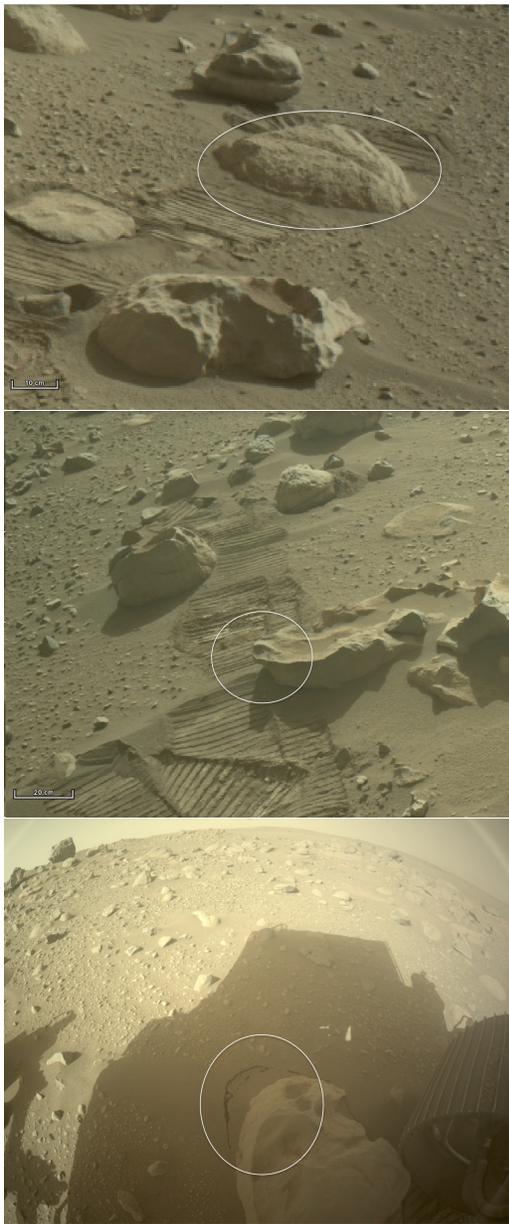


Figure 7. Rocks that caused steering faults on sol 896 (top), 897 (middle), and 899 (bottom). Some soil settling is visible only in the sol 899 image, not in the earlier sols' images.

controlled while it slows down to a stop, or while its held in place, until the brakes can be closed. This active control prevents the motor from backdriving and, thus, avoids hard-brake events. The main parameters modified were relating to current limits and persistence, decreasing the persistence needed for a CMAX_SLOW fault to be below the stall fault persistence, and increasing the CMAX_SLOW fault current limit so that it would effectively replace stall faults. These parameter updates were put in place on Sols 1063 and 1065, and driving resumed on Sol 1066. After these parameter updates, multiple instances of CMAX_SLOW faults were observed that would have triggered a stall and likely resulted in a hard-brake, if not for the parameter updates, confirming that the parameter updates were successful in preventing hard-brake events.

Delayed Brake Engagement

In high-rate data from a CMAX_SLOW fault on Sol 1108, the STEER_LR motor encoder telemetry indicated an additional 40 encoder counts (approximately 100 degrees) of motor rotation after the timer for brake closure had expired and motor control was released. Further investigation of other high-rate data indicated that similar instances of unexpected rotation after brake closure were seen in telemetry from 6 additional drives prior to the one on Sol 1108.

Investigation of brake closure timing on the Vehicle System Test Bed (VSTB) rover confirmed that mobility brake engagement time on M2020 could exceed the limit set for MSL, due to a combination of lower brake engagement current on the M2020-build actuator and the grouped braking configuration on M2020. Once brakes are commanded closed, motor control FSW waits for the expiration of a timer set by the parameter “tclose” which indicates the time to wait for brake engagement, after which control of the motor is released. In repeated testing, it was found that an individual motor’s brake would close within the time limit set with high probability, but once in the grouped braking configuration, the back-EMF of a brake drop-in can extend the time to re-engage for the paired brake due to increasing current in the series path. This phenomenon led to one or both brakes in a grouped pair unintentionally remaining in a disengaged state for a short time after motor control is stopped, which in the presence of external loads allowed for the unexpected back-drive seen in flight telemetry.

The fix was to increase the parameter for the “tclose” timer to wait for brake engagement before releasing control from its original value of 110 msec to 250 msec, a value which was determined from a dataset of measurements of the maximum time for brake engagement gathered in the VSTB. The parameter change was deployed in flight on Sol 1213.

Complex Terrain

Starting on sol 1096, we entered terrain that was too complex for the Autonav parameter settings that had been so successful during the first three years of driving. As a result, several sol’s drives ended with a NO_PATH fault in which Autonav was unable to find a way to finish the requested drive. Some earlier terrains had also been strategically assessed as being potentially challenging for Autonav performance, but this was the first to actually cause problems (see Figures 8 and 9).

Once it became clear that Autonav would be unable to find a safe path over a significant distance, Mars 2020 Project Management created a high priority task for the Robotic Operations team: develop strategics to tune and/or improve Autonav performance, and look for alternatives to the strategic route that would enable us to drive longer distances once again. However there was challenging terrain all around us, including tall sand ripples in the river valley to the north for the next kilometer. So Perseverance spent the next two months driving only as far as human Rover Planners could safely see each sol. That resulted in much shorter drives between sols 1096 and 1159, when 35 drives averaged only 42.89 meters/sol with $\sigma = 22.63$ meters/sol (compared to all 143 Autonav drives which averaged 129.59 meters/sol with $\sigma = 81.01$ meters/sol). Figure 8 illustrates the shorter distances via the more densely packed end-of-sol white dots spanning that part of the drive. We also found a way to descend into the flat river valley sooner than previously planned, once we had driven past the deepest ripples. Perseverance drove 1.7 km in this human-planned mode for over two months.



Figure 8. Complex Terrain prevented Perseverance from reliably employing Autonav from sol 1096 to sol 1159 (20 March 2024 - 24 May 2024). The blue line indicates the pre-planned strategic route that we had hoped would be viable for Autonav, the white dots indicate where the rover stopped at the end of each drive sol. Perseverance ultimately entered a flat and relatively clear part of the ripple-laden river valley on Sol 1159, the first time a Mars Rover has explored the floor of a large Mars riverbed.

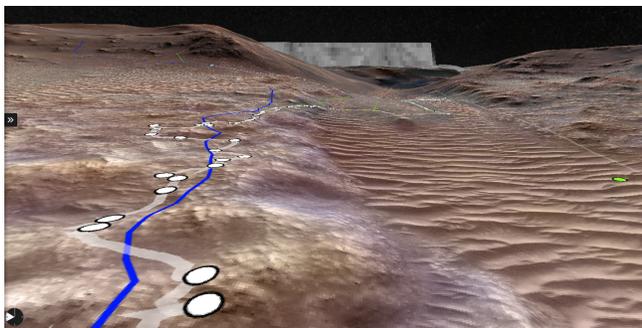


Figure 9. Perspective view of the challenging terrain in Figure 8, starting at the Sol 1096 location in the lower left of the image.

At the same time we formed a new strategic team comprised of Rover Planners, Mobility Downlink analysts, and current and former team members who originally developed the Enhanced Navigation (ENav) Flight Software [7]: the *Autonav Roundtable*. This team pursued multiple strategies, including: collecting stereo imagery and terrain assessments from the newly challenging terrain, characterizing all current and previous Autonav-related mobility faults, reviewing and identifying FSW parameters needing adjustment, restoring a Monte Carlo Autonav Terrain simulation framework, and re-assessing the Autonav FSW design and implementation.

After careful parameter and design review, we ultimately raised the planning tilt limit from 16 to 20 degrees and updated our predictive slip model accordingly, reduced the

clearance required under the belly pan from 25 cm to 20 cm after confirming the extra 5 cm of margin was not required, increased the average current limit across all drive motors from 1 to 2 Amps, and updated our slip calculations to include distance driven while rotating during Turns in Place.

We also proposed, approved and implemented nine FSW changes to improve Autonav performance and reduce the likelihood of future mobility faults. Those changes are expected to be incorporated onto Perseverance when the version S8.1 flight software release is uplinked sometime after June 2025.

The Autonav Roundtable, founded in April 2024, continues to meet, monitor and recommend changes as of October 2024.

High Slopes

During the first three years of operation, Perseverance was driven on mostly-flat terrain. In light of this, parameter for the maximum Tilt angle that Autonav would allow was set conservatively to 16 degrees. Anticipating the need to drive on higher slopes in and beyond the area of Complex Terrain mentioned above, this was increased to 20 degrees on sol 1124 (18 April 2024).

Figure 10 shows the beginning of our planned climb toward the Crater Rim region of Jezero Crater. The first uphill drive took place on Sol 1244 (19 August 2024); this and the next drive both stopped short with NO_PATH mobility faults due primarily to the high tilt of this region. So the Autonav Tilt limit was raised to 24 degrees on sol 1246. This new value and other related parameters were set only after careful

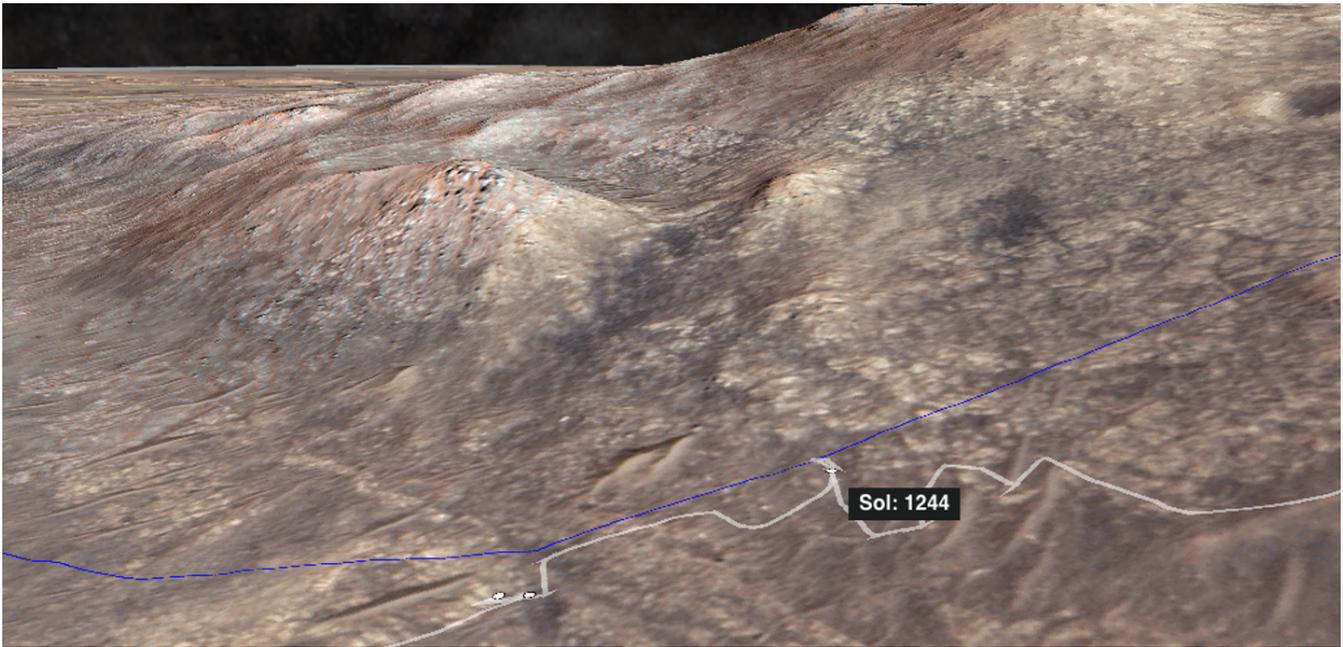


Figure 10. Sol 1244 was the first Autonav drive planned to climb a 20 degree continuous slope.

assessment of the parameter change impact on performance, running thousands of terrain simulations in a Monte Carlo framework that had been assembled during ENav FSW development. Those simulations showed continued safe and successful behavior, justifying the parameter change for use in flight.

We made these changes to improve performance while not increasing risk to the vehicle. A separate parameter defines the maximum tilt the vehicle is allowed to reach during drives, and remained set at 25 degrees as it had been during earlier drives. That limit is part of a reactive fault protection process checked at 8 Hz against the current tilt limit as measured by the accelerometers in the Inertial Measurement Unit (IMU).

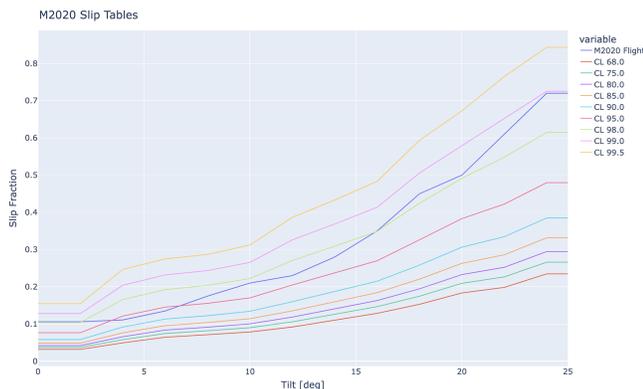


Figure 11. Onboard Slip Tables. All tables are derived from a combined dataset of MSL slip data and the first 1250 sols of M2020 slip data with different Confidence Levels applied to distributions of each 2-degree tilt bin. The default M2020 Flight table has been in use since landing.

The parameter changes were successful. No more NO_PATH faults occurred after sol 1246, though some drives did end early due to higher than expected slip on the steep terrain.

Between sols 1251 and 1264, 4 out of 5 drives faulted early due to excessive slip. 3 of the 4 faults occurred during the Autonav drive portion. The sustained high slip was due to a combination of the high tilt slopes, the drive direction being almost entirely pitch-up (up-slope as opposed to cross-slope), and loose, sandy terrain.

The abundance of higher degree slopes in the Crater Rim region is also expected to result in higher slip values than what the mission experienced over the first 3 years. ENav models expected rover slip using a “Slip Table” - a collection of parameters maintained in FSW which describe how much rover slip to expect at different rover tilts. This slip model determines the size of the “wheelboxes” that cover all possible ending wheel positions after driving a short distance (nominally 1 meter). Wheelbox size in turn impacts a number of predictive safety assessments made during autonomous driving including maximum wheel-drop and rover bellypan clearance, two of the most common reasons the rover has triggered NO_PATH faults during the mission. A single slip table, created using a combined dataset of MSL flight and M2020 testbed slip data, has been in use for the entire mission so far.

In anticipation of the desire to drive on high-tilt, high-slope terrain; a new process creating *Local Slip Tables* has been implemented to select different slip tables when driving that more accurately model rover slip in specific regions of terrain. The Strategic Route Planning team identifies stretches of terrain along the Crater Rim where the rover is expected to traverse homogeneous terrain. After entering such a region, the slip data collected on that terrain is used to choose from a larger set of slip tables now maintained onboard (Figure 11) which model different expected amounts of slip. In regions where slip is consistently lower than what is modeled by Autonav’s default slip table, using a lower slip table improves ENav’s slip modeling and reduces the likelihood of nuisance NO_PATH faults. In regions where slip is consistently greater than what is modeled by the default table, using a higher slip table allows the ops team to set correspondingly higher

slip limits. Prior to the use of Local Slip Tables, Autonav’s reactive slip limit was 60%. The new process will allow Autonav to drive with limits up to 80%. This should improve Autonav’s efficiency, particularly on steep but benign slopes which have high slip but low obstacle density. Local Slip Tables will begin in-flight use in October 2024.



Figure 12. Hi-RISE map with GBLC annotations. Green represents the approximate traversable region defined by Strategic Route Planning corridors. Yellow is the area where GBLC can be run. Red is the boundary of the map files.

4. GLOBAL LOCALIZATION

Maintaining position knowledge over long drives is critical to ensure the rover actually achieves the intended result of the drive. Long autonomous drives are planned using orbital imagery: humans draw a series of waypoints along with “Keep-in” and “Keep-out” zones to indicate safe and potentially hazardous areas, and the the rover autonomously finds its way while obeying those constraints. For safety’s sake the rover FSW maintains a model of position uncertainty while driving and uses it to grow/shrink the bad/safe areas in its world model to compensate. Eventually that uncertainty growth will mark all possible paths as too dangerous, and will end the autonomous drive early.

Our team has developed a new capability for autonomous *Global Localization* (GBLC), to allow the rover to re-localize itself in orbital Hi-RISE maps using NavCam panoramas [8], [9]. Global Localization not only updates the rover’s position knowledge to better align with the orbital map, it also allows us to reset our current uncertainty back to the scale of the orbital map (typically 1-2 meters). The ability to autonomously shrink uncertainty is an enabling capability for very-long-range drives (over 1km), limited only by orbital map size. [9] describes the benefits of resetting uncertainty during a long traverse. The image processing is performed on a separate Heli Base Station (HBS) processor within the rover, which includes a 2 GHz Snapdragon 801 CPU that was added to provide an interface between the Ingenuity helicopter and the rover’s FSW. We transmit the current position, orbital maps and rover-acquired images to the HBS, and it returns an offset and residual uncertainty information to the rover’s FSW.

Rover FSW was updated to include Global Localization commands and capabilities in its S8.0.0.3 Mobility component update on sol 1151 (15 May 2024).

This capability marks the first use of a commercial co-processor to supplement rover FSW processing on a slower

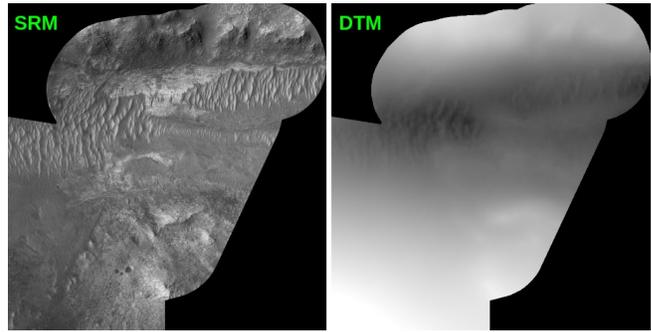


Figure 13. An example Orbital Appearance Map (Surface Reflectance Map (SRM)) and Orbital Elevation Map (Digital Terrain Model (DTM)) used during Global Localization.

(133 MHz), rad-hardened processor (RAD-750). It will be demonstrated in three phases [9]. The first phase demonstrated the ability to “sandbox” our software separately from the Ingenuity interface code, and completed on sol 859 (21 July 2023). The second phase exercised all the HBS software using previously acquired data, and showed that the flight results matched expectations on sol 914 (15 September 2023). The third phase will demonstrate the complete capability, taking images of the nearby terrain, performing Global Localization and making the results available to the rover FSW.

The third phase consists of a multi-use flight activity which can be run after any Perseverance drive. Following a drive’s completion, a full 360 degree panorama will be taken using the rover NavCams. The images from this panorama will be transferred to the HBS along with Global Localization files kept onboard the rover’s flight computer. The image processing and rover localization is then performed on the HBS, and its results are sent back to the main rover FSW which updates the rover’s position knowledge and uncertainty if the localization between the Hi-RISE map and panorama succeeded. This phase is scheduled to begin in late 2024.

There was a complication with the development of the full closed-loop GBLC capability. Our rover engineering model, the Vehicle System Testbed (VSTB) was not available for testing during the initial development of the S8.0.0.3 Mobility FSW. During later system testing we discovered two issues that would complicate our use of GBLC in flight operations. One was that the rover FSW failed to update an onboard frame transformation that represents where the onboard World Map Model exists in the current Rover Frame, and another was that a command used for checksum validation of files going between the HBS and main CPU always ended up modifying the computed checksum stored in those files. Happily, we were able to develop FSW patches [10] to address those problems and get them approved for use in a few days. It was too late to incorporate those fixes into the S8.0.0.3 release, but they were added to S8.1 release scheduled for deployment in mid-2025.

Global Localization requires maintaining a set of onboard orbital maps covering the mission’s traverse path. These maps are created based on “traversal corridors”, which are created and maintained by the Strategic Route Planning Team and denote the regions where the rover is expected or able to drive. Figure 12 shows the section of the Hi-RISE map used to create the example Global Localization maps shown in Figure 13 that were created for use onboard the vehicle.

5. ABRASION SOL PATH OPTIMIZATION

Target Evaluation

Analyzing a target for abrasion suitability is by far the most challenging type of arm target analysis performed by the robotic operations team. First, safety of the drill placement used for the abrasion must be completed. Adequate clearance between the arm hardware and the terrain must be ensured, and surface topography constraints for the terrain under the drills stabilizers and abrading bit verified. The abrading operation requires preloading the drill with 350 N of force. An analysis must be done to ensure that the robotic arm is capable of safely applying that load in the pose required to reach the target. In addition to evaluation of the safety of the abrading operation itself, there are a variety of additional arm activities to be evaluated. A Facility Contact Sensor (FCS) touch may be necessary to refine target location knowledge before abrading. The Gas Dust Removal Tool (gDRT) is used to clean the abraded patch. This activity consists of four different tool placements, each angled differently with respect to the surface and with a compressed gas puff at each placement. And finally, WATSON reconnaissance images are taken which are used to pick proximity science targets within the abraded area on a subsequent sol. Each of these six placements (FCS touch, four gDRT placements, and WATSON imaging), must have adequate terrain clearance for a potential target to be viable for abrasion.

In addition to ensuring that all of the abrasion activities can safely be performed on a candidate target, a predictive clearance assessment must be completed for the PIXL and SHERLOC instruments. These instruments require very small standoffs from the surface (2.55 cm and 4.8 cm respectively) generally making maintaining adequate terrain clearance challenging. Placing on abraded surfaces only makes things more challenging. The abrasion lowers the surface that the instruments need to be placed relative to, making clearances to everything outside the abraded patch worse. In addition, the abrasion creates tailings piles around the edge of the abrasion, further reducing instrument clearance. A ground-in-the-loop cycle is always taken between abrading and PIXL and SHERLOC placements. This is due to the fact that it is not possible to predict exactly how the abrasion activity will modify the topography down to the very small clearance margins available. In addition, there are occasionally catastrophic rock fractures or large shifts in terrain that generate topographies that are well outside the bound of a normal abrasion and do not allow PIXL or SHERLOC placement. See Figure 14 for an example.



Figure 14. Before and after 40 cm standoff WATSON images of the Chiniak abrasion attempted on sol 562. Note the significant rock fracturing and changes in surface topography due to the abrasion operation. Credit NASA/JPL/Malin Space Science Systems.

While PIXL or SHERLOC are never placed without knowing

the final topography, it is necessary to ensure there is a reasonable likelihood of having sufficient clearance before performing an abrasion in the first place. There is no point in spending the time and resources to create an abrasion we cannot collect proximity science data on it. For this predictive analysis, it is first necessary to determine the maximum depth abrasion that would still allow PIXL and SHERLOC placements. This analysis is fairly crude, but provides a reasonable guiding approximation. This depth is compared to the abrasion depth necessary to remove all surface topography within the abraded region, thus leaving a flat surface behind. Ideally this abrade depth is less than the instrument limited abrade depth (meaning PIXL and SHERLOC will likely be able to be placed). If not, then either risk will have to be accepted (either shallowing out the abrasion, or potentially not being able to perform instrument placements after the abrasion), or an alternative target identified.

All of the evaluations described above need to yield successful results in order for a candidate target to be viable for abrasion. Performing this significant number of evaluations has historically been time consuming and error prone. If any one evaluation fails, that rules out that target and operators have to start over on a new target. Therefore there are many cases when much of this process is repeated multiple times before a viable target is found. For the initial phases of the mission, this was dealt with by dedicating an entire planning cycle toward doing this target analysis. The rover would arrive at a new workspace, and rather than commanding the abrasion activities in the following planning cycle, the abrasion would be deferred until the subsequent planning cycle allowing time for suitable targets to be found. However, rather than doing nothing during that first planning cycle, WATSON images of possible abrasion targets were gathered. Designating targets using a WATSON images reduces the lateral arm placement uncertainty. This helps make the predictive PIXL and SHERLOC evaluations more likely to pass as it increases the chances that the abrasion will end up exactly where desired, and thus less terrain around the target needs to be included in the analysis. This resulted in each abrasion taking three planning cycles. The first was devoted to WATSON reconnaissance imaging of potential abrasion targets and target evaluation. In the second, abrasion, gDRT, and WATSON reconnaissance imaging of abraded patch were performed. During the third, the abraded patch proximity science was performed including PIXL, SHERLOC, and close approach WATSON imaging.

Optimization

To streamline operations and improve efficiency, a large effort was undertaken to automate the vast majority of the abrasion target analysis process, culminating in a new tool called Target Wizard. Using this tool, operators specify a few details about the candidate target, and then at the click of a button most of the necessary evaluations are done. This reduced the time to complete the evaluations from hours to seconds. This development meant that a planning cycle dedicated to target analysis was no longer necessary. However, immediately jumping into the abrasion also meant that WATSON imaging of potential abrasion targets would no longer be available. This would make identification of small pebbles in the vicinity of the abrasion patch more difficult and would necessitate using a much larger lateral placement uncertainty (generally going from about 11 mm to about 29 mm for the predictive instrument placement evaluations). This larger uncertainty would make it more difficult to find suitable targets. It would not make much difference for large flat rocks, but would have significant impact for smaller, rougher

targets. A small change to risk posture was introduced to increase the likelihood of locating suitable abrasion targets without having WATSON images in hand. Rather than using the standard 3-sigma lateral uncertainty for the accuracy of the drill placement (99.7% confidence) the uncertainty metric was reduced to 1.5-sigma (87% confidence). This was done only for performance related evaluations where the consequence is nuisance fault, increased risk of not being able to place PIXL or SHERLOC, or increased risk of having a shallow abrasion that does not abrade away all surface topography. For hardware safety evaluations (like terrain clearance) the standard 3-sigma uncertainty was retained. This makes target evaluation significantly more likely to pass (lateral uncertainty of about 17 mm for the predictive instrument placement evaluations), but still not quite as good as if WATSON imaging was available.

With these updates, the initial planning cycle dedicated to target evaluation was removed from the abrasion baseline, reducing the activity from three planning cycles to two. The option to fall back to the three cycle version was preserved for particularly challenging workspaces where the reduced lateral uncertainty provided by WATSON imaging is required to find a viable target. The optimized abrasion sol path was first exercised for the sol 1151 abrasion of Old Faithful Geyser. The abrasion was conducted on a planning cycle in which a mobility software update was being performed and thus no driving could be planned. This happened to occur on the planning cycle before a three-sol weekend plan in which arm activities and driving can both be planned. This presented the two planning cycles necessary for abrasion activities. This was an opportunistic abrasion that would not have occurred without the sol path optimization. Under the old sol path, one additional planning cycle would have been needed which would have necessitated sacrificing drive sols in the middle of a drive campaign. Between sol 1151 and sol 1256, six abrasions were completed using the optimized sol path. There were no cases where it was necessary to fall back to capturing WATSON images for abrasion target designation. There were no nuisance faults during any of these abrasions, and PIXL and SHERLOC were successfully placed on all abraded patches.

6. SAMPLING CHALLENGES

The Sampling and Caching Subsystem (SCS) has successfully acquired rock core, regolith, and atmospheric samples and prepared abraded rock surfaces for 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 15. The Corer's two stabilizers are preloaded onto rocks during abrading and coring, while the Corer feed translates the bit of a rotary-percussive drill into contact with the rock. The percussion mechanism is critical to the drilling operation, using percussive energy to chisel and chip away at the rock to produce either abraded patches or core samples. After abrading, the gDRT is used to remove dust and rock powder from the abraded patch to clean the surface for instrument observations. As of sol 1266, the SCS has created 32 abraded patches, puffed with the gDRT 176 times, and sealed 28 sample tubes: 22 rock core samples, 2 regolith samples, 1 atmospheric sample, and 3 Witness tube samples which document the SCS's exposure to the Martian elements over time (see Figure 16 for sample images) [11].

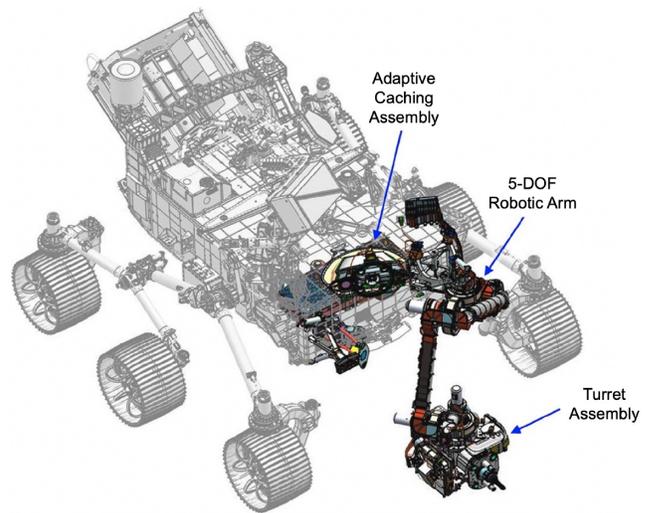


Figure 15. Mars 2020 Rover with SCS components highlighted.

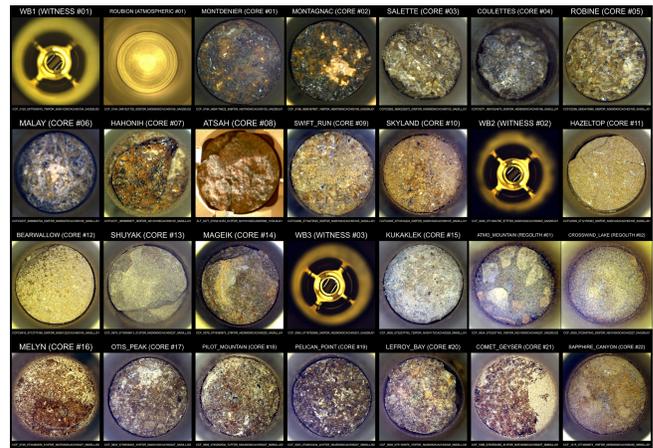


Figure 16. Cachecam or Mastcam-Z images of all 28 sealed samples in order of acquisition. Credit: J. Maki.

The SCS is operated by the Sampling and Caching (SNC) team, who since landing have evaluated 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 [12]. SNC works together with the Strategic Sampling Operations (SSO) team who test, validate, and deliver the sampling products used by the SNC team to efficiently operate the SCS. These teams have also faced and overcome various challenges in successfully operating the SCS. Early into the mission, the team noticed discrepancies in the gDRT's gas budget that, if true, would threaten the team's ability to complete certain science objectives. Later, the team encountered a hardware failure in the percussion mechanism during ground testing that led to a more detailed understanding of percussion life in flight operations. Tackling these problems has led to improved cognizance over the SCS and increased confidence in the SCS's ability to complete the remaining science mission.

gDRT Gas Budget Anomaly

The gDRT is responsible for clearing away dust from Martian surfaces prior to proximity science operations. The gDRT is employed in two scenarios: on natural, un-abraded surface patches, where a single puff of gas directed normal to the surface is performed, and after an abrasion (e.g., after the Eremita Mesa abrasion shown in Figure 18), where four puffs are typically performed (three at an angular offset from the surface normal, followed by one puff directed normal to the surface). In the case of abraded surfaces, the gDRT clears away dust after the weathered outer layer of the rock is removed through abrasion, exposing the unaltered, underlying rock. For natural surfaces, the gDRT is used to blow away loose dust. In both cases, this dust removal is critical for ensuring the PIXL and SHERLOC instruments can perform high-resolution imaging and analysis on clear surfaces. The gDRT relies on a limited supply of nitrogen gas stored in a pressurized tank, which was filled with 172.8g of nitrogen before launch. Each puff is a short opening of the gDRT run valve (see schematic 17). An entire abrasion activity is expected to use approximately 0.564g of gas, while a natural surface puff typically uses about 0.2535g of gas.



Figure 18. Eremita Mesa abrasion patch from sol 1257, following gDRT puffs. The gDRT was used to clear dust from the exposed rock surface after the abrasion process, enabling proximity science instruments PIXL and SHERLOC to analyze the freshly revealed, unweathered rock beneath the Martian surface layer.

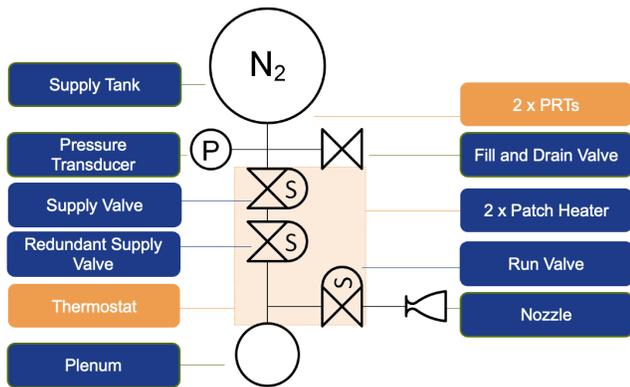


Figure 17. gDRT schematic showing the supply tank which holds the nitrogen gas supply, the plenum which is filled with a small amount of gas before every operation, the supply valves which control gas exchange between the supply tank and the plenum, and the run valve which controls gas exchange out of the plenum and onto a surface target.

On sol 254, an unexpectedly high gas usage estimate was noted during an abrasion activity, prompting an investigation into the gDRT gas budget and the accuracy of its tracking methods. This raised concerns about the longevity of the nitrogen supply in the gDRT system, which is critical to ensuring that sufficient gas is available to complete the mission's science objectives.

Three methods were originally used to estimate the gas usage: Method 1, which relies on pressure and temperature sensor readings to estimate the gas density and the amount of gas remaining via instantaneous readings - using the pressure measurement taken right before and after the puffs to calculate the mass difference in the supply tank; Method 2, which utilizes the known mass of nitrogen at launch and combines it with ground models to estimate gas depletion based on valve actuations and expected leakage rates; and Method 3, which expands on Method 2 by incorporating pressure data to refine gas usage calculations.

Of these, Method 1 showed the largest discrepancies in gas consumption estimates (see Figure 19), leading to suspected

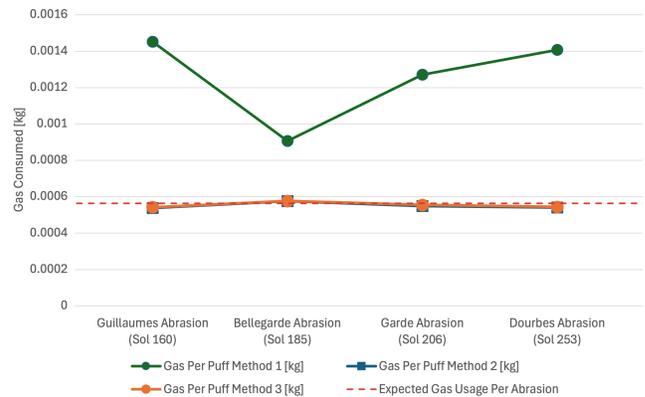


Figure 19. Initial instances of the gDRT gas budget anomaly, highlighting discrepancies in gas consumption estimates. Method 1 shows significant higher amounts of gas consumed per abrasion compared to the expected amount of 0.564g per abrasion. These discrepancies led us to investigate the accuracy of gDRT gas tracking methods.

errors in the pressure transducer readings, as Method 1 heavily depends on these readings to calculate gas density. The investigation also considered other potential sources of error, including issues with temperature measurement and valve timing, though these other potential sources were not found to have a significant impact on the discrepancies.

Closer analysis of gDRT data indicated that most pressure transducer measurement variability occurred during gDRT puffs. To increase the availability of reliable data, regular weekly gDRT reads were initiated, focusing on periods without active gas consumption. During these periods, gas density is expected to remain relatively constant since valve leakage is assumed to be negligible over short durations. Therefore, the analysis of gas density during read periods became the focal point of the investigation.

The analysis of these read data points revealed that while they

were generally more consistent compared to data acquired during gDRT puff activities, there was still a noticeable trend of increasing error at lower temperatures. The pressure transducer's readings at low temperatures were causing the gas density estimates to be artificially reduced, resulting in the overestimation of gas consumption seen in Method 1. This finding suggested that pressure transducer errors at low temperatures were a significant contributor to the discrepancies in the gas usage estimates.

In addition to pressure transducer errors, thermal lag between the gas itself and the platinum resistance thermometer (PRT) readings was investigated as another potential error source. The PRTs, located on the tank boss (see Figure 20), measure the temperature of the tank boss rather than the internal gas temperature. Particularly during rapid environment changes, as have been seen during gDRT valve operations, the PRT readings were found to lag behind actual gas temperature changes in the supply tank. To account for this, the rate of temperature change (dT/dt) was incorporated into the gDRT pressure modeling. The inclusion showed statistical significance, confirming that temperature lag was indeed impacting the accuracy of gas consumption estimates.

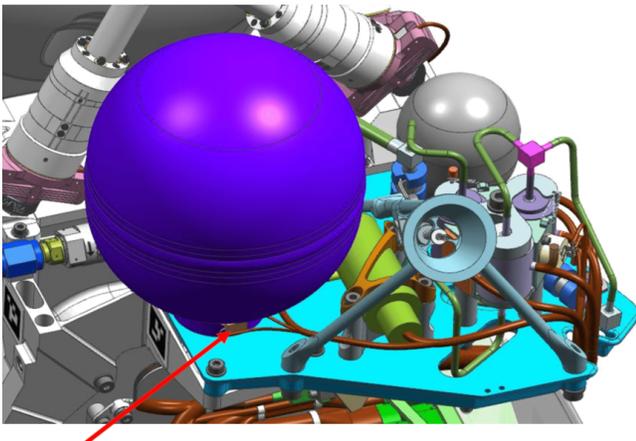


Figure 20. CAD model of the gDRT system with an arrow indicating the location of the PRTs on the tank boss. The PRTs provide critical temperature data that is used in conjunction with pressure transducer readings to estimate gas consumption. Their placement on the tank boss rather than measuring internal gas temperature was a key consideration in investigating the gas budget anomaly.

As a result, the decision was made to filter out data points captured with high absolute dT/dt values or low temperatures. These findings led to changes in the tracking of gas consumption. A new method was introduced that tracks gas usage based on “latest density” calculations. This method recalculates the gas density between gDRT operations using all valid data points collected between gDRT puffs, excluding error-prone data points captured at very low temperatures or during high absolute rates of temperature change (dT/dt). Regular gDRT reads provided a stable basis for these new calculations, keeping gas density relatively constant, offering a more accurate estimate of gas consumption over time.

This approach significantly improves the accuracy of the gDRT gas budget, aligning more closely with expected gas usage trends while overall exhibiting reduced cumulative error in gas consumption estimates (see Figure 21). This refined method confirmed that the gDRT was functioning as intended, with no evidence of excessive gas leakage or

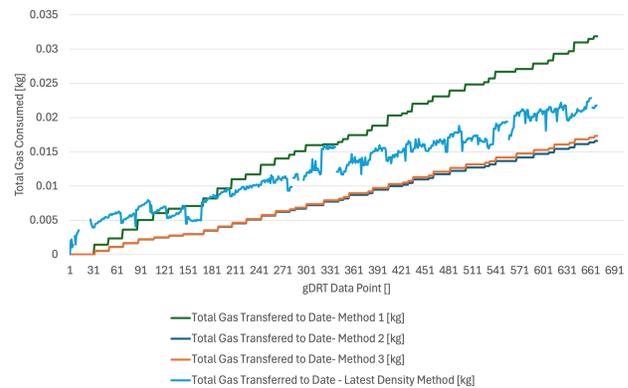


Figure 21. Comparison of gDRT gas consumption estimates using the new “latest density” method versus previous methods (Methods 1 - 3). The new method reduces the error observed in Method 1 by correcting for temperature reading errors and temperature lag, aligning more closely with the estimates from Methods 2 and 3. This new method is believed to provide a more realistic estimate of gas usage. Missing data points in the latest density curve represent filtered-out values due to low temperature and/or high absolute dT/dt , improving overall accuracy of the estimate.

hardware malfunction. The new gas usage tracking method has provided greater confidence in the longevity of the nitrogen supply for the gDRT system, ensuring continued support for proximity science operations throughout the extended mission. The updated method has been incorporated into ground tools used by the SNC team, enabling more precise monitoring of the gas budget for the rest of the mission.

Percussion Mechanism Life

In November of 2022, a Corer percussion mechanism used in Earth-based testing experienced a hardware failure. This event triggered an anomaly investigation to better understand why the failure occurred and to learn how to prevent a similar issue from occurring in flight. As part of this investigation, the SNC team performed a deep dive into the use of the flight percussion mechanism over the mission and projected use of the mechanism into the future. This data was then used to inform the investigation’s test campaign.

The percussion mechanism consists of an actuator that drives a hammer up and down a drive shaft. The hammer makes contact with the anvil, which the bit sits against (see Figure 22). As the hammer is driven down into contact, the kinetic energy from the hammer is transmitted into the anvil, through the bit, and into the rock. This chiseling behavior slowly breaks up the rock over the course of the drilling operation. The speed of the hammer, and thus the frequency of percussion, can be changed during the drilling operation to maintain a constant rate of penetration. However, higher frequency percussion is associated with increased energy passing through the hardware, leading to the thought that high frequency percussion could accelerate life and contribute to an early failure. As a result, it was important to the team to not only understand the total duration of percussion used, but also the amount of time spent percussing at various frequencies.

The expected percussion life was initially tabulated based on a pre-launch ground test campaign in which a variety of rock types were abraded and cored in a flight-like manner. This data, combined with an estimate of the total number of

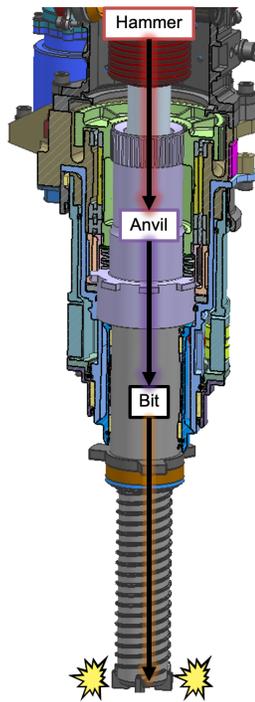


Figure 22. An internal view of the percussion hammer striking the anvil, which then transfers the percussive energy to a coring bit.

rocks to be drilled over the mission, resulted in a breakdown of the total duration of percussion in predefined percussion frequency bins. A life test unit was developed and successfully operated to three times the expected mission life without experiencing a failure.

However, for the anomaly investigation’s follow-on testing, there was a desire to re-evaluate the projected percussion life for each frequency bin based on use to date in surface operations. The results of these projections are shown in Figure 23. Surface use of the percussion mechanism to date, in blue, has yet to exceed the original life metric, in green, for all percussion frequency bins. For many percussion bins, the projected surface life, in red, is significantly less than the original life metric, indicating that fewer soft and very hard rocks have been encountered in flight than in the ground test program. However the projection exceeds the original life metric in the mid-range frequency bins of 31 to 37 Hz, indicating that surface operations have drilled medium to hard rocks more often than originally predicted. In addition, the overall projected percussion life duration is less than the original life metric, indicating that the margin added to these bins to account for failed drilling attempts and diagnostic activities has been underused thus far in flight. The increased use in bins 31 to 37 Hz is further explained by the variability in drilling durations, particularly for coring (see Figure 24). There is a minimum amount of low frequency percussion used in every coring attempt, but the individual use per bin is variable, even for cores collected from the same rock.

In general, the evidence showing that projected surface life is below the original life metric was welcome news to the anomaly investigation team and has been used to identify the appropriate frequency ranges and durations to use in follow-on testing. Further refinement of projected percussion use

will be performed over the remainder of the surface mission.

7. SENSOR RECALIBRATION

The Perseverance Rover employs a 2-meter, 5-degree-of-freedom Robotic Arm (RA) with a “Turret” end effector containing a Coring Drill, the gDRT, and multiple science instruments (PIXL, SHERLOC, WATSON). Together this equipment is used to perform abrading, sampling, and proximity science activities on the Martian surface.

The RA also contains a 6-degree-of-freedom (6-DOF) force / torque sensor (FTS), which is mounted at the distal end of the robotic arm, between the turret actuator output and the turret interface plate onto which the RA’s turret is mounted. The RA FTS is used to measure reaction loads on the turret hardware, enable quasi-static closed-loop control of forces / moments during critical sampling operations (e.g. drill stabilizer preloading for coring or abrading, corer docking for bit exchange), and provide fault-protection that ensures turret / arm hardware safety during all RA motions. This 6-DOF FTS is designed to operate across a significant range of loads ($\pm 800\text{N}$ / $\pm 200\text{Nm}$, nominal load range) and temperatures (-70 to +22 degC), and required a complex two-stage calibration process to meet its accuracy requirements. This calibration process was reported in detail in [13].

Anomaly & Root-Cause Identification

On Sol 521, a nominal motion of the RA was halted prematurely and mid-motion by active force-sensor fault protection. In this case, the RA FTS measured an excessive Fz load of -93N beyond the turret gravity wrench (Figure 28). This fault condition would typically indicate that the turret had made (inadvertent) contact with another object (e.g. rover hardware or terrain), but the entire arm and turret assembly was visually confirmed to be in free-space (i.e. well away from all collision hazards) at the time of the fault.

Analysis of RA FTS telemetry data yielded the conclusion that this fault was the result of three compounding factors:

Factor 1 – A large temperature delta existed between the RA turret and the RA force sensor (40 degC delta). This created a real load on the RA FTS, primarily in the Fz channel, due to hardware interactions (e.g. twist capsule covers) located between RA turret actuator output and the turret interface plate along load paths that are not through the RA FTS.

Factor 2 – The Fz channel was known to have an existing offset due to accumulated sensor drift over the time since the calibration was created (based on pre-launch data collected at JPL in 2019, approx. three years prior). The previous motion, on Sol 517, showed an offset of -50 N on the Fz channel.

Factor 3 – The FTS temperature was outside of its calibrated range by -15 degC. The FTS was measured at -75 degC, while its calibrated range for the active RA FTS calibration was -60 to +22 degC.

FTS Recalibration

In light of these revelations, it was determined that the appropriate corrective action was to recalibration the RA FTS.

Recalibration of the RA FTS followed the same strategy as outlined in [13]. We used the existing calibration as a starting point (since it relies on a broader range of calibration loads than can be measured on Mars), and perform least squares optimization to scale the existing per-axis calibration, plus

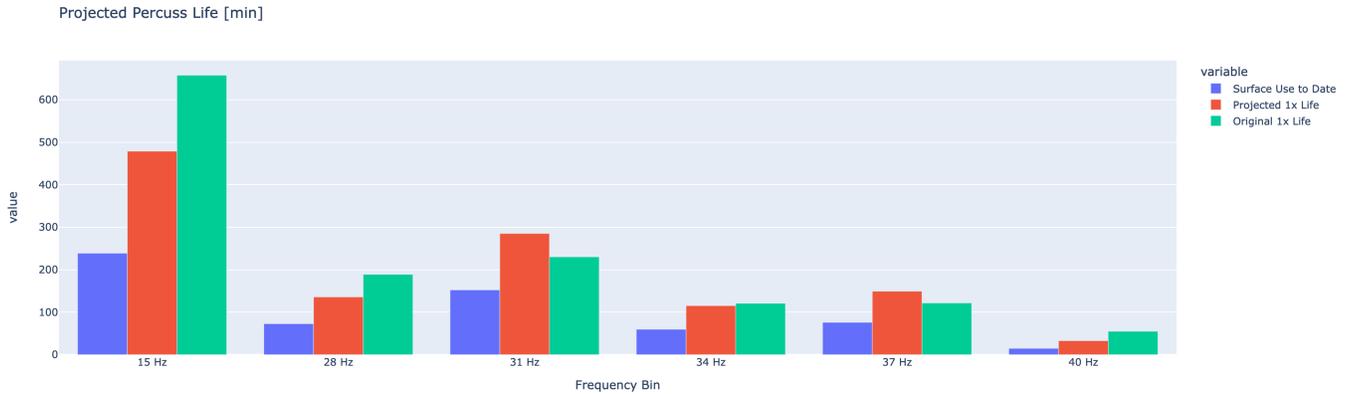


Figure 23. Comparison of the total percussion life, in minutes, collected over the surface mission, the expected percussion life based on pre-launch ground testing, and the projected percussion life based on the rate of surface usage to date. Percussion life is divided up into frequency bins, centered around each bin’s labeled frequency (i.e. 31 Hz covers 29.5 to 32.5 Hz).

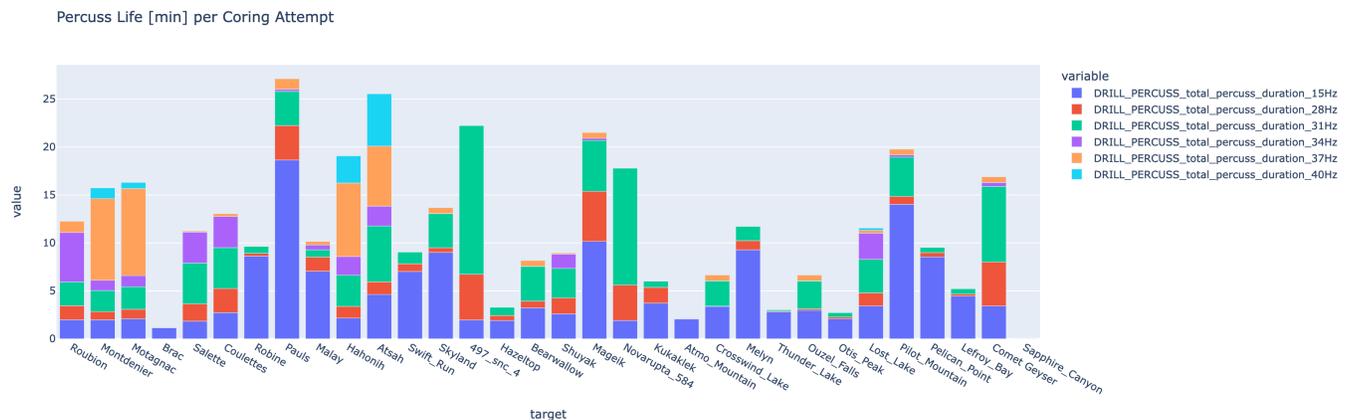


Figure 24. The percussion life, in minutes, accumulated per core, where each stacked bar is colored by the time spent in each frequency bin during the collection of that core.

further tweak some parameters in the calibration to minimize the sensor error relative to a pose-dependent gravity model for the turret weight. This optimization is done exclusively using newly collected data on Mars at multiple temperatures and the full range of Mars gravity-load conditions.

Regarding Factor 1 (internal loads on the FTS exacerbated by temperature), it is difficult to characterize the actual magnitude of internal hardware interaction loads since such loads are pose- and temperature-specific and they are implicitly coupled with the behavior of the RA FTS calibration. This is accepted source of error in our calibration.

Regarding Factor 2 (strain-gauge drift over time), trending of the RA FTS strain-gauge sensor drift was ongoing, but Sol 521 was the first time in the mission that sensor offset drift was observed to contribute to a nuisance fault during RA motion. The 2nd-order nonlinear nature of the Fz calibration (as compared to the 1st-order linear nature of the calibrations for the other 5 force / moment axes) meant that sensor drift on this channel was exacerbated relative to the other channels. This drift is able to be corrected by collecting data as close to the present Sol as possible, and ensuring periodic updates to the calibration on an ongoing basis.

Regarding Factor 3 (extrapolation outside of calibration data), the 2nd-order nonlinear nature of the Fz calibration meant

that extrapolating Fz loads beyond the calibrated temperature range added increased uncertainty compared to the other forces / moment axes as well as a pronounced negative bias for increasingly negative temperatures. This extrapolation error is able to be minimized by collecting data at a broader range of temperatures than were collected during the pre-launch calibration campaign on Earth.

This recalibration data is shown in Fig. 26, using data collected on Sol 595-621 at three different temperatures, ranging from -70 to -18 degC. These were the most extreme temperatures (plus a mid-range temperature) experienced by then.

The updated calibration for Fz is shown in Fig. 27, but all force and moment channels were updated in a similar manner.

We recomputed the Sol 521 fault loads (and Sol 523 recovery motion) using both the extant / updated calibration in Fig. 28. The updated calibration has healthy margin to fault limits, as desired. The updated calibration was uplinked to the rover on Sol 883 and has been operating successfully since then.

8. STAFFING CHALLENGES

The surface operations roles staffed by the M2020 Robotic Operations (RO) team are Rover Planner, Mobility Down-

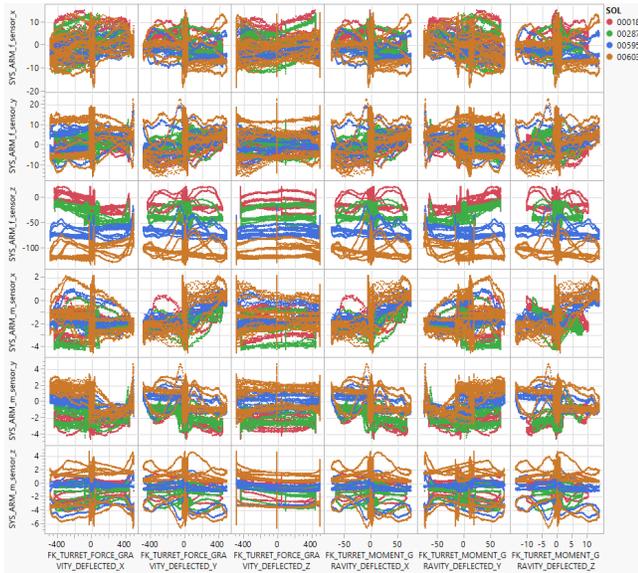


Figure 25. RA FTS sensor error over time (by Sol). RA FTS per-axis error (in N / Nm) collected during free-space motions of the RA turret, with no external applied load, versus per-axis RA turret gravity weight (in N / Nm).

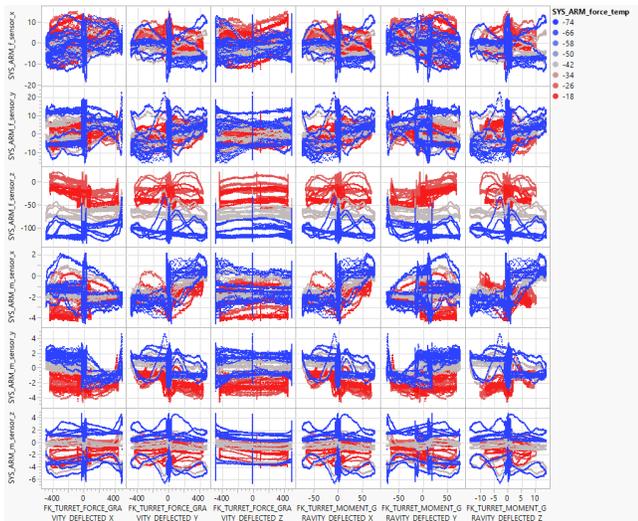


Figure 26. RA FTS sensor error over temp. (for Sols 595-621). RA FTS per-axis error (in N / Nm) during free-space motions of the RA turret, with no external applied load, versus per-axis RA turret gravity weight (in N / Nm).

link, Robotic Arm Downlink, SNC, Helicopter Integration Engineer (HIE), and Mechanisms. It is normal for there to be staffing loss and replenishment during the prime and extended mission as some individuals choose to leave the project to pursue other career-building opportunities within and outside of JPL. Over the first 35.5 Earth months since Perseverance landed on Mars, the RO team experienced an average departure of one person every 1.56 months. Over that period, individuals left the RO Team Chief, Rover Planner, Mobility Downlink, Robotic Arm Downlink, SNC, HIE, and Mechanisms roles at an average frequency of once every 36, 6, 9, 9, 4, 6, and 6 months, respectively.

Each RO role maintains a training program that certifies new

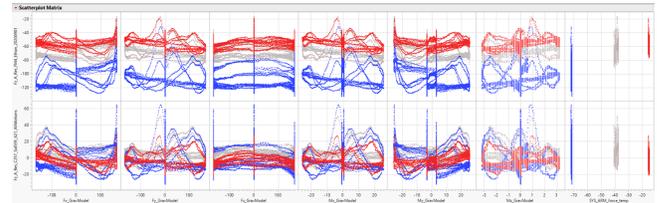


Figure 27. Recalibration of the Fz channel, with the original (top) and updated (bottom) calibration versus.

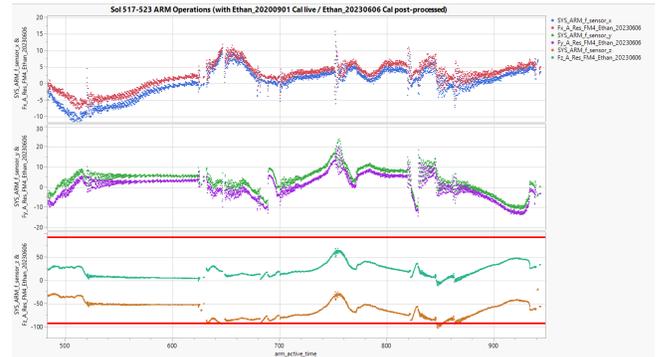


Figure 28. Sol 521 Fault Loads (plus recovery), using the original calibration as well as the updated calibration. RA FTS per-axis measured loads on Fx, Fy, and Fz (in N) versus time (in seconds). Fz fault limit violations are in red.

team members to backfill for departing team members. The time period required for training new RO team members generally varies from 2 to 12 months (at half time) depending on the specific role. The Mechanisms team has the shortest training period and the Rover Planner team has the longest. However, the training period for an individual can end up being longer for several reasons. Firstly, if the training class for a role is large, the trainees will invariably compete for opportunities to shadow tactical surface operation shifts (initially just observing, then taking on more and more of the actual duties while supervised by someone certified). Secondly, when surface operations are dominated by a specific activity for many sols, trainees will need to wait until that activity is completed to shadow other activities. For example, during a drive campaign to reach the next region of interest for Science Operations, opportunities for trainees to shadow on proximity science and sampling sols will be limited or nonexistent. Thirdly, training times can increase due to workload prioritization between RO and non-RO work. Non-RO work includes individuals that support other teams internal to M2020 and other JPL projects.

Roles with longer training periods proactively train several new individuals each year, anticipating there will be departures. Roles with shorter training periods tend to wait until a departure occurs to start a new trainee. When an individual decides to leave the RO team, a transition plan is generally implemented that may include starting a new trainee, closing out specific strategic work prior to the departure date, and/or transitioning strategic work to another RO team member.

In February 2024, there was a significant spike in the number of RO team members that left JPL. On February 7, 2024, JPL laid off 570 employees and contractors, nine of which were RO team members. Their departure was immediate with no opportunity for transition plans. One week later,

one additional RO team member had a preplanned departure from JPL for a total of ten departures within a one-week period, which corresponded to a 17.2% reduction in RO team staffing. The number of seats lost by the Rover Planner, Mobility Downlink, Robotic Arm Downlink, SNC, HIE, and Mechanisms roles was 1, 1, 4, 4, 1, and 1, respectively. The number of seats sums to twelve because two of the individuals that departed JPL were certified for two RO roles. M2020 cancelled tactical operations the day of the JPL layoffs, and the following day, all of the RO role leads confirmed they could adjust their staffing schedule to ensure all tactical shifts could be staffed for the remainder of the month.

The RO role which was impacted the most was Robotic Arm Downlink, losing four individuals which corresponded to 36.4% of its staff. With their reductions, the Robotic Arm Downlink team initially had no availability for any strategic work, such as new capability release, tool development, and anomaly investigations. Examples of specific tasks which were temporarily suspended included a Vehicle System Testbed (VSTB) robotic arm noise anomaly investigation (which precluded the use of the VSTB robotic arm), and work on a Go and Hover capability that could enable proximity science after an AutoNav drive without ground-in-the-loop to verify terrain safety. The Robotic Arm Downlink backfill plan returned a previously certified team member (who was working other JPL projects) to M2020 within eight days, cross-trained an individual already certified and actively working other RO roles (since the training time tends to be shorter for those already familiar with M2020 robotic operations), brought in a new team member to primarily focus on Go and Hover strategic work, and three team members reduced their time on other projects to increase their time on the Robotic Arm Downlink team. The previously certified team member that returned to M2020 was recertified for tactical operations approximately five weeks after the layoffs following several refresher shadow shifts.

The next most impacted role was the SNC team, also losing four individuals which corresponded to 23.5% of its staff. Two of the impacted individuals were fully certified and two were conditionally certified trainees. Conditionally certified SNC trainees can be staffed for surface operations when they are paired with a fully certified SNC team member. The largest impact to the SNC team was the loss of their deputy SNC lead, who was also the M2020 Corer Technical Authority. With these reductions, the SNC team had very limited availability for strategic work. Work was temporarily suspended on the anomaly investigation into the VSTB corer percussion no longer producing output force, which had precluded the use of the VSTB corer since late 2022. The SNC backfill plan returned three previously certified SNC team members (who were working other JPL projects) to M2020 within two weeks of the layoffs, and returned a previous Corer Technical Authority to M2020 to restart the VSTB corer percussion anomaly investigation. Two of the three previously certified SNC team members that returned to M2020 were recertified for tactical operations approximately five weeks after the layoffs, while the other returning SNC team member was tasked with strategic work.

Although the Mechanisms team only lost one seat, it corresponded to 25% of the team's core staffing and was significant because the team was already at its lowest staffing level for supporting surface operations, there was no opportunity for a transition plan, and the team was in the midst of a high-priority steer motor hardbrake anomaly investigation which had caused a Perseverance driving stand-down two days prior

to the JPL layoffs. After the third right-rear steer motor hardbrake of the mission on sol 1052, only 40% of its motor's hardbrakes remained with 72% of Perseverance's expected mission odometry still to be realized. Under its reduced staffing, the Mechanisms team surged their M2020 support by up to 4x. They efficiently selected, tested, and uploaded a new set of parameters to Perseverance, and on sol 1066, two weeks after the driving stand down started, Perseverance resumed unconstrained driving on Mars.



Figure 29. Final location of the Ingenuity Helicopter at Valinor Hills. Image courtesy NASA/JPL-Caltech/ASU/MSSS.

9. INGENUITY OPERATIONS

The Ingenuity Mars Helicopter's mission was officially declared completed on January 25, 2024, after 3 years of flights on Mars beginning with the first powered, controlled flight on another planet on April 19, 2021 and culminating in the helicopter's 72nd and final flight on January 18, 2024. Ingenuity flew a total of 17 kilometers, with a total flight time just shy of 130 minutes and a maximum altitude of 24 meters. During the landing sequence of its final flight, Ingenuity's rotor blades suffered damage rendering it incapable of any further flights. The exact cause of this incident is still under investigation, with the first publication expected soon. Its final landing site was named Valinor Hills (see Figure 29). Despite the damage to its rotor blades, Ingenuity's other various subsystems continue to operate relatively nominally, and command and data handling remains possible as before.

After Ingenuity was confirmed to be permanently grounded, its FSW team put together a final software patch that allowed the helicopter to run a default command sequence on boot-up. This sequence leveraged prior developments in using bash and Python scripts in parallel with command sequences to automate log engineering health telemetry and EVRs from the helicopter's computers—as well as daily images from both its cameras—in a compressed format. This data is intended to catalog the continuing status of the helicopter's systems in the Martian environment and potentially support future science and engineering efforts. This fully automated operations regime has been termed “quiescent operations” and the helicopter is projected to be capable of continued logging for a number of years before disk space runs out. A reusable rover sequence which collects file listings and limited telemetry is currently run on a weekly cadence without any need for operator staffing. In addition, quiescent mode

data is downlinked on an approximately monthly cadence, with RO's Helicopter Integration Engineering team (HIE) negotiating activity timing and delivering sequence files generated by Helicopter Operators (HOs) as usual.

We take this moment to once again thank all the people whose tireless labors gave Ingenuity wings.

10. SIMPLE PLANNER

As described in [1], Simple Planner is the flight and ground software functionality newly provisioned for Perseverance to control the scheduling and execution of flight activities. It is a wholesale replacement of the Curiosity and Mars Exploration Rover legacy "Master-Submaster" ("MSM") control construct described in [14], effectively closing that open loop. The functionality began use in October 2023 on Sol 934 and has impacted Robotic Operations behaviors and planning cadences, according to the design intent to harvest energy and time that would otherwise go unused, and to facilitate more aggressive plans. Every RO activity since Sol 934 has been controlled by Simple Planner.

Fault Recovery

Simple Planner is most impactful when energy intensive operations such as coring or driving fault, as it will forego all the heating and Rover Compute Element ("RCE") awake time for the remainder of the faulted activity and any later activity dependent on its success. On Sols 942 and 947, OBP saved ~1000W-h, or nearly 40% SOC, in each plan following successive faults in the lengthy sample acquisition and processing activities due to distinct idiosyncratic trips of fault protection for Hall Effect position sensors, later relaxed. In the Sol 949 3-Sol weekend plan, not only was sampling finished on the third attempt but a long drive of more than 300 meters also fit on the second of three Sols of the plan, where an early afternoon decisional downlink telecommunication pass did not curtail it. Without OBP, after a recharge Sol, the second sampling attempt would have occurred in the weekend plan and after faulting there, three more weekday planning cycles would have been needed to recover to the same point.

Event-Driven Operation

Simple Planner assumes a best-case duration estimate when scheduling serial activities and responsive heating and can also often recover time when upstream activities finish early. In practice, this means that nearly all Autonav drives start earlier than they would have using MSM and in the single, "grounded" representation of the plan, in which duration for worst-case margin and cleanup for each activity stacks. Simple Planner also shifts the nuisance risk for tripping a "safety deactivate" sequence, and so an adjustment to commanded cutoff is often made that allows driving a bit later. Together, these effects average tens of meters of additional distance per drive and have extended as far as 75m for individual drives. Collectively, during a time of relatively sparse Autonav productivity from October 2023 through September 2024, this has nevertheless added nearly 800 meters of drive distance.

Similarly, shorter warmup heating durations based on actually encountered and not worst-case theoretical temperatures often allow drives to start at or close to the earliest times available in this plan, rather than being locked into a later, unresponsive time. This can be useful when scientists elect to forego or shorten typical pre-drive Remote Science for lower interest, as determined from data only available in the

decisional downlink telecommunication pass. This vector of comparative utility is difficult to gauge but has been repetitive, reflected in more than a dozen instances of drive heating starting in the first second possible in a new plan. On Sols 1002 and 1226, Rover Planners also elected to end a drive by turning to more favorable headings better using the morning sunlight to preferentially heat drive actuator hardware in the next plan, where it was anticipated to materially impact the earliest drive start time. This yaw dependency can shift the onset of no-heat temperatures by several hours, and the strategy may see greater use in the future where the time or energy savings effectively add drive time.

Tolerance to delay has also largely eliminated a failure mode that recurred eight times using MSM between Sols 142 and 777, faulting numerous Robotics Operations activities – warmup heating remaining below a target temperature at a predicted preheat end time. During an instance of the FTS characterization activity described in Section 7 on Sol 1250, near the coldest date and time of year, the warmup zone attached to the RA wrist and turret actuators, which is coupled to the massive drill housing, took nearly 30 minutes longer to reach a target temperature than reasonably calibrated physics-based ground models or any past flight experience would suggest. Before Simple Planner, the first mechanism actuation when the zone was not yet ready to use would simply have faulted. However, Simple Planner blocked waiting before starting the activity, and the constraints provided by operators on that and all downstream activities were sufficient to tolerate the delay. The RA activity and later drive were saved, and two recovery planning cycles avoided.

Latent Potential

The rover activity most amenable to productivity enhancement with Simple Planner is Autonav driving for distance. It is a "continuous" activity that can just keep going, utility is additive, and Autonav driving is configured to be gracefully stopped for the day at appropriate offset to the procedural nominal minimum state-of-charge threshold. Yet, as described in Section 3, during much of the initial year of use of Simple Planner, Autonav drives were either not attempted in apparently difficult terrain or were circumscribed until a new record of performance could become established on variable terrain. Only in the Sol 1003 3-Sol plan did operators attempt the sort of weekend drive plan originally envisioned with Simple Planner – nearly 800 meters reasonably possible across two Sols of driving. Unfortunately, that drive faulted on the first Sol prior to commencing Autonav. (That Simple Planner dutifully harvests energy after such overreach helps to justify the attempt by eliminating a power-based risk trade for attempting longer drives).

In the first year of use of Simple Planner, ~10% of the total production of the Radioisotope Thermoelectric Generator ("RTG") has been saved from reduced RCE idle time awake and reduced actuator and instrument warmup and maintenance heating. Notably, this is more than 20% of the "discretionary" production of the RTG budget not subscribed by "mandatory" loads in the core avionics, survival heating, Heat Rejection System pump and telecommunication passes.

However, considering only a realized energy balance, approximately half of this gain, or ~5% of the RTG's output, was not able to be "claimed" in execution, instead being shunted away in the manner Simple Planner provides, using targeted RCE awake intervals to avoid bleeding. On occasion, this is unavoidable when multiple faults occur in successive plans. However, in large part, this reflects throttling in the

ground modeling from effective conservatism and infidelities approximating fluctuating commanded loads as uniform linear rates, as well as systemic early activity completion in execution. This bias is best represented in the minimum actual state-of-charge encountered in flight over the past year – more than 12% higher than the value tolerated – despite regularly modeling interaction with it on the ground. That is equivalent to almost two hours of driving, or more than 200 meters. In many weekend plans, this discrepancy balloons much higher still, often to 30% or more by the third Sol, even in un-faulted plans. Of 45 3-Sol weekend plans with Simple Planner through September 2024, only a single one has handed over to the next plan at a state-of-charge below 80%. 35 of those have handed over above 90%. Coming into the week’s planning topped up tends to bequeath shunting at some point during the week, especially if any fault occurs.

For comparison, this ~5% of the RTG’s output shunted was more than twice the amount of energy that was used by Perseverance for all driving in the same period of time. (Apart from the ~5% of the RTG’s output that was inspectably – based solely on energy balance – realized as additional or faster productive activity over the period of time, a substantial but difficult to gauge amount of *additional* activity with Simple Planner is enabled by higher incoming state-of-charge buying headroom to a ground modeled minimum or handover state-of-charge violation that limits the set of activities attempted, regardless of how much shunting ends up occurring in execution. This effect alone has likely impacted the Sol path on occasion, though the complexities and snowballing impacts that could have contributed to alternative Sol paths under MSM are nearly impossible to project accurately).

Looking forward, as the team has grown into a level of comfort with Simple Planner and ironed out some initial kinks, the Robotics Operations team anticipates accessing some of this unused energy for driving and other activity in moving beyond the limitations of a simple linear-impact ground power model, allowing Autonav drives to “bounce off” the actual state-of-charge encountered in flight. From first principles, much more than a kilometer of driving in a weekend plan is possible.

11. LESSONS LEARNED

Success in extended mission by designing for the unknown

FSW developers are often admonished to minimize (or even eliminate) multiple code paths when implementing system behaviors, to simplify the amount of Verification and Validation (V&V) testing that needs to be done overall in preparation for launch. But in an exploration setting unexpected problems arise and must be addressed. Although in some cases completely new FSW might be needed (along with a months-long V&V and review process), as we have discussed it is often possible to implement changes simply by updating a few system parameters and/or small FSW patches over days, not months. Additionally, building in the capability and processes to collect sensor (re)calibration data within the active mission ensures that we are able to maintain performance of key systems / sensors even as their behavior changes over time. We have been quite successful in addressing unexpected problems this way, thanks to coding strategies developed over generations of Mars rover development.

Integrated System Verification & Validation and ground software dependencies delay new technology deployment

The first in-flight demonstration of Global Localization (phase 2) took place over a year ago, yet it still has not been incorporated into tactical operations. The reason is that there are a large number of Ground Data System (GDS) tools that were impacted with an update to onboard software, and integrated system Verification & Validation (V&V) must be performed for FSW updates. Global Localization software was demonstrated on Mars on the HBS on Sol 859 (21 July 2023). This required local V&V which is fairly efficient. The FSW updates required to make use of the HBS co-processor were completed before that demonstration and used an approach for making localized FSW updates using components [15], which reduces risk of impact on other FSW modules. However, due to detailed integrated system testing and review and ground software updates needed to deploy a new dictionary for the component update, the new FSW was not deployed on Mars until eight months later. The first attempt to demonstrate integrated use (phase 3) was planned to occur on sol 1188 (23 June 2024), but due to a problem in a system-wide ground validation tool it was further delayed for over three months while the ground tools were updated, reviewed, and deployed. Similarly, tools supporting human localization on orbital maps were not written with autonomous updates in mind and may also need significant time for update, review, and re-deployment.

Prepare for recalibration

Although we have enjoyed substantial success calibrating cameras and instruments on Earth prior to launch, as missions continue many years beyond their lifetimes it can be critical to retain the ability to recalibrate sensors on the surface. Section 7 discussed recalibration of the Force-Torque Sensor, using a procedure that was developed post-launch. It is important to ensure sufficient measurement capabilities and ground-truth data are available onboard to enable post-landing recalibration, including calibration fiducials (for optical sensors), calibration targets / materials (for instruments), and known mass or stiffness structures (for load sensors). For critical systems, redundant but heterogeneous sensors (e.g. position encoders and resolvers) can be combined with physical structures (e.g. joint hardstops) to provide similar benefits.

Include all teams in flight software scheduling

M2020 is preparing for regression testing of the next release of FSW (S8.1), which is scheduled for upload in mid-2025. As with previous Mars rover missions, the M2020 FSW team is managed by the Engineering Operations (EO) team. The RO team is new to M2020, so unlike previous Mars rover missions, M2020 FSW developers are spread across both the EO and RO teams. In previous M2020 FSW releases since landing, it had not been a major concern that the EO team solely determined the FSW schedule, including code freezes. But given S8.1 contains 14 RO Engineering Change Requests which RO FSW developers are responsible for implementing, RO has found it challenging to meet code freeze deadlines set solely by EO and has had to negotiate code freeze extensions on two occasions. Going forward, RO and EO Team Chiefs have agreed that RO will be a part of scheduling code freezes for future FSW releases, which will take into account the RO FSW development and testing scopes of work.

Grouping brake channels is not so simple

Grouping together brake channels that are always activated together can help with reducing the electrical footprint of the brake control avionics, which can save space and money. However, this grouping will affect the brake performance and behavior in multiple ways, such as brake closure time and handling of faults. When the brakes do not operate as expected, it can lead to excessive wear and risk early motor or brake failure. This needs to be taken into account on future missions when designing brake avionics, fine tuning brake control parameters, and implementing stop methods for fault scenarios. One cannot assume that the brakes will behave in the same way as individually activated brake channels.

Future missions can have ripple effects

Perseverance has ten Wheel Steer Actuators (WSA), six for driving and four for steering. After the launch of Perseverance, the M2020 WSA engineering models that underwent pre-launch life testing were used in the construction of the M2020 VSTB, which is used for flight-like testing of new capabilities, uplink products, FSW releases, and patches. Due to a recent need to support future missions, additional life testing of WSAs is required and is currently underway. But no spare units were available, so during the summer of 2024, several previously life tested WSAs were harvested from the M2020 VSTB to support the additional WSA life testing. Unfortunately, their absence leaves the M2020 VSTB in a state where mobility is precluded. Since a portion of mobility regression testing of new FSW releases must be performed on the M2020 VSTB, an unintended potential impact of the additional WSA life testing is a delay in the completion of S8.1 FSW mobility regression testing and S8.1 FSW readiness for flight.

ACKNOWLEDGMENTS

This work is supported by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the with the National Aeronautics and Space Administration (80NM0018D0004). Thanks to the Mobility Downlink Team (including Darwin Chiu and Xianmei Lei) for mobility trending plots and data. Includes Pre-Decisional information in reference to Mars Sample Return – For Planning and Discussion Purposes Only.

REFERENCES

- [1] V. Verma, M. M. Maimone, D. M. Gaines, *et al.*, “Autonomous robotics is driving perseverance rover’s progress on mars,” *Science Robotics*, vol. 8, no. 80, 2023.
- [2] *Nasa’s perseverance rover scientists find intriguing mars rock*, Jul. 2024. [Online]. Available: <https://www.jpl.nasa.gov/news/nasas-perseverance-rover-scientists-find-intriguing-mars-rock/>.
- [3] V. Verma, M. Maimone, E. Graser, *et al.*, “Results from the First Year and a Half of Mars 2020 Robotic Operations,” in *2023 IEEE Aerospace Conference*, 2023.
- [4] V. Verma, M. Maimone, T. Del Sesto, *et al.*, “Robotic operations for the first sample depot on mars,” in *2024 IEEE Aerospace Conference*, IEEE, 2024, pp. 1–12.
- [5] J. Maki, K. Farley, K. Stack, *et al.*, “The mars 2020 three forks sample depot,” *LPI Contributions*, vol. 2806, p. 2875, 2023.
- [6] *NASA Exploring Alternative Mars Sample Return Methods*, 2024. [Online]. Available: <https://www.jpl.nasa.gov/news/nasa-exploring-alternative-mars-sample-return-methods/>.
- [7] M. McHenry, N. Abcouwer, J. Biesiadecki, *et al.*, “Mars 2020 autonomous rover navigation,” *Proceedings of the 43rd AAS Rocky Mountain Section Guidance, Navigation, and Control Conference, Breckenridge, Colorado*, Jan. 2020.
- [8] J. Nash, Q. Dwight, L. Saldyt, *et al.*, “Censible: A Robust and Practical Global Localization Framework for Planetary Surface Missions,” in *International Conference on Robotics and Automation (ICRA)*, IEEE, 2024.
- [9] V. Verma, J. Nash, L. Saldyt, *et al.*, “Enabling long and precise drives for the perseverance mars rover via onboard global localization,” in *2024 IEEE Aerospace Conference*, 2024, pp. 1–18. DOI: 10.1109/AERO58975.2024.10521160.
- [10] E. Benowitz and M. Maimone, “Patching flight software on Mars,” in *Workshop on Spacecraft Flight Software*, Laurel, MD, USA, Oct. 2015.
- [11] 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. DOI: 10.1007/s11214-020-00783-7. [Online]. Available: <https://doi.org/10.1007/s11214-020-00783-7>.
- [12] 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.
- [13] E. W. Schaler, J. Wisnowski, Y. Iwashita, *et al.*, “Two-stage calibration of a 6-axis force-torque sensor for robust operation in the mars 2020 robot arm,” *Advanced Robotics*, vol. 35, no. 21-22, pp. 1347–1358, 2021. DOI: 10.1080/01691864.2021.1946424. [Online]. Available: <https://doi.org/10.1080/01691864.2021.1946424>.
- [14] S. Kuhn, “From determinism to emergence: Systems engineering a step change in execution on mars,” in *2019 IEEE Aerospace Conference*, 2019, pp. 1–10. DOI: 10.1109/AERO.2019.8741741.
- [15] M. Maimone and A. Holloway, “Load and patch: Improving hot patch capabilities in curiosity’s flight software,” in *Flight Software Workshop*, Laurel, MD USA, Feb. 2022. [Online]. Available: <https://youtu.be/htMkASxX-ro>.

BIOGRAPHY



Vandi Verma is the Development Manager in the JPL Office of Technology Infusion and Strategy, and the Chief Engineer of Robotic Operations for Mars 2020. Robotics capabilities she has worked on are in use on the Perseverance, and Curiosity rovers, and in human spaceflight projects. She worked on the Mars Exploration Rovers and Ingenuity helicopter operations.



Mark Maimone is a JPL Principal in Autonomous Planetary Rover Navigation, and Mars 2020 Robotic Operations Deputy Team Chief, Mobility Technical Authority, and member of the Rover Planner and FSW development teams. He received NASA Exceptional Achievement Medals for designing and implementing GESTALT self-driving Flight Software for MER and MSL; contributed

to the Mars 2020 self-driving Flight Software; served as MSL Deputy Lead Rover Planner, Lead Mobility Rover Planner and Flight Software Lead; developed downlink automation tools for MER/MSL; and is on NASA's Lunar Terrain Vehicle Insight team. Mark has a Ph.D. in Computer Science from Carnegie Mellon University.



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 the lead of the Sampling and Caching operations team.



Ellen Thiel received her B.S. in Mechanical Engineering in 2016, and her M.Eng. in Aerospace Engineering in 2017, both from Cornell University. She started at JPL in the Mechanical Engineering division, working on a variety of projects, including M2020 EECAM, Mars Sample Return, and Sentinel-6, before moving to the Operations division in 2021. There, Ellen currently works as

the M2020 Mechanisms Deputy Team Lead, and serves as a mechanisms Technical Authority for both M2020 and MSL.



Noah Rothenberger is a Robotics Systems Engineer in the Robot Ops and V&V group. She received her B.S. and M.S. in Mechanical Engineering with a focus on Robotics and Computer Vision at ETH Zurich. At JPL, Noah is part of the Robotic Operations team, working as a Sampling and Caching Operator on the Mars Perseverance Rover. She previously served as an operator for the Mars

Ingenuity Helicopter. Additionally, she conducts research in the field of computer vision and mapping, developing advanced algorithms and image analysis tools.



Joseph Carsten is currently a Mars 2020 Robotic Arm Technical Authority as well as a member of the Rover Planner team. He has also served as a Rover Planner for the MER and MSL missions, and as an arm operator for the Phoenix Lander mission. Joseph has developed flight software for MER, MSL, and Mars 2020, contributing to the areas of robotic arm control, sample handling, and au-

tonomous navigation. He received his M.S. in Robotics from Carnegie Mellon University in 2005.



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.



Ethan Schaler is a Robotics Mechanical Engineer in JPL's Extreme Environment Robotic Systems group. He supports Mars 2020 operations as a member of the Rover Planner team, and is currently a Robotic Arm Technical Authority. His past work includes verification and validation of the Mars 2020 Robotic Arm and calibration of its force-/torque-sensor. Ethan is a 4-time NIAC

Fellow, and the PI of robotics concepts for the Moon and Ocean Worlds. Ethan holds a Ph.D. in Electrical Engineering from the University of California, Berkeley, with a focus on miniature sensors, actuators, grippers, and robots.



Evan Graser is currently a deputy lead of the Mars 2020 Robotic Operations Rover Planner team and Strategic Route Planning Team. He has also served as a Systems and Mobility Downlink engineer for the MSL mission, and was previously the Mechansims Technical Authority for the MSL mission. He received his M.S. in Aerospace Engineering from the University of Colorado Boulder in

2017.



Nadya Balabanska graduated MIT in 2020 with a Masters of Engineering in Computer Science, with a research focus on motion-planning algorithms. Since then, she has been working in JPL's Operable Robotics Group as a member of the technical staff.



Stephen Kuhn earned his B.S. in Mechanical Engineering at Carnegie Mellon University in 2008, and earned his M.S. in Aerospace Engineering from the Massachusetts Institute of Technology in 2010. He held the role of flight systems engineer for the CHIMRA mechanism on the Mars Science Laboratory mission, rover planner on MSL and Mars 2020, and held other systems engineering roles on MSL and the Mars 2020 missions at JPL. He is the On-Board Planner flight systems engineer for Mars 2020.



Harel Dor is a Robotics Systems Engineer for the Robot Operations Group. They have worked on various simulation, modelling, and operations tools for MSL, M2020, and Mars Helicopter, and lead development of the Proximeter terrain assessment tool and Periscope arm placement refinement tool. They currently serve as Rover Planner Arm Technical Lead for MSL and the Tactical Operations Lead for Mars Helicopter. They received their BS in Applied Physics and Computer Science from Caltech in 2020.