Mars 2020 Simple Planner

Operations Rollout

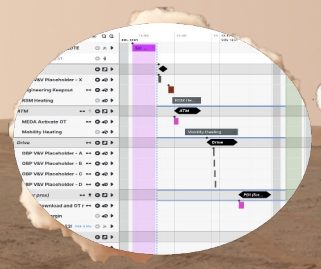

Nick Waldram

Jet Propulsion Laboratory
California Institute of Technology

March 19th 2025

JPL Clearances
URS 332091
CL #25-1175

© Copyright 2025, California Institute of Technology. Government sponsorship acknowledged.

Talks on the Mars 2020 Simple Planner

Topic	Speaker	Date
Overview of Simple Planner	Moffi	5th December 2024
Onboard Planner: Flight Software	Gaines	4th February 2025
Onboard Planner: Trusted AI on Mars	Reich, Chien	18th February 2025
Simple Planner: Ground Tools for Operations	Connell	25th February 2025
Simple Planner: Systems Engineering Operations with Autonomy	Hazelrig	11th March 2025
Rollout of the Simple Planner	Waldram	19th March 2025

You are here

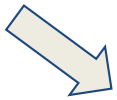

Missed a talk?

Recordings archived on JPLTube
Slides posted at https://ai.jpl.nasa.gov/public/projects/m2020-scheduler/

MARS 2020 PERSEVERANCE

Introduction

Simple Planner (SP) is the flight and ground system that enables the Mars 2020 Perseverance Rover to **adjust to unexpected state**, such as Martian temperature fluctuations or battery performance and **activity execution feedback**, such as activities failing, ending earlier or later than expected.

Simple Planner development began in 2016, and its first use was October 5th, 2023

This talk will recount the story of SP's deployment so far, which has involved incremental adoption of increasingly autonomous capability.

Key Takeaway: An incremental deployment approach, combined with enthusiastic engagement from personnel from across the Mission, has enabled increased autonomy without short-term sacrifices to science or operability. This increased autonomy has made the rover significantly more efficient.

Rollout Planning

Deployment of Trusted Autonomy into Operations

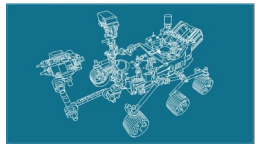

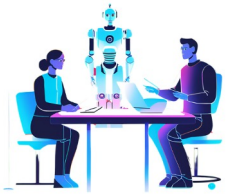

flight build

Design → Prototype → Analyze → Build → Test

Prototype ← Analyze (iterate)

Test → FTA, Train → Deploy

not flight version!

Kevin Reich's talk

Steve Chien's talk

Focus of today's talk

Integrating Autonomy into the Mission System

- Top priorities were operability, and project-wide buy-in
- **Unique challenge**: how to achieve a major paradigm shift *without* disrupting operations?
- **Solution**: deliberate gradual deployment of capabilities

“Crawl, Walk, Run”

- Crawl – pre-deployment activities intended to validate the SP design end-to-end (E2E) via high fidelity tests, executing both on JPL’s testbeds as well as onboard Perseverance itself

- Walk (SP1) – like an accordion. Activities can grow or shrink in duration, and start early or late, but order of activities is maintained

- Run (SP2) – optimize and increase the utility of a sol by opportunistically adding activities to the plan as allowed by energy, time, and other constraints, and by honoring more sophisticated dependencies

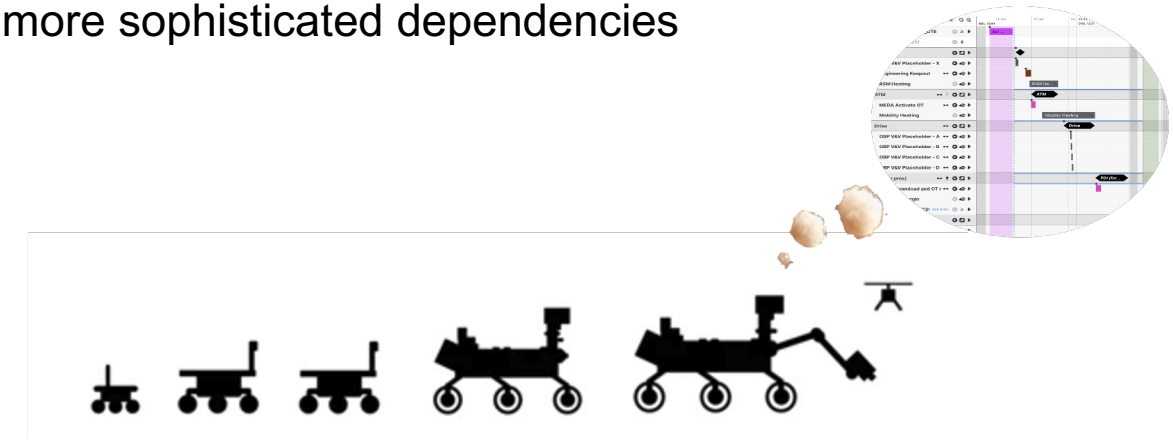

SP1

Accordion-like compression of execution timeline, as activities end early, and idle time is recouped.

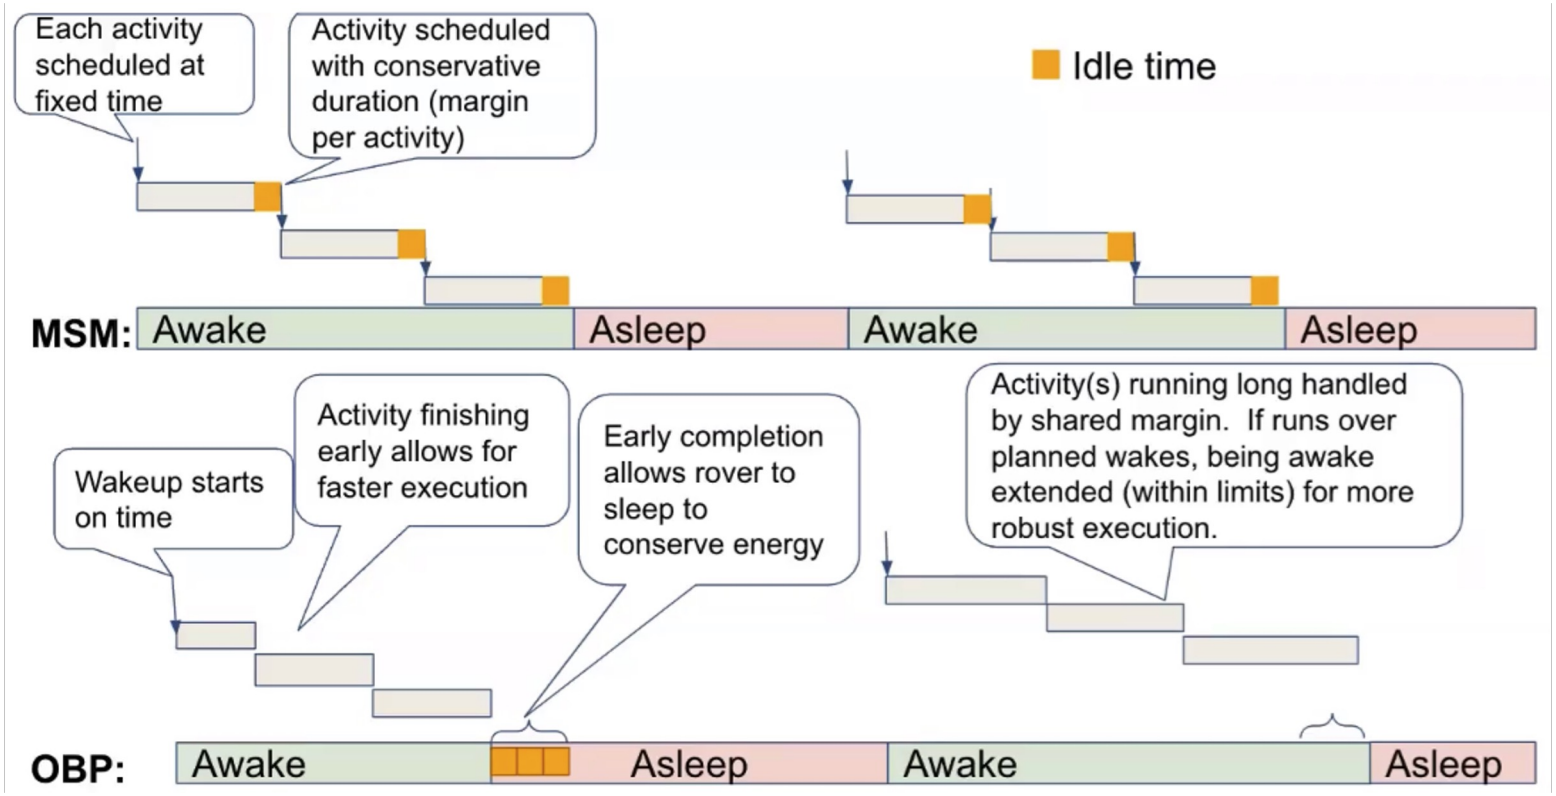

SP2

SP2 enables re-ordering of activities within user constraints, extending drives, and opportunistically including optional activities.

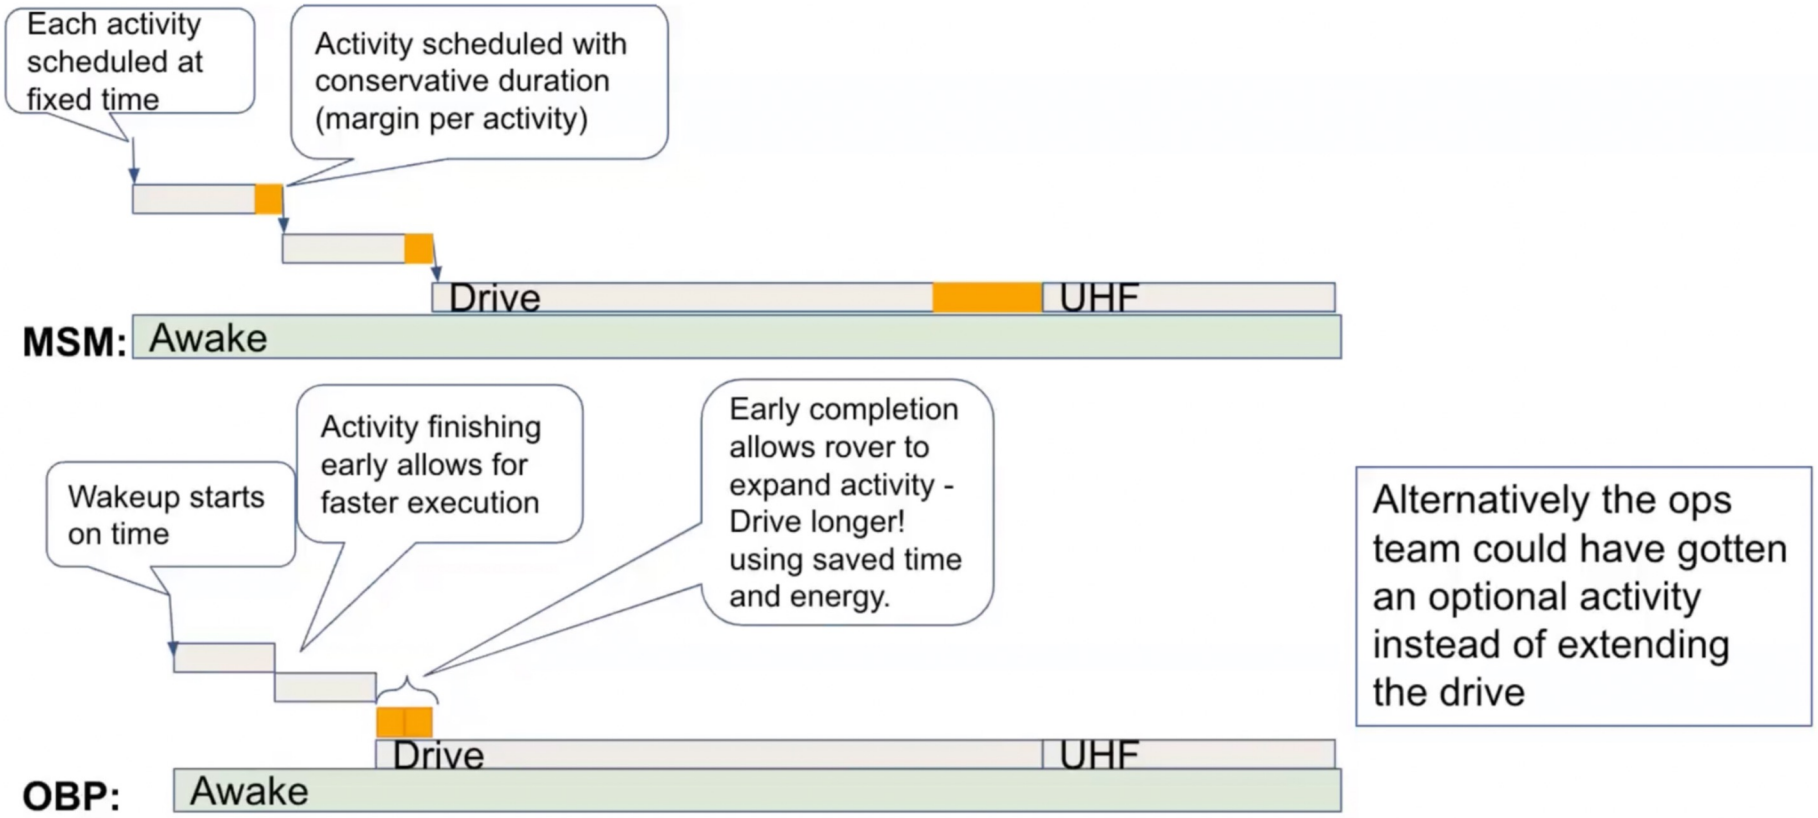

Utility, Capability, and Need

- In the lead-up to SP deployment, Simple Planner Working Group looked across traditional ops capabilities, made a list of anything that didn't have a solution using SP
- If we didn't meet this threshold across all ops use-cases, we'd need to revert to traditional ops paradigm
  - would require ground system to remain indefinitely backwards compatible, meaning specific "mode transition" procedures, ops training and familiarity with multiple paradigms (confusing)
  - risk of commanding errors, tool issues, longer planning timeline
- Threshold for ops readiness was "minimum viable product" – **need** the particular ops use-case to "play nice" with SP, without causing unreasonable burden on operators, but also without over-engineering/costing significant short-term development time

Utility, Capability, and Need

- On Board Planner (OBP) FSW has many **capabilities** which were not, and will not be, "switched on" in SP operations on Mars 2020
  - examples:
    - data volume constraint optimization
    - expanding activities
    - preferred time
    - complex CNF dependencies
    - pause-for-UHF
    - switch groups
  - while the individual rationales for de-scoping these capabilities vary, what is common is the **utility** to ops did not surpass the cost (V&V, ops dev, thread tests)
- Certain specific needs of the system are not inherently aligned with the principles of flexible execution, due to design or hardware limitations, and had to be spliced into SP (examples on following slide)

Utility, Capability, and Need

Hardware limitation

- Sample Handling Arm (SHA) heater on/off switch mechanically applies more load to arm (PFR 67066). To maintain SHA force sensing accuracy budgets, prevent tripping fault protection, and prevent hardware damage, need to FORCE rover to stay awake across multiple OBAs (“unnatural” for OBP)

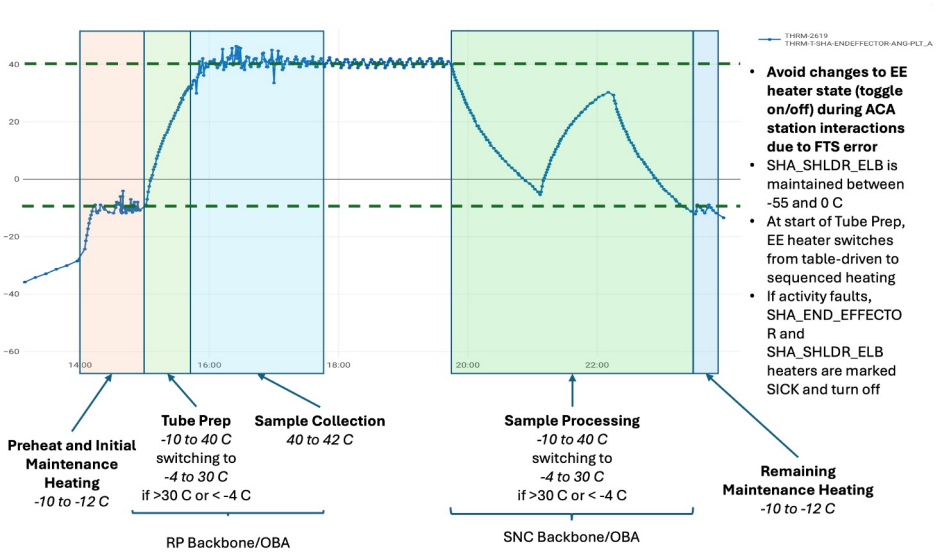

Thermal design limitation

- SHERLOC instrument has complex thermal design, few nodes, and thin margin between AFT min and max, thus heating execution flexibility (key behavior in SP) not well tolerated by instrument

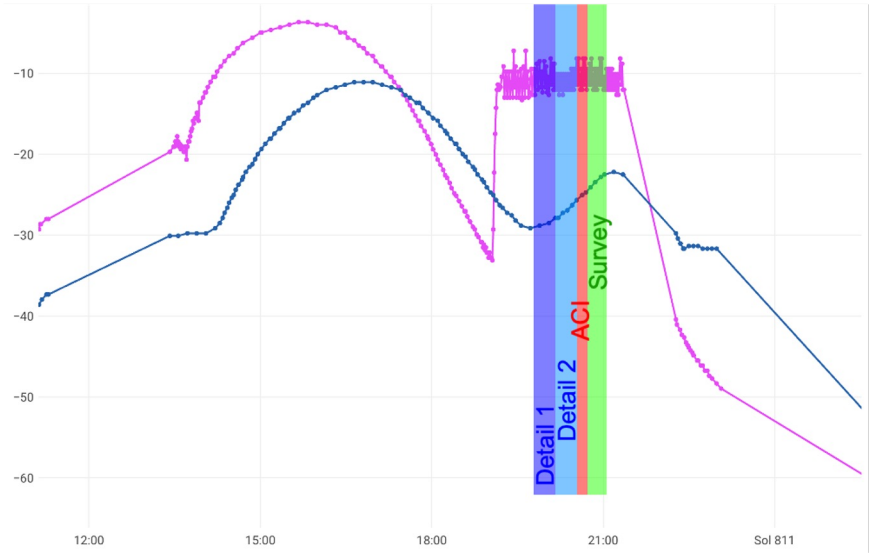

Working Group Evolution and Triage

- The rollout of SP capabilities in ground and flight software into operations followed the incremental “crawl, walk, run” template, from pre-deployment testing through SP2
- Along the way, the utility of certain capabilities were re-evaluated against the needs and desires of operations, and based on cost and complexity, some features were descoped
- Throughout this process, working groups navigated these decisions, which drove the development, triage, and bug fixes necessary to support ops in real time, including:
  - Simple Planner Working Group – broad charter, intersection of Instrument, Science, Engineering, and Robotic Operations
  - Simple Planner Teamtools Working Group – ground software development and support for the brand new “SP-only” tools such as SP Dashboard, As Run RML, snc_bridge_seq.py
  - MOS-APSS Weekly Tagup – triage for APSS ground software tool development

First-Time Activities

First-Time Activities (FTAs)

- FTAs formally exercised increasingly realistic and complex OBP control of the spacecraft in the mission operations environment and **on Mars**
  - In contrast, the Verification and Validation (V&V) campaign of OBP FSW predominantly leveraged software-only simulation to determine whether the software was
    1. comprehensive in its coverage of nominal scenarios and edge cases
    2. responsive to the demands of software requirements and high-level expectations from the mission level

- The primary function of the FTAs was as commissioning activities to verify that FSW worked as expected on the flight system, and in doing so, provide validation that our vast array of V&V was applicable to flight.
  - Testbed venue execution (flight hardware) as well as execution on the flight vehicle both corroborated V&V

- FTAs were not performed in an “ops-like” way – plans were strategically crafted well ahead of time, then slotted into traditional ops process.

First-Time Activities (FTAs)

- Part 1 – “toe dip” into OBP control of flight vehicle for ~30min. Demonstrated basic functionalities of scheduling, starting activities early or late, aborting activities that run long, and deactivating an OBP plan. Used “dummy” activities only (didn’t command mechanisms)

- Part 2 – incorporated comm windows and shutdown/wakeup with OBP in control, and preheating for select activities both while awake and asleep (“dream mode” heating). Also used dummy activities

- Part 3 – OBP controlled spacecraft for an entire sol. Executed real engineering and science activities

- **FTAs completed with success!**

“FTA Part 2” Storyboard

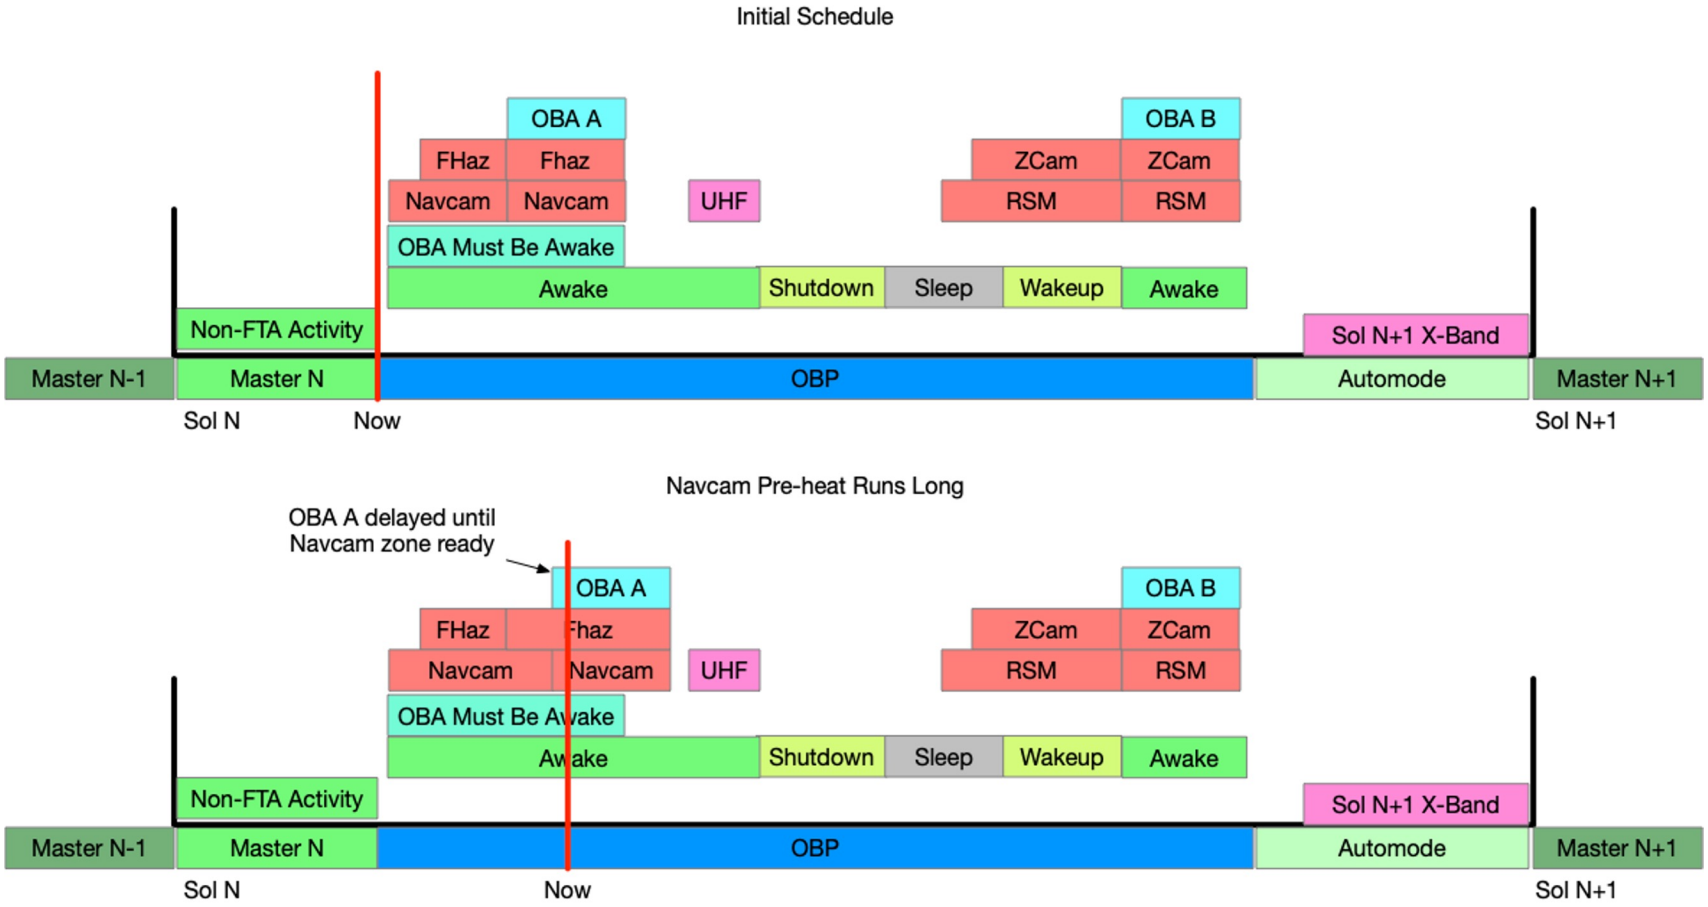

MARS 2020
PERSEVERANCE

Procedures, Training, Operability

Tabletops - tactical planning with ground scheduling tools (not E2E, no FSW in the loop)

Tabletop Date	Focus
4/13/23	3 different plan types - creation of constraints by ops team
6/15/23	Campaign Implementation (CI) → Tactical Planning, created a 3 Sol plan which contained activities most instruments, and arm and mobility backbones, participation of Instrument Operations (IO) teams
6/22/23	Ground software acceptance testing performed in SP mode
6/29/23	Coordination between the PILOT (plan coordination) and PULs (instruments)

Tabletop Date	Focus
8/10/23	Operations flow to transition from OBP → MSM and MSM → OBP operations modes
9/7/23	Implementation of sampling sol in OBP mode, included identifying process updates and role responsibilities
9/7/23	Sequencing changes in OBP mode, key differences between MSM and OBP COCPIT plans, practice “translating” the OBP plan into sequencing

“Flight School” Training

- Seven sessions with over 10 hours of material
- Taught live in sessions during July 2023 and recorded for future training
- Over 150 M2020 staff participated
- Sessions included:
  - Simple Planner Foundations
  - Science Operations and Instrument Operations Overview
  - Campaign Implementation
  - Simple Planner Planning
  - Tactical Uplink
  - Tactical Downlink
  - Anomalies

OBP Chair Certification and Staffing

- OBP Chair – brand new role in both uplink and downlink, new tools, processes, and procedures
- Became subject matter expert for OBP scheduling and execution behavior, constraint-based planning in tactical operations, validation of new OBP products in downlink and uplink, and new SP tools (SP Dashboard, AutoRML)
- Responsible for generating the as-run COCPIT plan
- Designed to be short-lived – lasted for ~6 months, off-loaded responsibilities to other engineering roles
- Having an SP expert in the room helped bridge the gap in training between traditional ops and new paradigm

Coordination with Instrument and Science Operations

- Coordination with Instrument Operations and Project Science was critical to successful SP deployment
- 2/2023 - 5/2023 the SP team held weekly technical interchange meeting (TIM) with the Instrument Operations teams
  - Develop the SP concept(s) of Instrument Operations
  - Work through changes in SP operations; ensure that science intent would be achieved with these changes
  - Communicate and assist with training in the new SP operations concept
- Several detailed briefings also kept Project Science leadership informed of development “as it occurred”
- Science and instrument teams were heavily involved in preparatory exercises (e.g., Tabletops, ORTs, flight schools)
- Ground Software SP minimum viable product and critical features and bug fixes were prioritized so that the above activities could use staged SP-like software whenever possible

Communication and Messaging to Science Team

- SP can adjust times but YOU still control when your activities are allowed to happen
- SP just adjusts the time of activities, YOU control the content (as always)
- The new system will NOT degrade science (and in fact there are some ways it can get you more of what you want)
- **This IS worth the effort of a big paradigm change because of the power savings – efficiency SP provides becomes *increasingly important* as RTG power output declines (as expected)**

End-to-End Testing

Mission System Testing

Increasing fidelity

Test Type	Definition
Operational Readiness Tests (ORTs)	Ops system level tests. Flight-like timelines, staffing, facilities, processes, procedures, protocols, products and tools using a high fidelity test venue to emulate the flight system. **Project-level review (ORR) to assess outcomes.**
Super Thread Tests/Mini ORTs	Ops system level testing for **limited but key capabilities** across **select teams**. Flight-like staffing (but limited to select teams), process, procedures, products and tools using a high fidelity test venue to emulate the flight system (as needed). Flight-like timeline and flight facilities as required.
Capabilities Tests/Thread Tests	Process, products, and tools that cross teams (usually does not include flight-like staffing, timing, facilities, and testbed), at varying levels of fidelity.

Super Thread Tests

- **Super TT-20 (12-14 July 2022)** focused on remote sensing and included uplink planning
  - VSTB (high-fidelity, RAD750 flight computer) execution of SSIM-validated planfile, and downlink
  - Super TT-20 had limited strategic or lookahead plan development
  - Was constrained in the types of activities included

- **Super TT-22 (14-17 November 2022)** focused on including a more complete operations cycle as well as more rover activities
  - Super TT-22 included
    - Campaign Implementation (lookahead planning)
    - Tactical uplink planning
    - VSTB execution of planfile
    - Downlink assessment
    - Second uplink planning session
  - The plan included remote sensing as well as an autonav drive (first drive with OBP in control)
  - TT-22 was a more flight-like exercise of uplink and downlink tooling and involved targeting in ASTTRO for robotic activities

SP-ORT

- 2-14 August 2023

- Flight-like timelines, staffing, facilities, processes, procedures, protocols, products and tools using a high fidelity test venue (VSTB) to emulate the flight system.

- Ambitious undertaking in terms of time, planning, team buy-in – within the heritage of Mars surface operations at JPL, exercising an ORT during operations is rarely, if ever, attempted. Required ops to plan on “off days.”

- High level objectives:
  - Process Objectives -- Include as many teams/tools/processes as possible on uplink and downlink.
  - Activity Objectives -- Further characterize performance of OBP running on VSTB (flight computer and mechanisms), exercising more key scenarios and more “flight-like” plan content.

- Campaign Implementation → Tactical →Run on VSTB → Generate downlink products

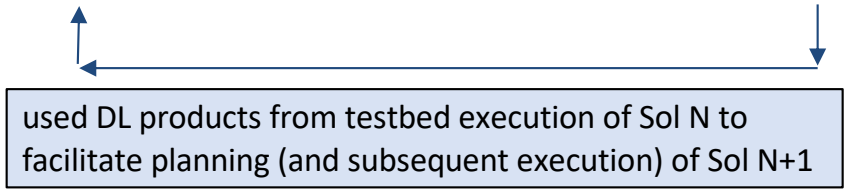

SP-ORT – Building the first plan

MARS 2020
PERSEVERANCE

SP-ORT Objectives

Process Objectives	Applicable Sol	Met?
Ensure that RP tooling supports crafting OBP plan	1261 and 1262	Yes
Test the translation of GDS/other tools to a plan file, through testbed execution	1261 and 1262	Yes
Exercise DL tooling to ingest DPs and display SP tools, and test OBP data compatibility with other non-OBP tools/processes	1261 and 1262	Yes
Exercise the OBP role in both downlink and uplink, including analyzing OBP health and safety, delivering the as-run COCPIT RML, and assisting with building/validating the command products	1261 and 1262	Yes
Ensure incon generation, delivery, and propagation succeeds for both OBP and non-OBP products	1260 and 1261 and 1262	Yes
Exercise and drive to closure SP1 requirements	1261 and 1262	Yes

Activity Objectives	Applicable Sol	Met?
OBP to OBP handover	1261 into 1262	Yes
Rover awake at handover time	1261 into 1262	Yes
Heating over handover	1261 into 1262	Partial
Dream Mode Heating	1262	Yes
Test PIXL with IN_PROGRESS activity dependency	1261	Yes
Test Plan Setup File	1261 and 1262	Yes
Testing implementation of "gatekeeper" at grounded start time	1261 and 1262	Yes

Training Objectives	Applicable Sol	Met?
Participation from as many individuals per role as is feasible (either as prime or shadow/observer)	1261 and 1262	Yes
Exercise updated role procedures for SP operations	1261 and 1262	Yes
Introduce and use a common vocabulary and terminology for SP operations	1261 and 1262	Yes
Understand the key differences between M/SM and SP paradigms	1261 and 1262	Yes
Understand end-to-end how the ground operations tools and products generated are executed, what is occurring onboard the rover with OBP, and how the results are assessed and interpreted in downlinked data	1261 and 1262	Yes
Practice production of cross-team and role inputs and outputs correctly, efficiently, and within the timeframes required for tactical operations	1261 and 1262	Yes

SP1

SP1 Plan Content

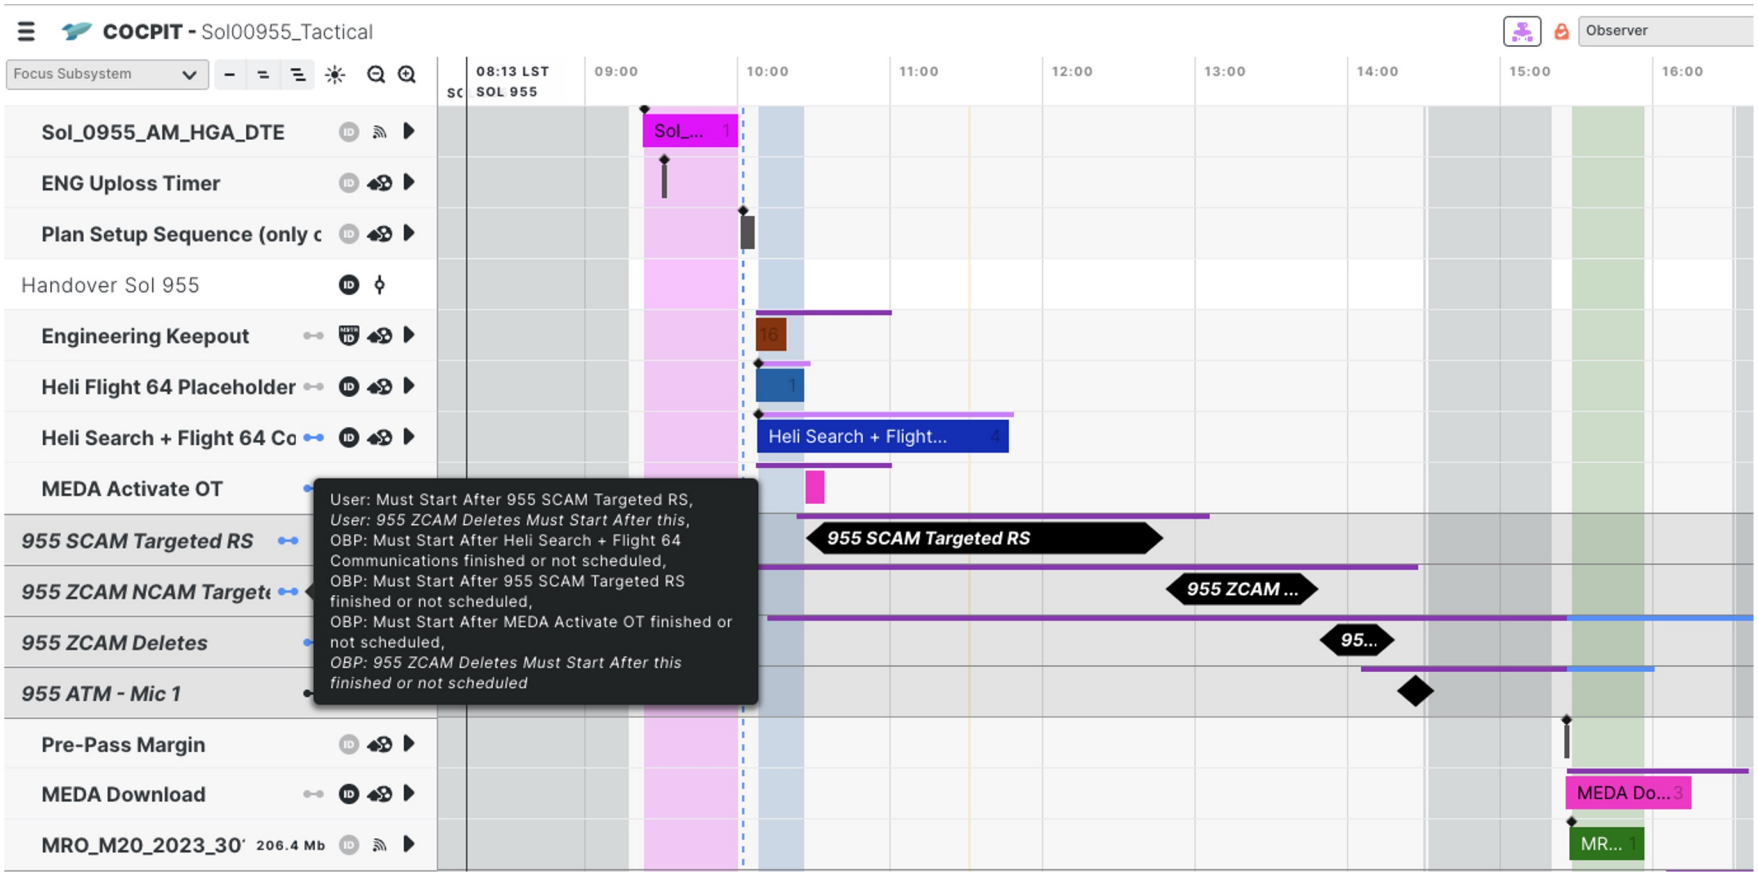

It's doing what we told it…what now?

- Once operators became accustomed to constraints, tool developers validated their ground model's constraint-encoding, and the mission harvested the energy and time-savings from the planner, the trust was established – this thing works and it's *better* than what we had before

- Other than responding to critical bug fixes and procedure updates, the highest strategic priority for Simple Planner post-deployment was to **lift** SP1 policies

- While SP1 meaningfully achieved many of Simple Planner's energy saving objectives, by overly constraining the plan, it limited the ability of the flight system and the user to exploit the capabilities of SP

SP2

SP2 Plan Content

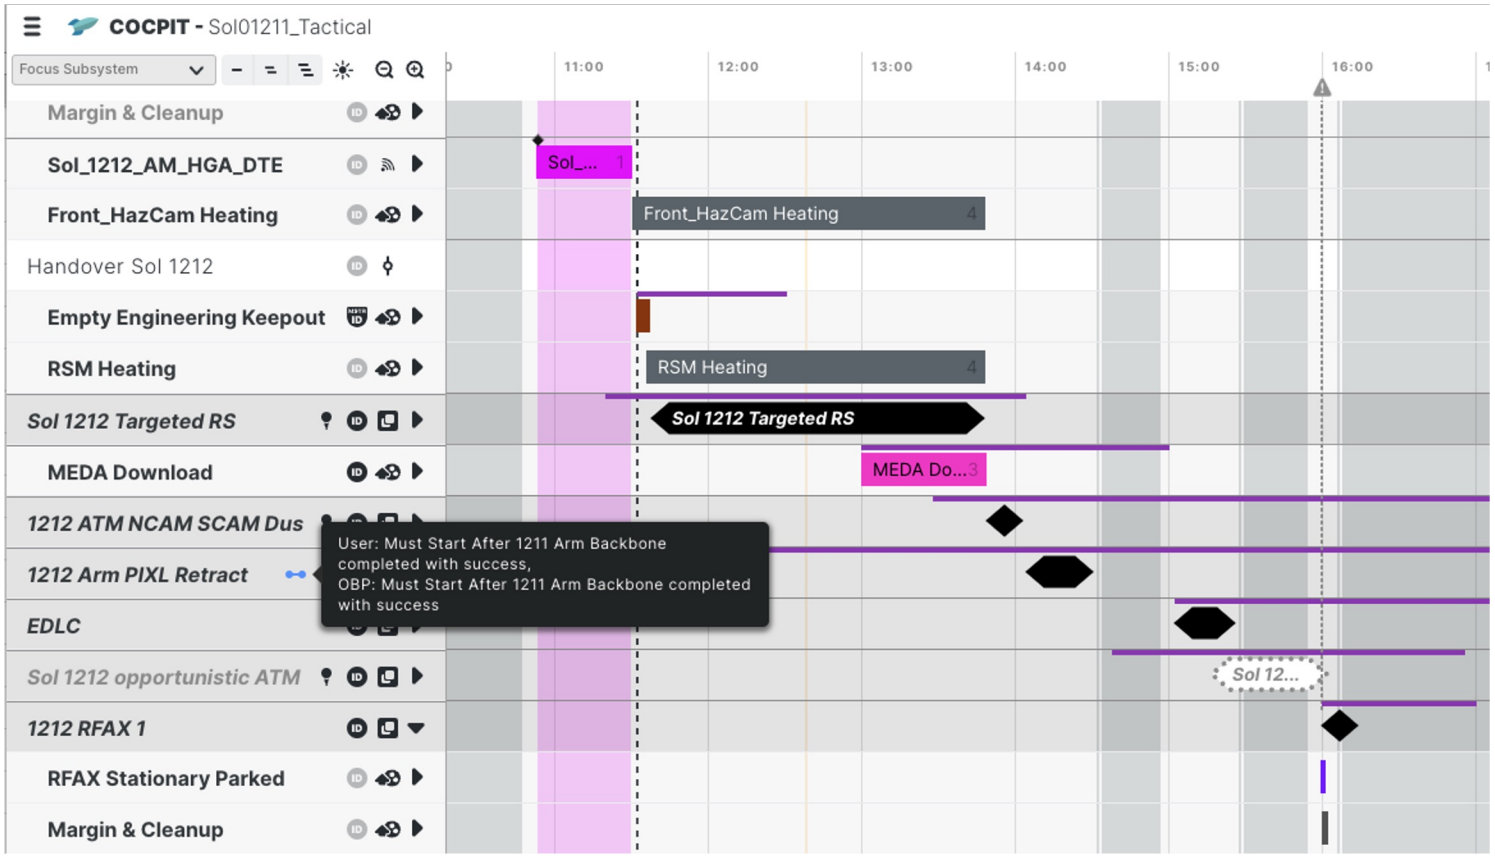

Early-Start Drives

- For drives, we often want to “drive as far as we can” – it behooves us to:
  - Start the drive as early as we can (once mechanisms are heated)
  - Drive as long as resources (battery) allow, up until a time-of-day cutoff limit
- Under SP, we can start drives early relative to traditional ops paradigm
  - In SP, margin not encoded in activity’s scheduled duration, unlike traditional ops
  - Prior activities ending much earlier can allow scheduler to place drives early
- The rover’s available power influences how much duration we can give to autonav drives when planning (can’t dip below 40% battery SOC), yet on Mars we find ourselves with a more optimistic power situation than in planning allowing a longer duration
  - In execution, sequencing can enforce cut-off if battery charge dips too low

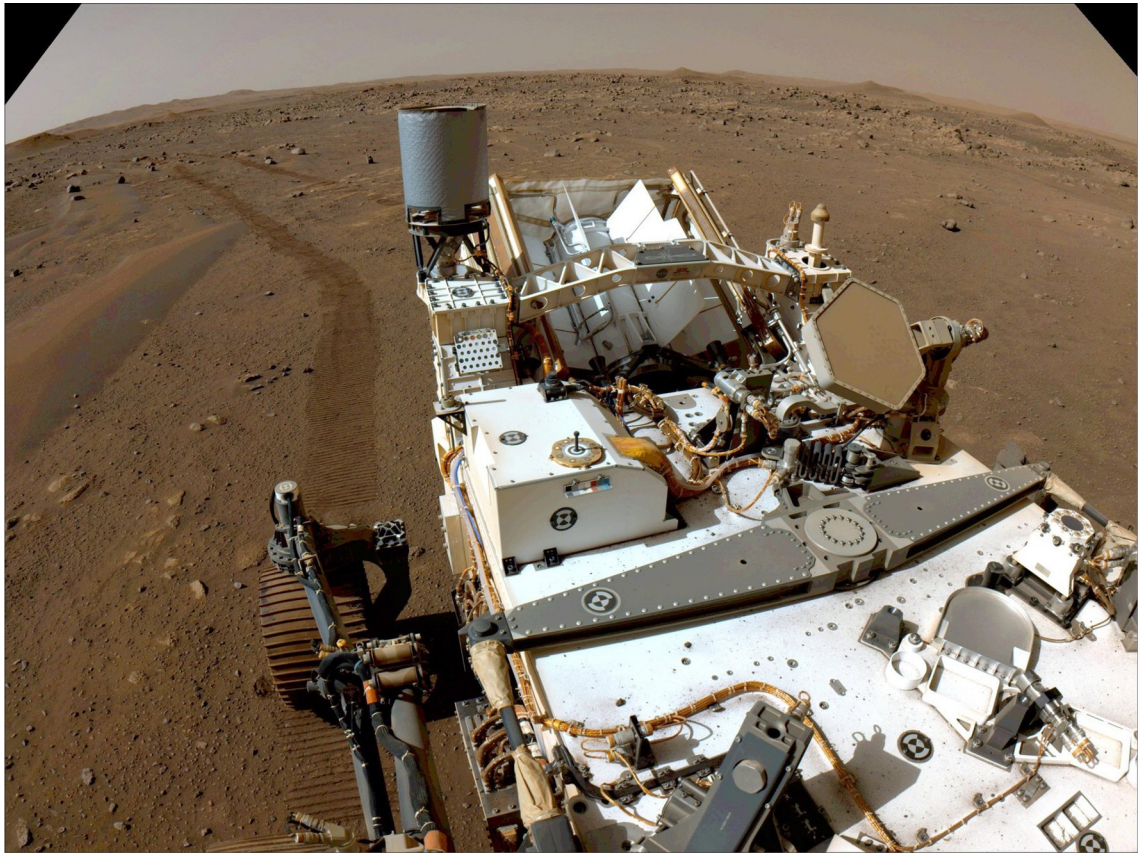

Optional Activities

- Optional activities allow desired (not required) activities to fit into the sol, should things onboard go **better than assumed in ground modeling** (more power/time than predicted)
- atmospheric science have proven to be a valuable application – typically run at the same time each day, if they fit – and overwhelmingly, they *do* fit
- **"*TL; DR: 'optional' is our friend with OBP*"** – scientist in response to a dust devil movie running 2 hours earlier than planned, at a time more conducive to acquiring the observation.

Adapting to Change

ISA Summary

- **25 total SP-related ISAs** (wide net, given that SP encompasses anything new in ops relative to traditional paradigm, which is a lot)
  - Six ISAs related to activity failures
  - Some cases in which activity failed to schedule at expected time but still able to be scheduled
  - Other contributing causes: power model inconsistencies, incorrect encoding of parallelism in ground model, idios closed UAI (mostly learning)
  - Far more ISAs for issues caught on ground, including integration validation (SSim), than on flight vehicle (yay!)
- **SP-related ISAs represent small fraction of total (159)**
- **Anomalies have not prevented operations from using SP on subsequent planning cycles**

1.5 OBP FSW issues
- One case in which an activity failed to schedule (flight software contributed but was not the only cause)
- One in which an off-nominal plan activation (via UHF forward link) uncovered a flaw resulting in plan file rejection

3.5 OBP Thermal Relation Flaws
- Thermal KIP deliverables underwent significant change from traditional operations, including new scripts and processes. Automation is as-of-yet imperfect and requires manual review (prone to operator error)
- Ground tool updates to add automated robustness to “brittleness”
- Thermal behavior under OBP is complex!

1 Ground Tool Process Issue

Total M2020 ISAs, 10/2023 thru present

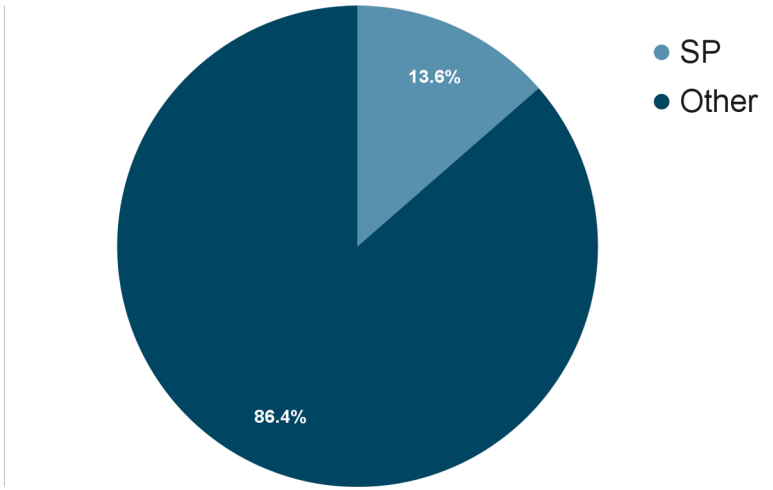

New Engineering Challenges

- We have gotten much better at conserving energy, but can be challenging to harvest this energy – conservative ground model tells us we'll dip lower than we actually do onboard

- We charge the battery when asleep, OBP will keep the RCE on to avoid charging past its limit of 99.999% – helpful in that we don't need to shunt excess energy, but introduction of "awake" periods draws the battery lower than shunting through RPA bleed resistors

- Systematic overmodeling presents a challenge – how do we harvest all the power gains from SP i.e. not shunt away savings?
  - reconciling power models helps
  - more optional activities helps
  - dipping to a lower battery charge helps (takes longer to charge battery)

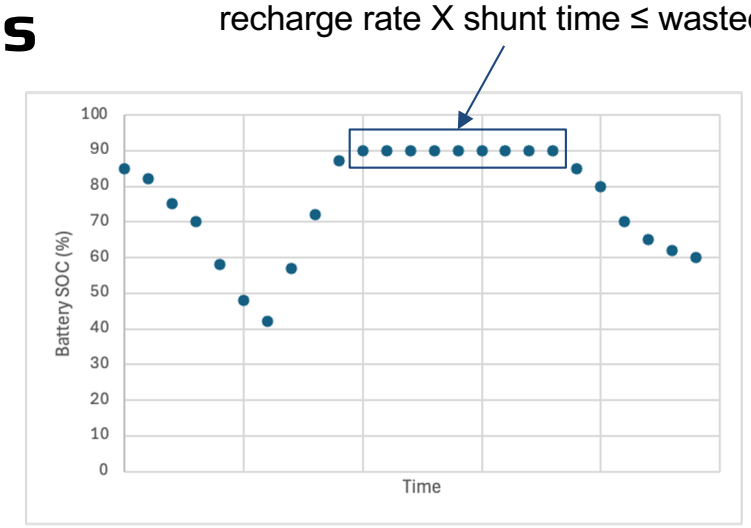

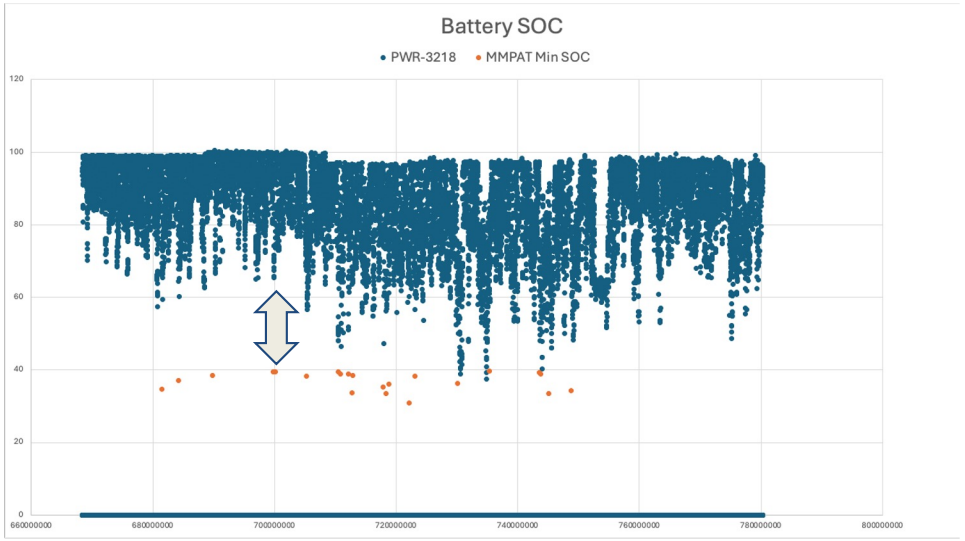

MARS 2020 PERSEVERANCE

Rollout Scorecard

Constraint Based Planning deployed August 25th, 2023

On Board Planner deployed October 5th, 2023

Simple Planner used for **257** tactical planning cycles
(as of January 29, 2025)

Tactical Timeline is averaging 6.5 hours

To date, Mars 2020 has not reverted from Simple Planner
(besides Solar Conjunction and FSW update)

Simple Planner remains enabled for Sample Sols

Science sites arrived at faster thanks to additional drive distance

25 SP ISAs (15% of total M2020 ISAs),
none required Simple Planner stand-down

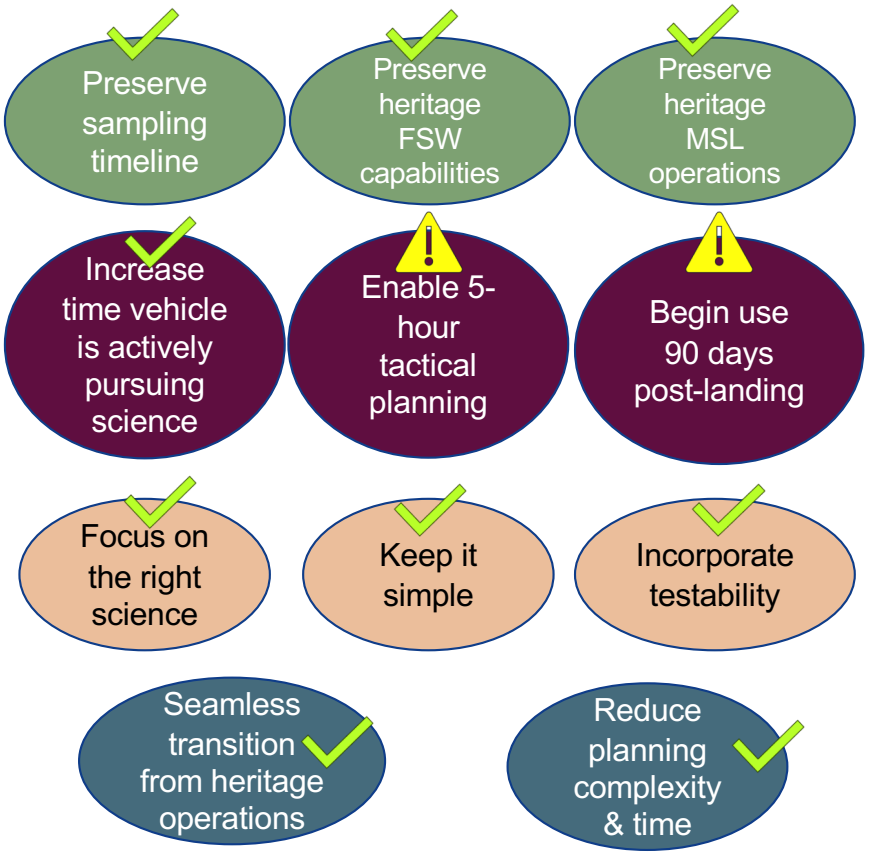

Closing Remarks

- Simple Planner has been the baseline for M2020 operations since rollout in October 2023
  - Anomalies have not required operations to revert to Master/Submaster paradigm
- As of 29 January 2025, SP has executed **257** plans covering **429** sols on Mars
  - **6917** onboard scheduling cycles
  - **7810** user activities executed
  - **13** km driven
  - **70,000**+ images acquired
  - **4** rock core samples acquired
- A deliberately incremental approach to deployment, combined with enthusiastic engagement from personnel from across the Mission, has enabled increased autonomy without short-term sacrifices to science or operability. This increased autonomy has made the rover significantly more efficient.

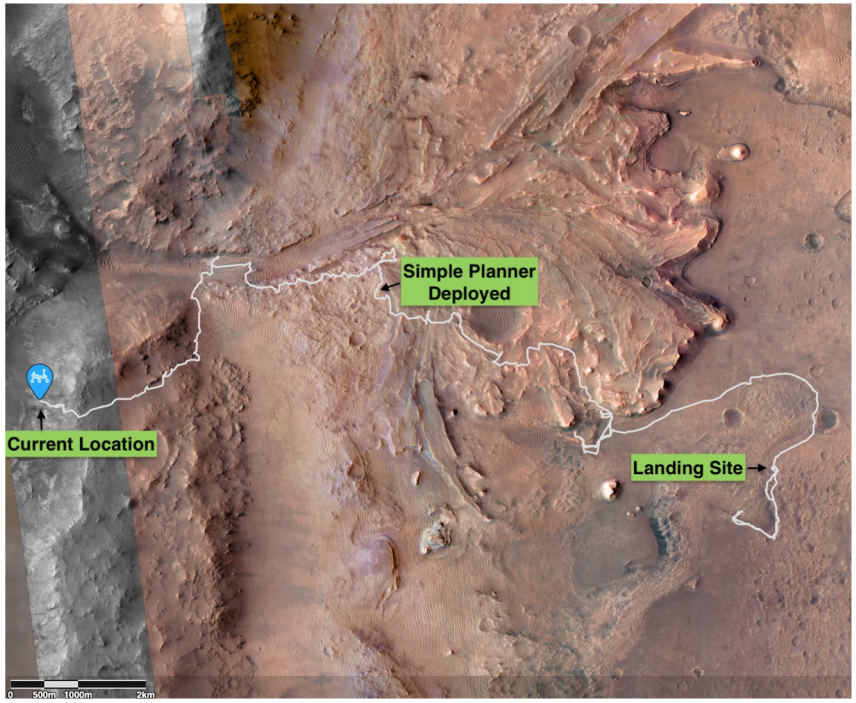

Jet Propulsion Laboratory
California Institute of Technology

jpl.nasa.gov