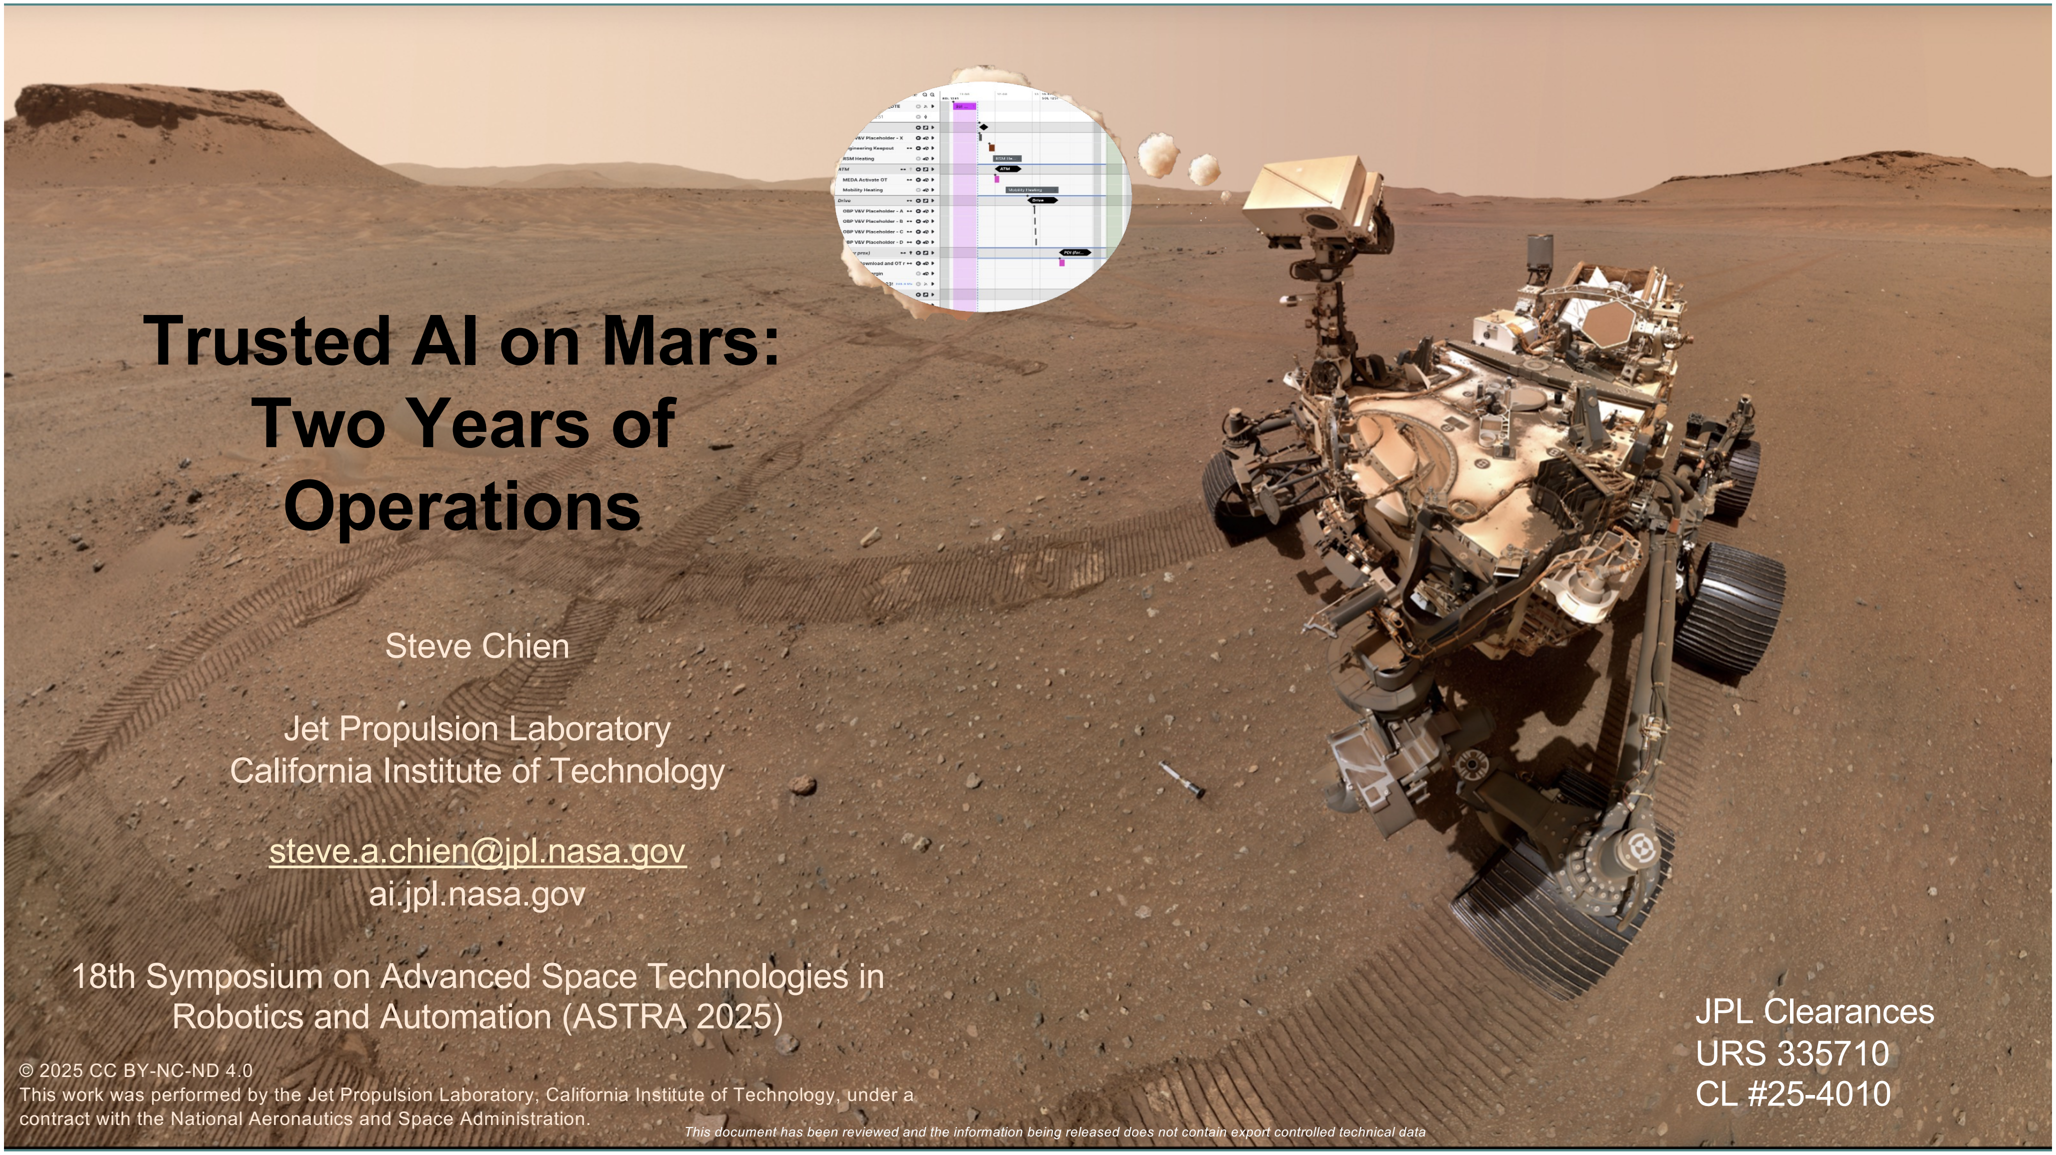

Trusted AI on Mars: Two Years of Operations

Steve Chien

Jet Propulsion Laboratory
California Institute of Technology

steve.a.chien@jpl.nasa.gov
ai.jpl.nasa.gov

18th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA 2025)

This work was performed by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

JPL Clearances
URS 335710
CL #25-4010

Introduction

Simple Planner is 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, with the Verification and Validation (V+V) test campaign beginning October of 2021, and deployment to operations October 5th, 2023.

Building and Deploying Trusted Autonomy is ***a full lifecycle process*** that begins with *conception*, continues through *design* and *prototyping*, product *build*, *testing*, *training*, and *deployment*.
[Potentially multiple iterations per multiple deliveries]

This talk describes how the Onboard Planner (OBP) moved through this process, from formulation, design, analysis and prototyping through testing as well as results from almost two years of operations.

- Search for signs of ancient life
  - Imagers, spectrometers, radar
- Collecting samples for potential return to Earth
  - Coring, regolith collection, sample handling
  - Prepare for human exploration
  - Weather, climate, radiation, dust
  - Oxygen generation

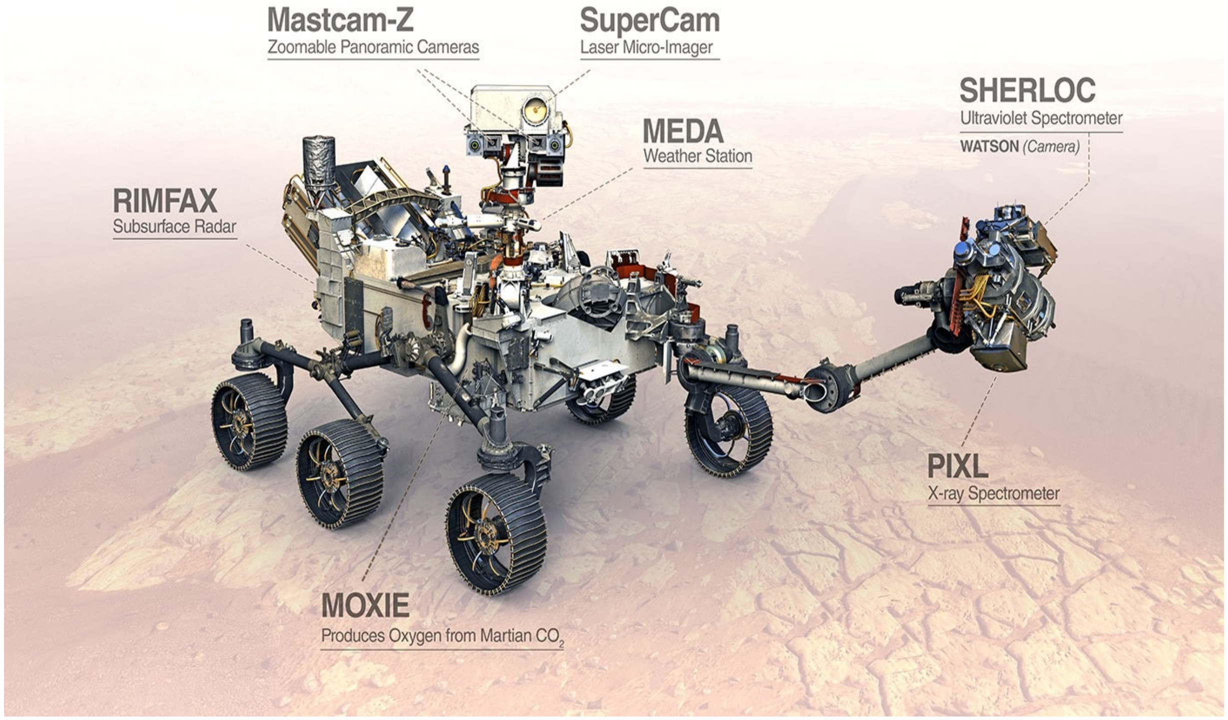

JPL Clearance CL#22-1884

Onboard Planner: Background

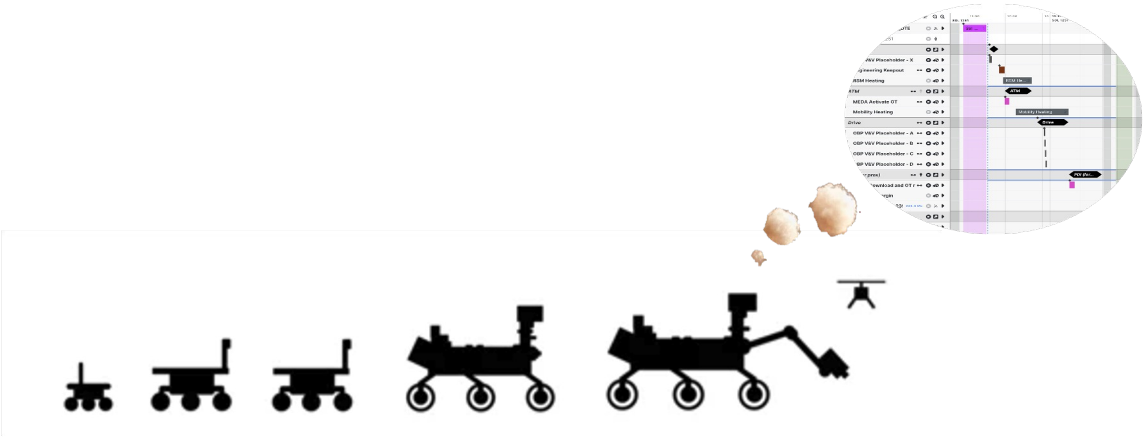

Productivity Challenge: Predicting Rover Resource Usage[8]

Activity resource estimation is difficult

Largely due to difficulty in **predicting activity duration and actual temperatures**

Resources: **time, energy, data**

Operations takes conservative approach

Typically overestimate and add margin

Can unnecessarily limit activity

Results in unused vehicle resources

Simple Planner approach

- **Move resource management on-board** by giving vehicle knowledge of energy state, pre-heat status, data volume, activity execution time ranges.
- **Enable more autonomy on vehicle** and let the vehicle **manage wake/sleep and energy usage.**

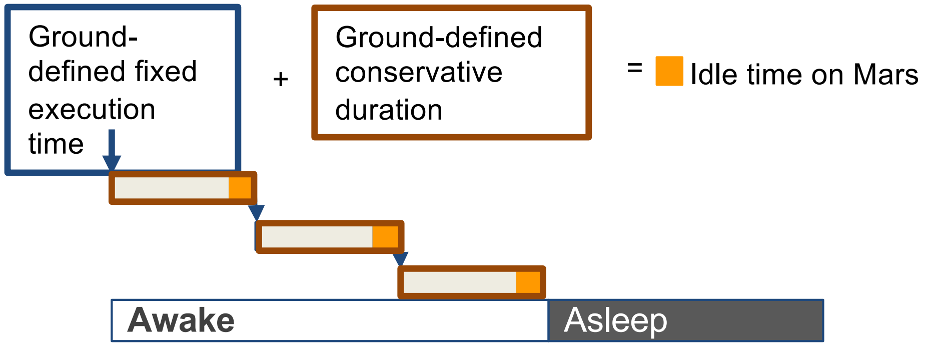

Ground-defined fixed execution time + Ground-defined conservative duration = 🟧 Idle time on Mars

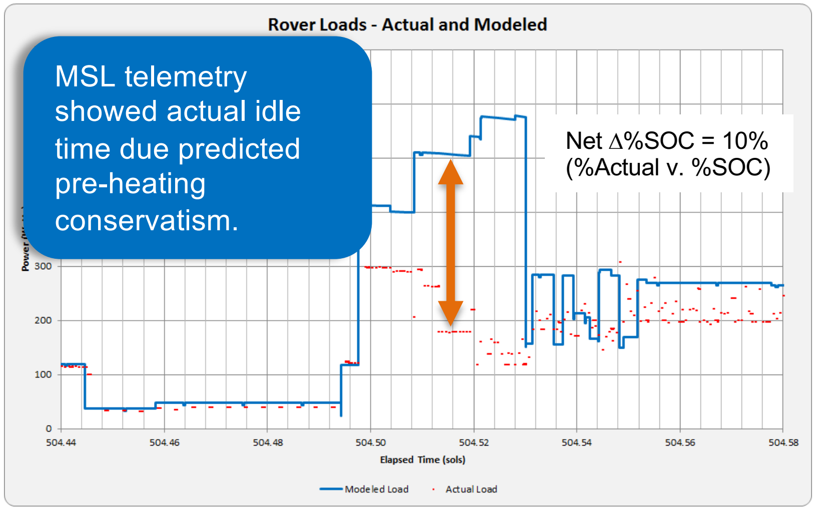

MSL telemetry showed actual idle time due predicted pre-heating conservatism.

Net Δ%SOC = 10% (%Actual v. %SOC)

JPL Clearance URS 330005 CL#24-6429

---

[8] Gaines, Doran, Justice, et al. (2016)

OBP Scheduling

Onboard Planner iteratively constructs Schedule

- Event driven scheduler invocation:
    - Plan activation
    - Wakeup
    - Activity vetoed, failed, or aborted
    - Activity ends earlier/later than scheduled time by **Δ**
    - Timer
- To minimize computation time, scheduler does not backtrack across activity placements
    - Ground awareness of search ordering for specific sols

See:

Gaines, D.; Rabideau, G.; Wong, V.; Kuhn, S.; Fosse, E.; and Chien, S. The Mars 2020 On-Board Planner: Balancing Performance and Computational Constraints. *In Flight Software Workshop*, February 2022.

Gaines, D.; Chien, S.; Rabideau, G.; Kuhn, S.; Wong, V.; Yelamanchili, A.; Towey, S.; Agrawal, J.; Chi, W.; Connell, A.; Davis, E.; and Lohr, C. Onboard Planning for the Mars 2020 Perseverance Rover. *In 16th Symposium on Advanced Space Technologies in Robotics and Automation,* June 2022.

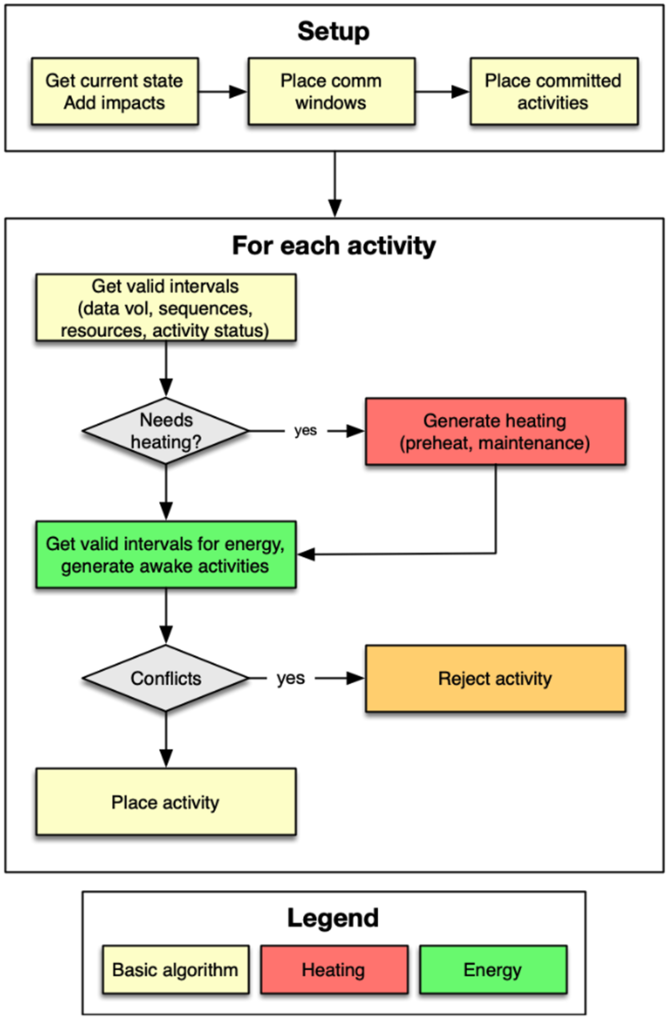

Setup

Get current state
Add impacts

Place comm windows

Place committed activities

For each activity

Get valid intervals
(data vol, sequences, resources, activity status)

Needs heating?

yes

Generate heating
(preheat, maintenance)

Get valid intervals for energy, generate awake activities

Conflicts

yes

Reject activity

Place activity

Legend

🟨 Basic algorithm

🟥 Heating

🟩 Energy

OBP Challenges - Thermal

OBP must implement thermal management for the Perseverance rover.

- Mechanisms onboard the Perseverance rover must be at a safe temperature before being used (such as to drive motors, arm motors, or instruments).
- Therefore OBP needs to schedule preheat and maintenance heating to support these activities.
- Preheat lead time (and energy consumption) can be affected by time of day and shared costs with other activities.

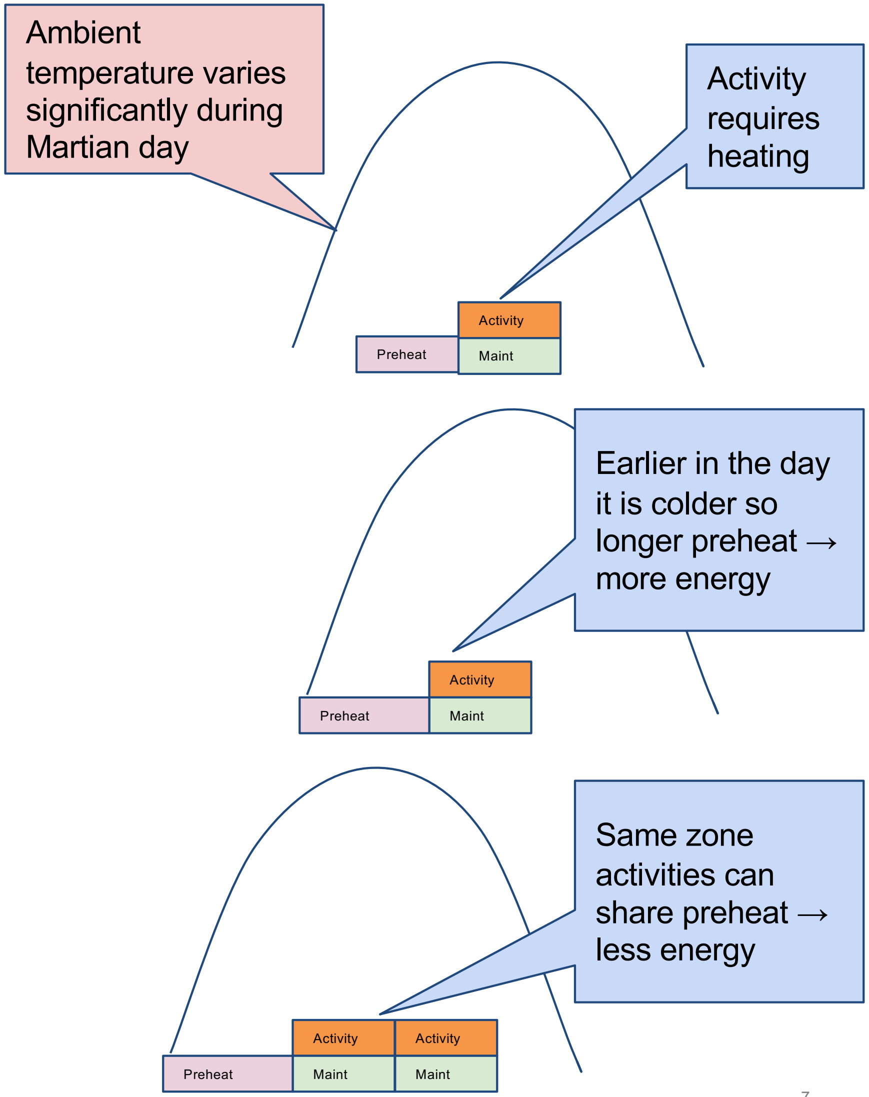

Ambient temperature varies significantly during Martian day

Activity requires heating

Earlier in the day it is colder so longer preheat → more energy

Same zone activities can share preheat → less energy

OBP Challenges - Energy

The Perseverance rover is energy constrained

- The rover just being awake → power negative state
- Most activities require that the rover (RCE) be awake in order to be performed
- Common pattern of operations is that the rover takes naps throughout the martian day, or *sol*, in order to conserve energy
- OBP must manage the wake sleep schedule of the rover

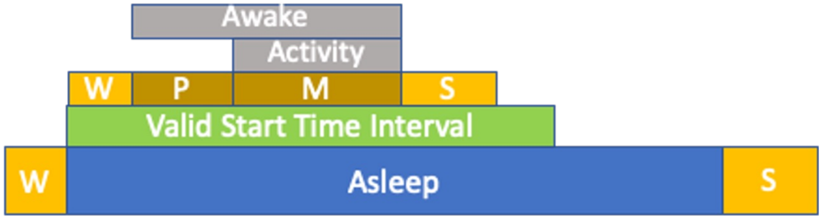

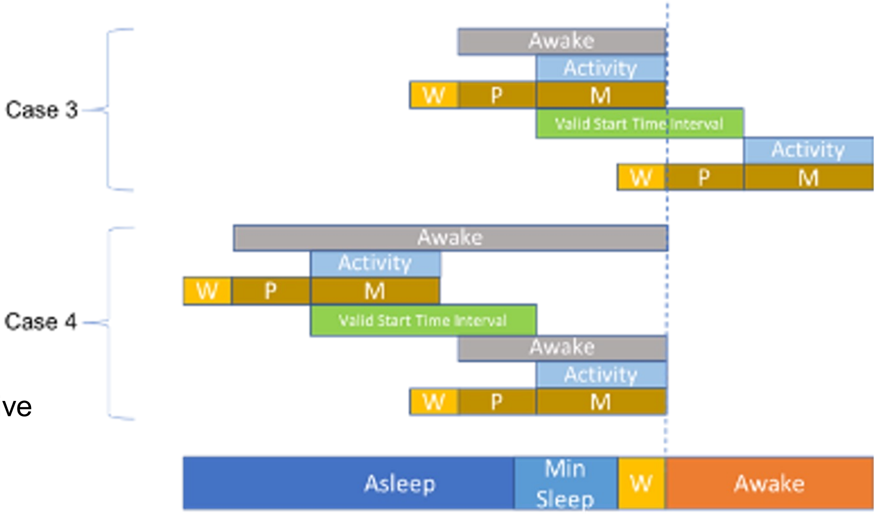

See:
W. Chi, S.Chien, and J. Agrawal. Scheduling with complex consumptive resources for a planetary rover. In International Conference on Automated Planning and Scheduling (ICAPS 2020), Nancy, France, October 2020.

Developing Trusted Autonomy: SW Development - Analyze (Informal Methods)

Trusted M2020 OBP – Informal Methods Computational Model of Scheduler Runtime

- Developed an in-depth computational complexity model for onboard planner runtime (Chi, Chien, et al. 2017).

$$((4d + 4dp) + 2*(4d^2 + 4d^{2p}))n +$$
$$(4r^2 + T + (4+4p)^2 + (16 + 32p + 16p^2) + 2*(32d + 64dp +32dp^2) + 2)n^2$$
$$+ (8 + 2*(64 + 192p + 192p^2 + 64p^3))n^3$$

where
n = Number of Activities (excluding preheats, maintenance, awakes)
d = Number of discrete intervals
T = Number of Timelines
P = Total Number of preheats

- In part driven by the complexity model - the root finding algorithm was eliminated in favor of inverting a preheat table, producing constant time and guaranteeing termination
- The complexity model highlights that cumulative rate timelines are dominate the other aspects of the algorithm at $O(n^2)$.
- With this knowledge, there is a more complex algorithm that would speed up the cumulative rate timeline to $O(n)$
- Termination is proved by showing that each portion of the algorithm is bounded in input.

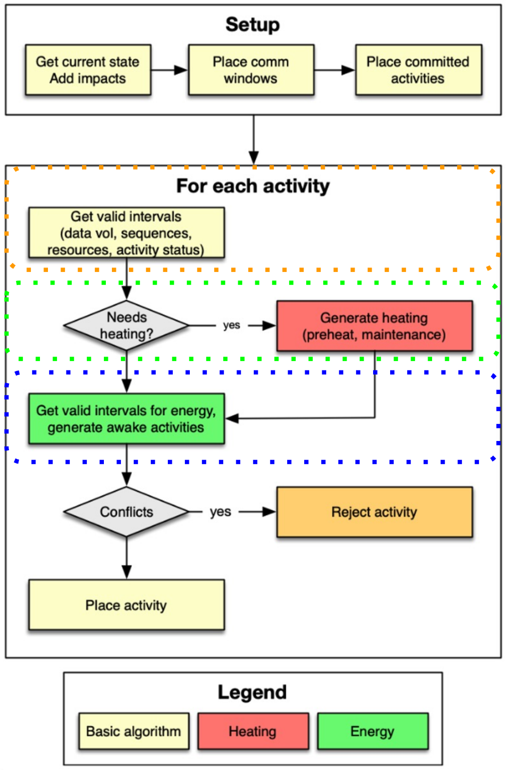

Trusted M2020 OBP – Informal Methods
Empirical Model of Scheduler runtime

- Developed an empirical model of onboard planner runtime*.
  - Uses parameters from scheduler call to estimate runtime
    - # of activities, preheat estimates, last invocation runtime, …
- Said model enables adaptive setting of the commit window to enable more effective rescheduling
  - However, gains were deemed not worth the added FSW complexity (incl. V+V) so the OBP as flown uses a static commit window
  - FE enables pulling forward activities when OBP completes rescheduling early

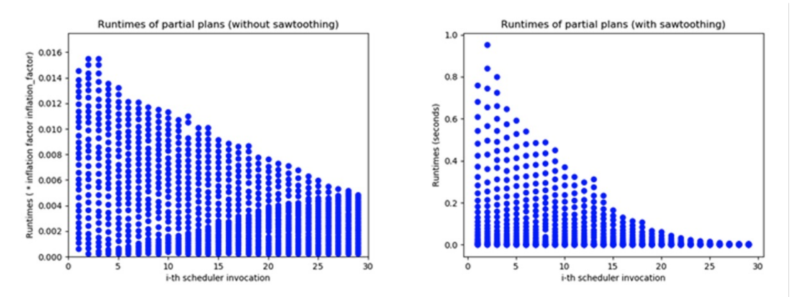

Runtimes decrease later in the plan with fewer activities to schedule

Linear feature-based predictive model (zoom in at right)

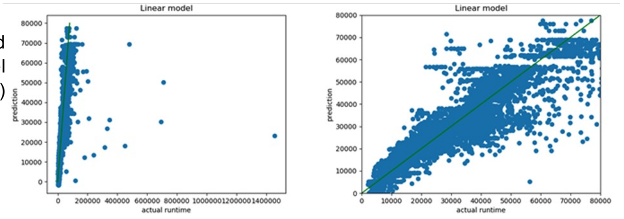

Prior invocation runtime model (zoom in at right)

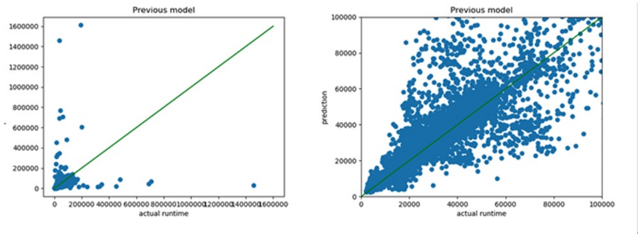

* - Bhaskaran, S.; Agrawal, J.; and Chien, S. Using a Model of Scheduler Runtime to Improve the Effectiveness of Scheduling Embedded in Execution. In Workshop on Integrated Execution (IntEx) / Goal Reasoning (GR), International Conference on Automated Planning and Scheduling (ICAPS IntEx/GP 2020), October 2020.

Trusted M2020 OBP – Informal Methods
Analysis of FE and Scheduling Invocation methods

- Developed an analytical model scheduler invocation and FE makespan gains [Agrawal et al. 2021]
    - This analysis supported the design decision for event driven rescheduling
    - This analysis supported some aspects of FE design

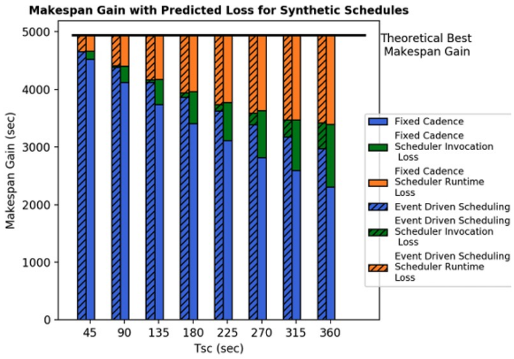

(b) Synthetic schedules are entirely serial: Makespan gain + predicted loss = theoretical best makespan gain.

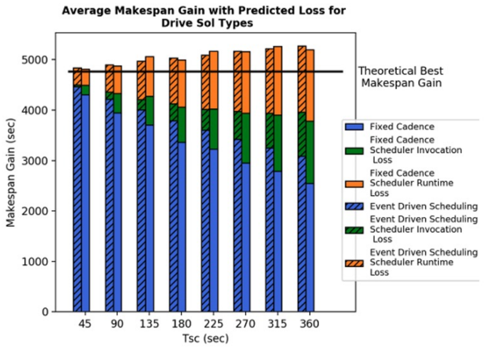

(b) Makespan gain + predicted loss ≠ theoretical best makespan gain (if we had an instantaneous scheduler). Real schedules are not all perfectly serial and have other execution constraints (e.g. preheats).

See:

Agrawal, J.; Chi, W.; Chien, S. A.; Rabideau, G.; Gaines, D.; and Kuhn, S. Analyzing the Effectiveness of Rescheduling and Flexible Execution Methods to Address Uncertainty in Execution Duration for a Planetary Rover. *Robotics and Autonomous Systems*, 140 (2021) 103758. 2021.

Trusted M2020 OBP – Informal Methods
Analysis of Energy Scheduling Algorithms

- Analyzed three competing methods of implementing energy constraints in scheduling [Chi et al. 2020]:
    - Max Duration
    - Probe
    - Linear
- This analysis was used to support the design decision to implement the heuristic probe algorithm for energy scheduling in the OBP

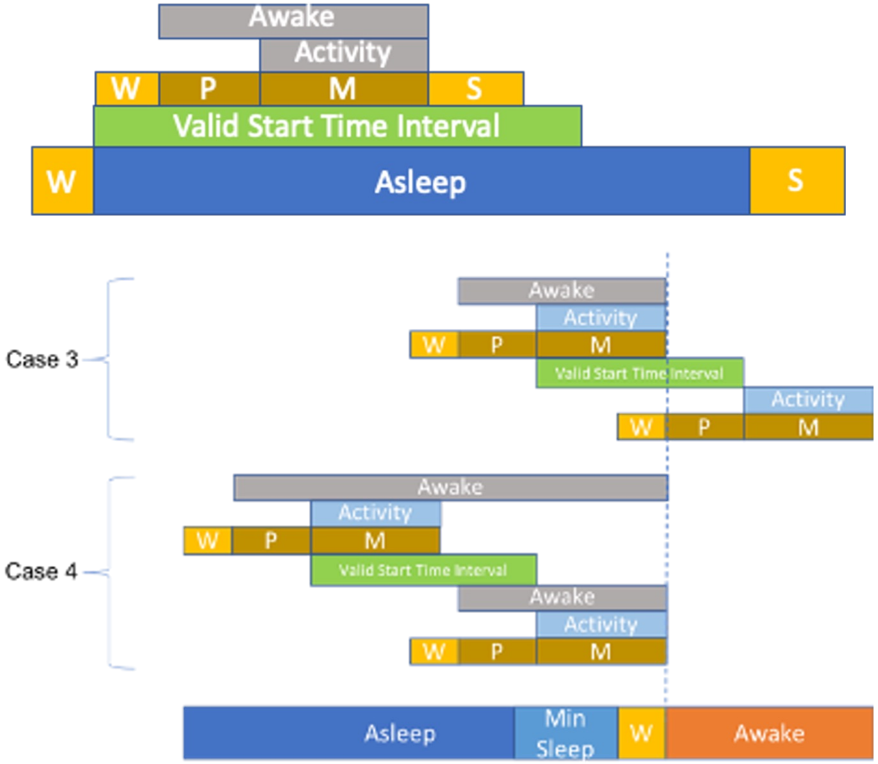

* - Chi, W.; S.Chien; and Agrawal, J. Scheduling with Complex Consumptive Resources for a Planetary Rover. In *International Conference on Automated Planning and Scheduling (ICAPS 2020)*, Nancy, France, October 2020.

Developing Trusted Autonomy:
SW Development - Build

M2020 FSW Processes → OBP

- Informal methods in software development includes:
    - code walkthroughs,
    - coding guidelines,
    - coding rules (see MSL process*)
    - design reviews, and
    - software documentation.

- Formal methods in software development includes:
    - use of static code analyzers (CodeSonar) as part of the M2020 software development process
    - runtime analyzers as part of unit test: Valgrind (memory, performance), AddressSanitizer (Memory)

- Testing
    - Unit test coverage analysis via Gcov

* Gerard J. Holzmann, “Mars Code,” Communications of the ACM, February 2014, Vol. 57 No. 2, Pages 64-73
10.1145/2560217.2560218

Developing Trusted Autonomy: Test Campaign

M2020 OBP Testing - Test Challenges

Challenge: Ensure system is and *will remain* safe in dynamic environment with OBP-in-control

- Complex **interfaces** with critical health and safety subsystems
    - Sequence control
    - Wakeup/shutdown control
    - Thermal control
    - Battery SOC estimator
    - Comm window manager
    - System fault protection

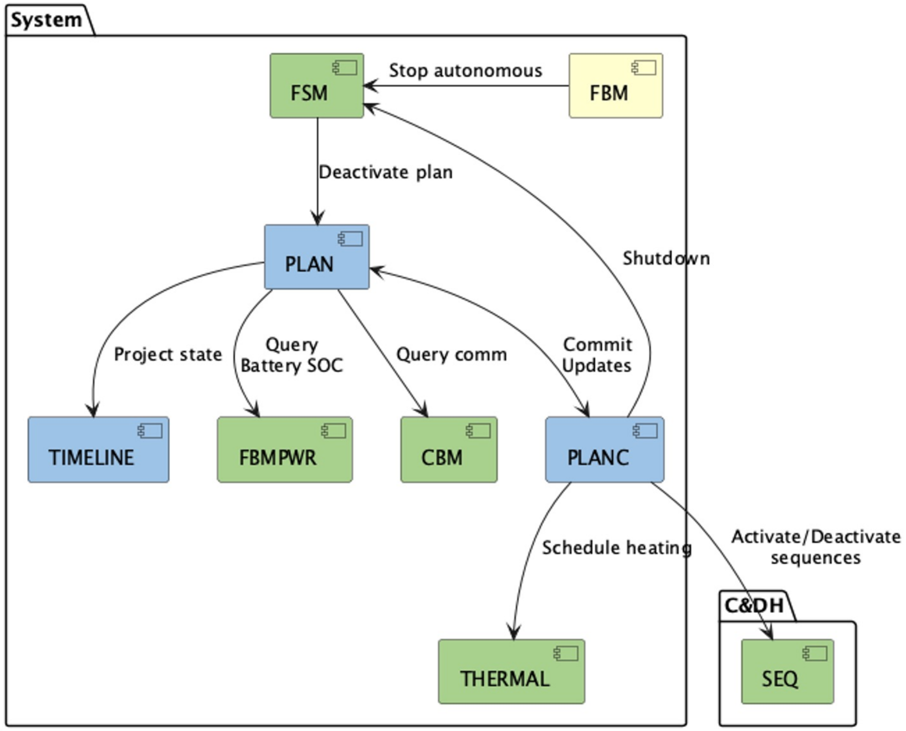

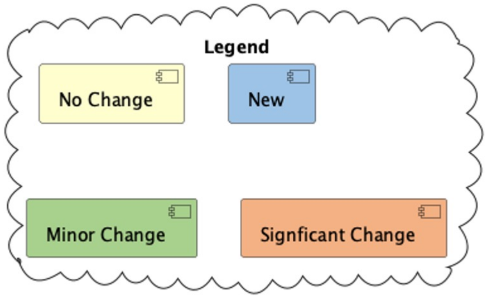

M2020 OBP Testing - Test Challenges

Challenge: Flexible scheduling/execution can lead to highly variable outcomes

- Schedule structure **sensitive to variation**
    - Activity duration
    - Activity completion status
    - Hardware temperatures
    - Energy available
    - Time available
- Periodic rescheduling creates many **branch points**
    - Schedule revised hourly while awake
    - Execution deviations trigger scheduling
- Numerous possible execution outcomes
    - Need to constrain to finite number of tests

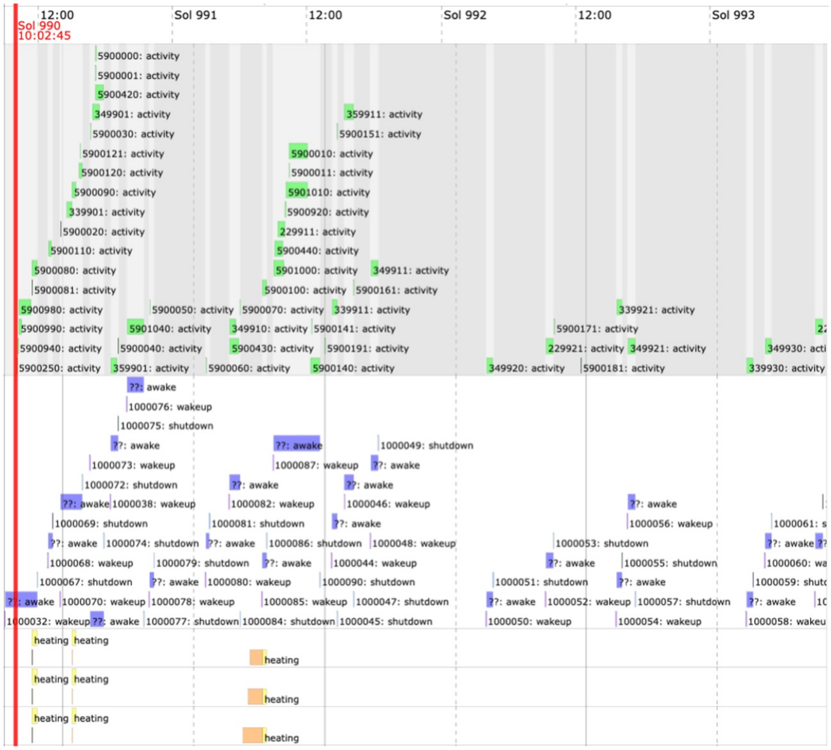

M2020 OBP Testing - Test Campaign Overview

Opportunity: Deploy OBP during stable surface operations campaign

- Test campaign performed after ~500 sols of Master/Submaster operations
- Extensive Mars surface operations experience within V+V team
- Regular interfacing with mission system through Simple Planner Working Group

Approach: Release capabilities in phases to allow parallel V+V and operations

Test Strategy: Verify specific capabilities in flight-like scenarios where practical

- Primary focus: **Functional Testing** (bottom-up)
    - Targeted individual capabilities
    - Verified L3 and L5 requirements and SWG (software guidance) artifacts
- Secondary focus: **Scenario Testing** (top-down)
    - Transformed as-flown Master/Submaster plans into OBP plans

Trusted M2020 OBP - Beyond the Test Campaign

Science Walkthroughs (weekly 2/2023-5/2023)

Super Thread Tests:

- Performed on *VSTB* July and November of 2022
- Partial validation of OBP use in operations

First-Time Activities (FTA):

- Performed on *Perseverance* May - July 2023
- 3 toe-dips, demonstrating basic capabilities

Operational Readiness Test (ORT):

- Performed on *VSTB* August of 2023
- Final validation of OBP use in operations

Flight School:
- Seven (7) sessions
- Over 10 hours of material
- Over 150 staff had completed prior to OBP operations

Discussion

- Onboard Planner represents a *paradigm shift* in that the flight system has considerable latitude to alter execution

- OBP team embraced *full life cycle* development of trusted autonomy
  - Informal methods
    - Termination and complexity analysis of scheduler; empirical analysis of runtime
    - Prototyping of alternative algorithms
    - Estimates of V+V effort for various algorithms → design decisions
  - Software practices - informal and formal methods
  - Continuous emphasis of going after biggest, lowest hanging fruit of productivity gains

- OBP testing campaign proved critical to infusing autonomy for a flagship mission
  - Heavy reliance on functional tests informed by flight use-cases
  - Less reliance on scenario-based testing than EO-1 and IPEX
  - Leveraged hundreds of actual plans run on Mars

- Very few anomalies (Flight Software or Operations) have been encountered in flight, with zero returns to MSM (except for conjunction), validating the trust that operators have placed in this new system

Lessons Learned – Plan Analysis – Ground Ops Vulnerability Aanalysis

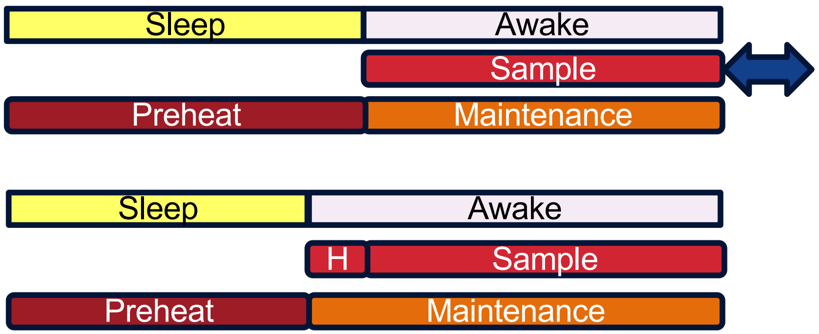
Sleep
Awake
Sample
Preheat
Maintenance

Sleep
Awake
H
Sample
Preheat
Maintenance

IF Heated activity < 15 minutes temporal flexibility

Cocpit adds helper activity starting 20m earlier
(energy cost = woke up earlier CPU energy
+ 20 m heater maintenance,
+ plus earlier heat may be colder)

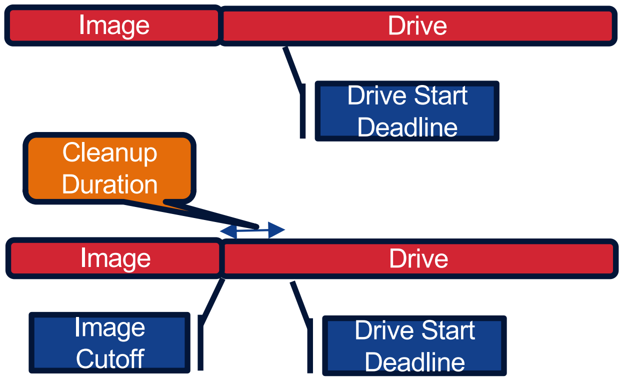
Image
Drive
Drive Start Deadline
Cleanup Duration
Image
Drive
Image Cutoff
Drive Start Deadline

IF need to protect later activity from earlier activity running long (e.g. unit resource conflict)

MSM commits to grounded time to provide margin

OBP reduces margin; and can use temporal flexibility if needed

If you had a risk model for duration, you could warn for brittleness

Example: drive start must not be pushed too late

overrun protection: assert cutoff for prior activities to prevent drive from missing last start from prior activity push

MARS 2020
PERSEVERANCE

Lessons Learned – Improved Science

- The SHERLOC instrument has a minimum thermal state and an ideal thermal state for optimal science operations
- In Master SubMaster (MSM) operations activities are driven by time; to ensure minimum thermal state

  conservative temperature predictions + conservative preheat durations

  → often SHERLOC is operated at beyond ideal thermal state

  (in addition to the energy waste from un-needed heating or keeping warm)

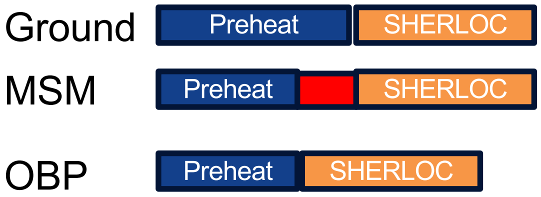

Ground Preheat SHERLOC

MSM Preheat SHERLOC

OBP Preheat SHERLOC

- In OBP operations,

  preheat can be started later if the rover senses that it is warmer that predicted

  If the preheat completes early, the observation can occur earlier
  (if designated temporally flexible)

  → This can result in measurements under better thermal state and therefore better science

Lessons Learned – Improved Science

- OBP can save energy when execution varies from projection (usually savings due to conservative ground projections)
- Often this saved energy can be carried over to subsequent sols and used for additional science; however limited battery capacity means this is not guaranteed
- Using this energy in the same sol can obviate this difficulty
  - Ground can specify "optional" activities that execute only if resources are available (usually from earlier savings)

- A recent informal analysis of Sols 1319-present (300+ sols) showed that 96% of optional activities during this time range executed successfully. These optionals represent ~ +7.7% addition to direct activities

- A more direct and valuable OBP contribution to science is by compressing the Sol path (e.g. by completing execution with more energy remaining, "recharge" sols are avoided and the primary science chain (often across many sols) executes more rapidly

Science Discoveries under OBP

- NASA Says Mars Rover Discovered Potential Biosignature Last Year
  https://www.nasa.gov/news-release/nasa-says-mars-rover-discovered-potential-biosignature-last-year/

- A sample collected by NASA's Perseverance Mars rover from an ancient dry riverbed in Jezero Crater could preserve evidence of ancient microbial life. Taken from a rock named "Cheyava Falls" last year, the sample, called "Sapphire Canyon," contains potential biosignatures, according to a paper published Wednesday in the journal Nature.
- The Sapphire Canyon sample was acquired on Sol 1215, which was part of a 3 Sol OBP plan covering Sols 1214-1216. The data from Sols 1178, 1188, 1201, 1202, and 1217 supporting the sampling discovery were all taken under OBP control.

NASA's Perseverance rover discovered leopard spots on a reddish rock nicknamed "Cheyava Falls" in Mars' Jezero Crater in July 2024. Scientists think the spots may indicate that, billions of years ago, the chemical reactions in this rock could have supported microbial life; other explanations are being considered.

Credit: NASA/JPL-Caltech/MSSS

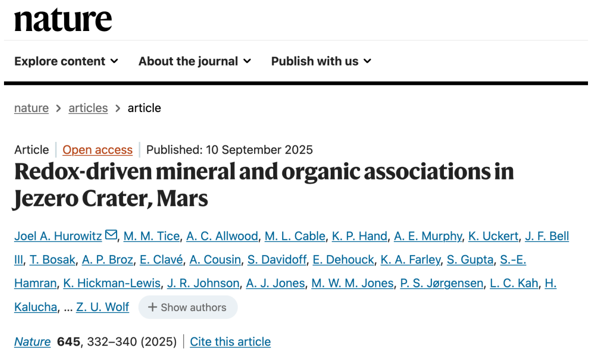

Article | Open access | Published: 10 September 2025

Redox-driven mineral and organic associations in Jezero Crater, Mars

Joel A. Hurowitz, M. M. Tice, A. C. Allwood, M. L. Cable, K. P. Hand, A. E. Murphy, K. Uckert, J. F. Bell III, T. Bosak, A. P. Broz, E. Clavé, A. Cousin, S. Davidoff, E. Dehouck, K. A. Farley, S. Gupta, S.-E. Hamran, K. Hickman-Lewis, J. R. Johnson, A. J. Jones, M. W. M. Jones, P. S. Jørgensen, L. C. Kah, H. Kalucha, ... Z. U. Wolf

Nature **645**, 332–340 (2025) | Cite this article

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 25 September 2025, OBP has executed **384** plans covering **656** sols on Mars:
  - **10,688** onboard scheduling cycles
  - **11,864** user activities executed
  - **17.4** km driven
  - **105,000**+ images acquired
  - **8** rock core samples acquired

- The M2020 Simple Planner approach to Trusted Autonomy has been overwhelmingly successful!

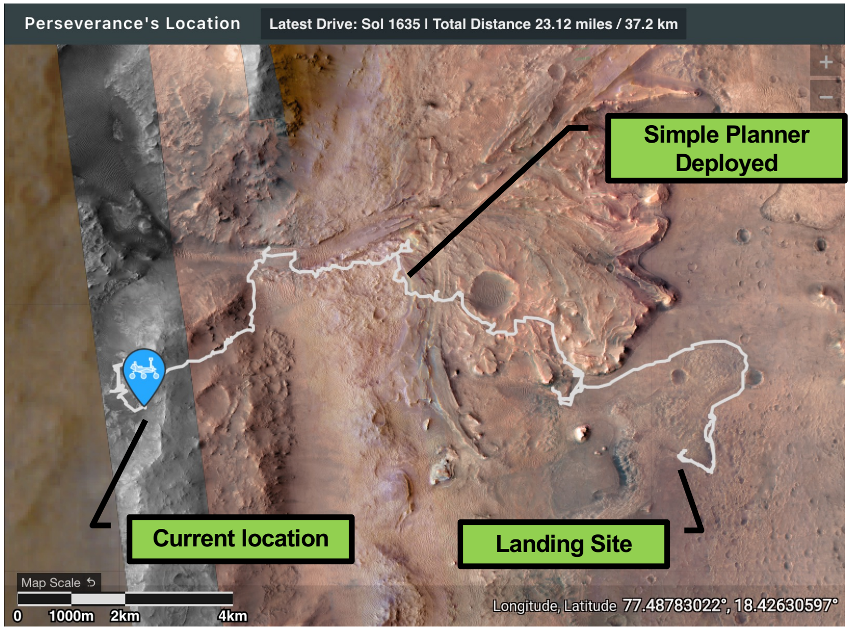

Map generated 29 Sep 2025.

Onboard Planner

Gaines et al. 2023 ASTRA
Verma et al. 2023 Sci Rob

AEGIS Autonomous Targeting

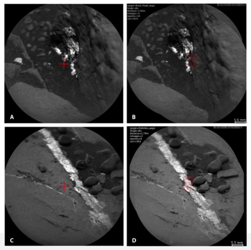

Francis et al. Sci Rob 2017

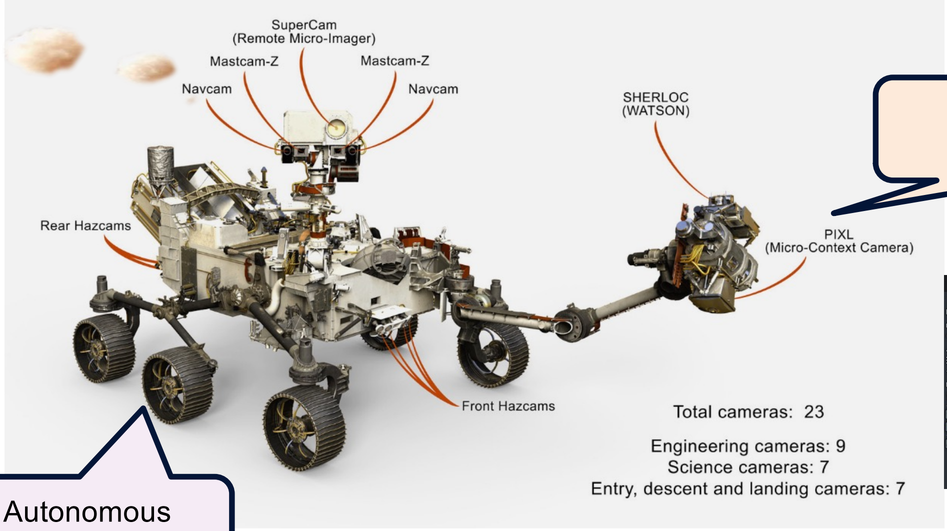

Total cameras: 23

Engineering cameras: 9
Science cameras: 7
Entry, descent and landing cameras: 7

PIXL Adaptive Sampling

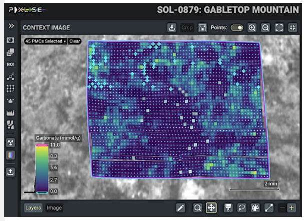

Lawson et al. 2025 Icarus

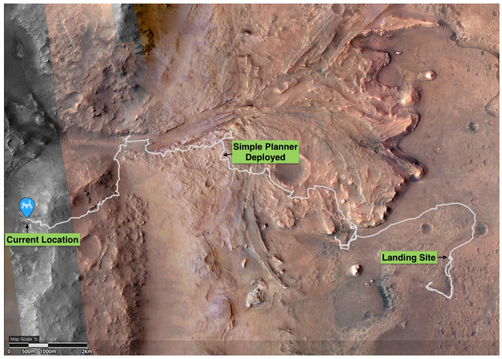

Enav Autonomous Driving

Verma et al. 2023 Sci Rob

Acknowledgements

OBP FSW Team

- Dan Gaines
- Gregg Rabideau
- Eddie Benowitz
- Vincent Wong
- Amruta Yelamanchili

Funding

- Mars 2020 Project
- Europa Lander Feed-Forward
- 5x R&TD
- Directors Innovation Fund
- OTIS

OBP V+V Team

- Dan Gaines
- Stephen Kuhn
- Elyse Moffi
- Shreya Parjan
- Kevin Reich
- Ansel Rothstein-Dowden
- Danny Tran
- Nick Waldram
- Sean Wenzel

5x R&TD Team

- Steve Chien
- Wayne Chi
- Jagriti Agrawal
- Sarah Bhaskaran

M2020 Operations Team

Original SP Element Lead: Elyse Moffi
SP FSW Lead: Dan Gaines
SP Flight Systems Engineer: Stephen Kuhn
Deployment SP Element Lead: Bekah Siegfriedt
Current SP Element Lead: Nick Waldram
Operations for SP: James Hazelrig and James Biehl
SP Science: Raymond Francis

What’s Next?

URS332197
CL#25-1222

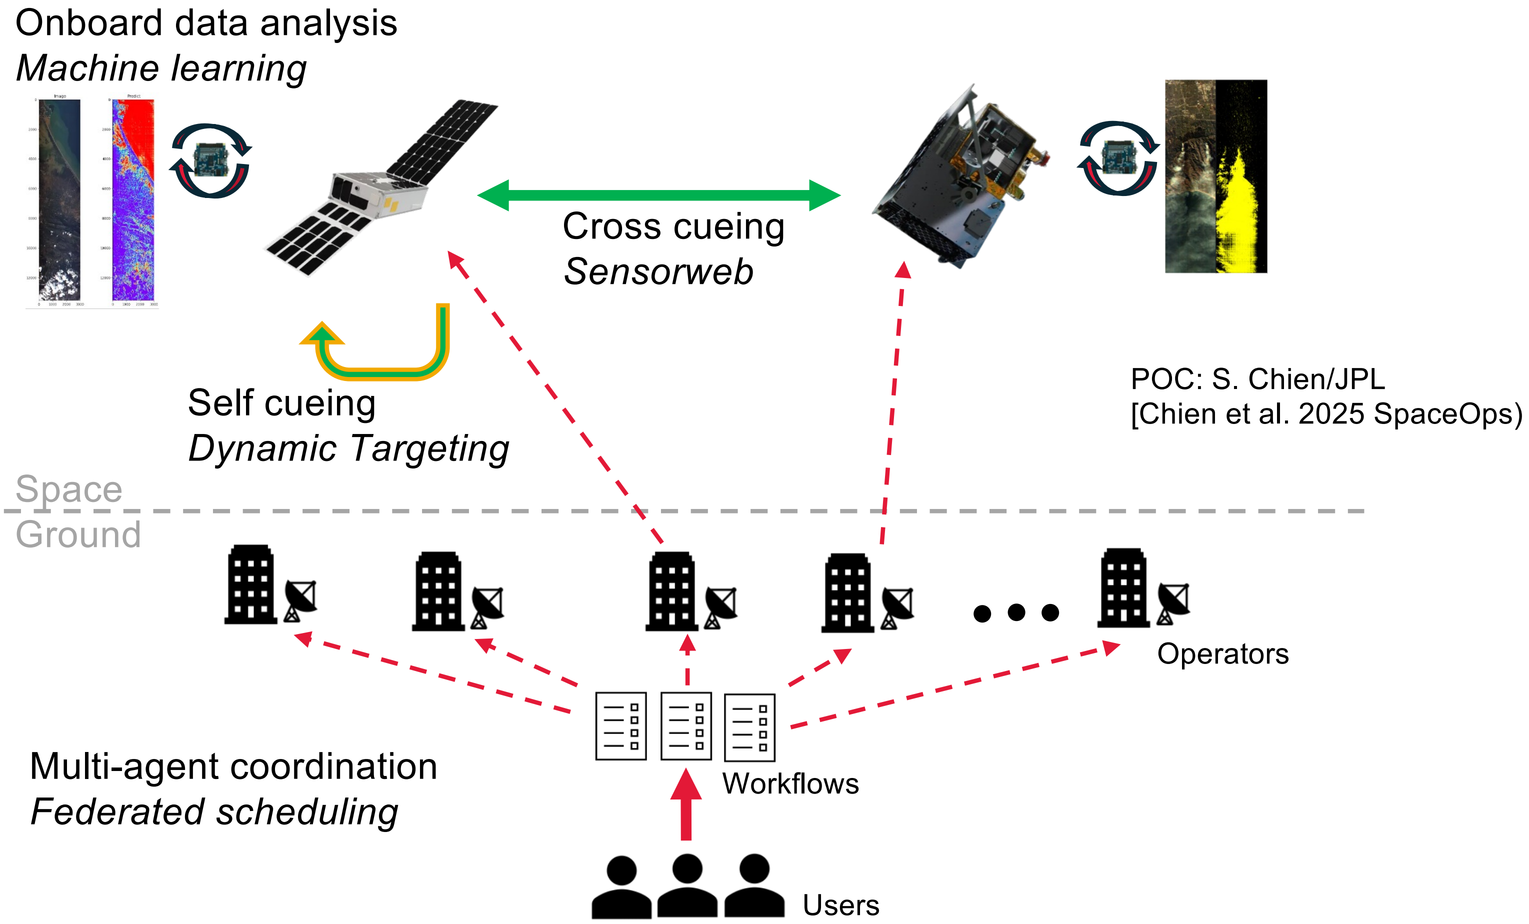

Onboard data analysis
Machine learning

Cross cueing
Sensorweb

Self cueing
Dynamic Targeting

POC: S. Chien/JPL
[Chien et al. 2025 SpaceOps)

Multi-agent coordination
Federated scheduling

FAME in Space Demonstration

First flights – Fall 2024. Up to 50 assets included. Duration: 3 years (planned).

URS324900
CL#24-2153

AI is Central to Ambitious Space Mission Concepts to *Find Life Beyond Earth*

Europa Lander Mission Concept
Matanuska Glacier Field Trials (2022)
Decision Theoretic Planning!
[Chien et al. 2022 ASTRA, Wagner et al. 2023 JAIS, Russino et al. 2023 JAIS]
Robotics – Sampling Autonomy: [Bowkett et al. 2023 IEEE Aero]

JWST CO_2 @ Europa:
[Villanueva et al. 2023 Science]
[Trumbo et al. 2023 Science]

URS324900
CL#24-2153

AI is Central to Ambitious Space Mission Concepts to *Find Life Beyond Earth*

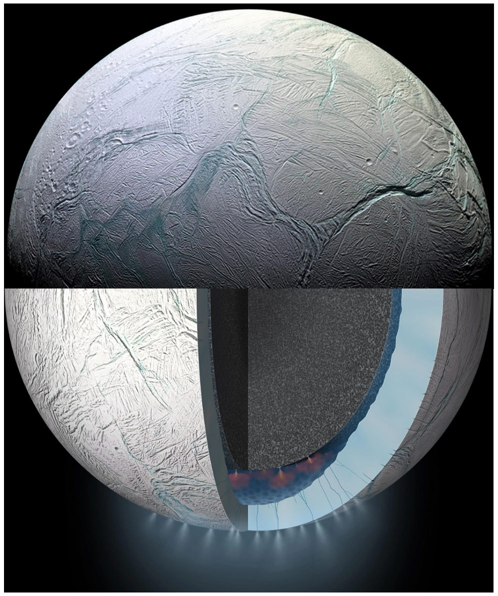

Enceladus CO_2 [Brown et al 2006 Science]
Enceladus Phosphates [Postberg et al. 2023 Nature]

Exobiology Extant Life Surveyor (EELS)
[Vaquero et al. 2024 Science Robotics]

Multiple tests on Athabasca Glacier (2022, 2023)

Jet Propulsion Laboratory
California Institute of Technology

---

jpl.nasa.gov