

## Scheduling NASA's Deep Space Network in the Artemis Era

Mark D. Johnston

*Jet Propulsion Laboratory/California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California 91109*

### Abstract

Scheduling NASA's Deep Space Network (DSN) is inherently challenging under nominal circumstances. However, the additional stress of supporting the Artemis 1 mission, along with the expectations for future flights in the Artemis series, introduces another level of complexity and demand. The Artemis 1 mission was accompanied by a number of smallsats, which further complicated the task of providing communications and navigation support to all the launching missions. Even if future missions do not include DSN-supported smallsats, the primary Artemis scheduling support remains complicated by uncertainties in the launch date and the exact launch time within a specified window — a window during which the support geometry can change significantly. This complexity is further compounded by post-launch trajectory updates, and by the need to balance any impacts on other DSN users, including those who could potentially be seriously impacted by last-minute schedule changes. This paper outlines the nominal DSN scheduling process and tools, and then details how the Artemis 1 mission requirements (and the accompanying smallsats) impacted that nominal process. Lessons learned from this experience have been collected and converted into a set of potential future improvements designed to both lessen the impact and streamline future support timelines. In this paper we will describe the rationale behind the proposed improvements, how they will better support Artemis missions, and the architecture and design considerations, specific to the DSN, that come into play.

**Keywords:** Deep Space Network; scheduling; heuristic search

### 1. Introduction

NASA's Deep Space Network (DSN) is a critical infrastructure for space exploration, serving as the primary communication and navigation backbone for dozens of interplanetary missions. Consisting of large 34-meter and

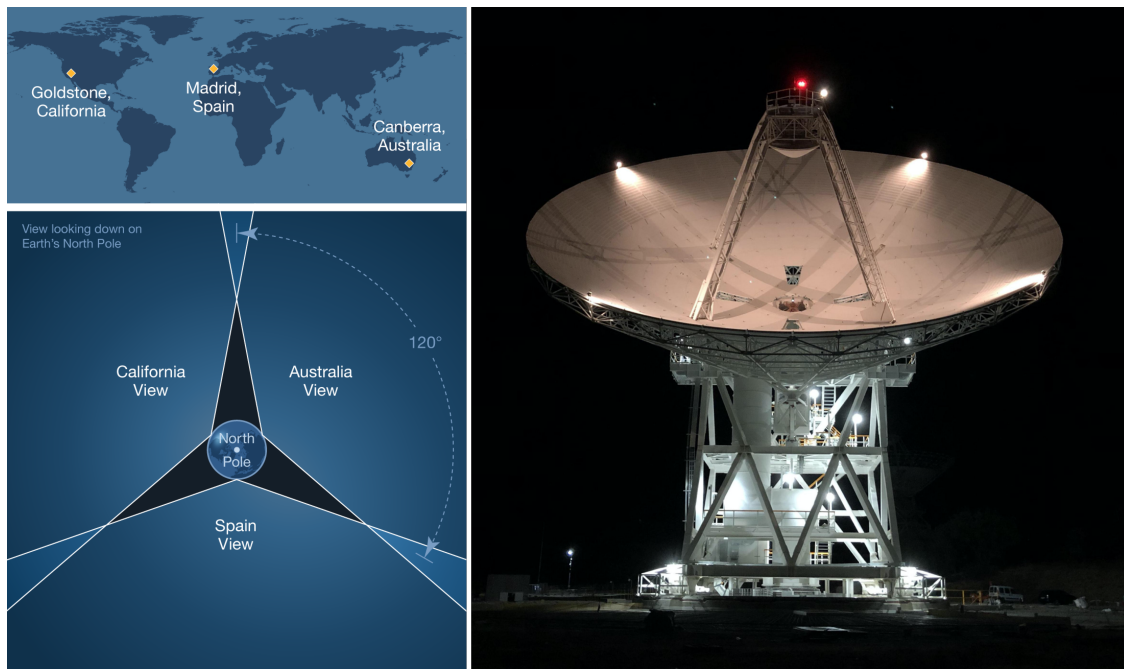


Fig. 1 The Deep Space Network station locations (upper left), approximate fields of view showing how a spacecraft is nearly always visible from at least one station (lower left) [1], and the newest DSN antenna, DSS-56, which entered service in January 2021 in Madrid, Spain (right).

70-meter antennas located at three globally distributed sites — Goldstone, California (USA); Madrid, Spain; and Canberra, Australia — the DSN ensures continuous coverage and connectivity for spacecraft traveling beyond Earth's orbit. It is the largest and most sensitive telecommunications network in the world, enabling not only mission operations but also scientific research in radio astronomy and planetary radar.

The DSN currently supports approximately 50 spacecraft missions using a total of 14 operational antennas, and overall demand is expected to double within the next decade. Operated by the Jet Propulsion Laboratory (JPL) in Pasadena, California, the DSN's scheduling process is a complex and highly constrained endeavor. Missions require a stable and well-defined schedule weeks to months in advance to plan their onboard sequencing of activities, ensuring that crucial data downloads and command uplinks align with spacecraft operational constraints — especially when long light-travel-time delays are involved.

### *1.1. The Artemis Program: A New Era of Lunar Exploration*

The Artemis program is NASA's ambitious effort to return humans to the Moon and establish a sustainable presence, ultimately paving the way for future human exploration of Mars. Unlike Apollo, which was a series of short-duration missions, Artemis is designed as a multi-phase program that will build an enduring lunar infrastructure. Each mission in the series brings increasing complexity and DSN demands, making efficient scheduling a mission-critical factor.

- Artemis 1 (2022) was an uncrewed test flight that validated NASA's Space Launch System (SLS) and Orion spacecraft. It successfully completed a lunar orbit and returned to Earth, testing propulsion, navigation, and re-entry systems. This mission set the stage for crewed deep-space flights but also exposed key challenges in DSN scheduling, particularly regarding uncertain launch windows and frequent last minute updates to the schedule.
- Artemis 2 (planned for 2026) will be the first crewed flight, carrying astronauts on a lunar flyby to demonstrate Orion's life support and deep-space operational capabilities. This mission is shorter in duration, but will introduce even stricter DSN scheduling requirements, as real-time voice, telemetry, and command support will be essential for crew safety.
- Artemis 3 and later missions aim to land astronauts on the lunar surface, focusing on the lunar south pole—a region of immense scientific interest due to its permanently shadowed ice deposits. This mission introduces a Human Landing System (HLS) to transfer crew from Orion to the Moon's surface and back, further increasing the operational complexity of DSN support.

### *1.2. DSN Scheduling in the Artemis Era*

The DSN scheduling process operates on a rolling weekly basis, typically 16 to 20 months ahead of execution. The first step involves integrating and deconflicting service requests from all users, submitted ahead of a common deadline. This phase is one of the most labor-intensive aspects of DSN scheduling. It takes several days of manual adjustments to resolve conflicts, followed by another week of negotiations among mission teams before a mutually agreed-upon baseline schedule is published.

With Artemis, this already complex system faces a new set of challenges:

- Preparing for potential launch date or time changes, sometimes with little notice, causes DSN schedulers to have to work on multiple potential timelines (contingency schedules) simultaneously, each representing a set of fully negotiated launch contingency tracking schedules
- In-flight changes to asset availability (such as equipment or antenna outages) may require rapid and last minute changes to the schedule
- Artemis 1 launched with a number of DSN-supported cubesats that further increased the complexity of DSN scheduling and preparation. While current plans do not call for cubesats tracked by DSN to be launched with Artemis 2, it remains a possibility for which the DSN has to be prepared.

As for many missions, the Artemis launch window is not instantaneous — there is an interval, typically a few hours in duration, such that launch may occur at any time within the window. Because of the changes in trajectory over the course of the launch window, the opportunities for contacting DSN antennas can vary significantly depending on exactly when launch occurs. In some cases, contact opportunities that exist at the start (open) of the launch windows move and shrink, and no longer exist at the end (close) of the window. The times and durations of the potential contacts also can

vary depending on the exact instant of launch. As a result, it is necessary to pre-plan for this variability in geometry over the course of the launch window, since there is not time once launch actually occurs to regenerate and reload all of the products necessary to reflect any changes.

The existing DSN scheduling process is significantly stressed by these Artemis support requirements, which affects the teams responsible for the work as well as the other DSN mission users. Improvements to the scheduling system and processes will allow the DSN to better adapt dynamically to last-minute changes, while ensuring minimum disruption to existing deep-space users.

### 1.3. Objectives of This Paper

This paper explores:

- The impact of Artemis 1 on DSN scheduling and lessons learned from its execution.
- Challenges expected in Artemis 2, 3, and beyond, including increased complexity that may arise from crewed missions and lunar surface operations.
- Proposed improvements to DSN scheduling, including automated deconfliction, AI-based decision support, and other performance and feature improvements
- The implementation status of new scheduling tools and methodologies, identifying completed improvements and remaining work to go.

## 2. Background and Context

NASA's Deep Space Network (DSN) has played a foundational role in space exploration for over 60 years, providing critical communications, navigation, and scientific support for deep-space missions. It has been the lifeline for dozens of interplanetary spacecraft, enabling the success of iconic missions such as Voyager, Mars rovers, and the New Horizons Pluto flyby. Historically, the DSN also played a pivotal role in the Apollo program, relaying communications and telemetry during humanity's first journeys to the Moon. Today, as NASA embarks on the Artemis program to establish a long-term human presence on the Moon, the DSN faces new scheduling challenges that will test its ability to efficiently allocate its limited resources.



Fig. 2 Artemis 1 launch as seen from the DSN Darkroom (Deep Space Operations Center) at JPL.

## 2.1. Overview of the DSN Scheduling Process

The DSN scheduling process is managed by the DSN using the Service Scheduling Software (SSS or S3), a web-based application that facilitates collaborative development of the DSN schedule (see references [2–5]). While SSS is operated by DSN, it is actively used by scheduling teams from different missions who must negotiate and concur on a final schedule that accommodates the competing needs of multiple spacecraft. These teams are geographically diverse, with representatives at different facilities on the US East and West coasts who have a 3-hour difference in their working hours.

### 2.1.1. Scheduling Workflow

The scheduling process begins about six months before execution, when each mission submits its expected requirements for a given week. These requirements typically include:

- Number of tracking passes needed over what time intervals
- Acceptable tracking durations and time gaps between passes
- Required frequency bands (S, X, or K/Ka-band) and other ground system configuration requirements — this constrains which antennas are available to support the mission, as not all antennas are able to support all bands
- Specific timing constraints and mission-critical operations
- for longer tracks, a range of valid durations, and the option to split a track into segments with a specified minimum duration and corresponding separation or overlap
- service-dependent setup time and teardown time that has to be included on all tracks (and on any split segments)
- for certain tracks, the option to schedule on either a single antenna, or on a simultaneous array of two (or more) antennas

DSN supports a Multiple Spacecraft Per Antenna (MSPA) mode of operation, where up to four spacecraft that are within an antenna's beamwidth can simultaneously downlink on a single antenna, with only one permitted to uplink at a time. While MSPA scheduling improves network efficiency, it adds complexity in that the individual spacecraft operations have to be fully and carefully coordinated.

Once mission requests are submitted, the DSN scheduling engine component of SSS integrates them and runs a series of AI-driven search algorithms to resolve conflicts without deleting any mission requests. There are a large number of factors considered, including:

- tracking time must only be scheduled on an antenna when a spacecraft is in view of that particular antenna, and when the antenna is available (not in maintenance or other downtime); however, pre-track setup and post-track teardown time can be scheduled out of view
- scheduled activities on one antenna must not overlap with tracking the same spacecraft on another antenna (except for above-mentioned arrays or special navigation tracking, or for critical event redundancy)
- different spacecraft cannot generally be simultaneously scheduled in overlap on the same antenna — such a configuration is called a *facility conflict*

Despite the use of AI, oversubscription is inevitable, and many conflicts must be manually resolved by expert schedulers who use their knowledge of acceptable reductions to edit the schedule. This process takes days of intensive work, after which the schedule is finalized and released for fine-tuning and mission-to-mission negotiation. The baseline schedule is typically finalized 16 to 18 weeks before execution, providing missions with the stability needed to develop spacecraft activity sequences.

Once the baseline schedule is in place, any modifications require concurrence from all affected missions, as changes can have cascading impacts on spacecraft operations. The normal process for this is for one user to submit a change proposal, with any antenna track changes or deletions suggested; all affected users must either concur to the proposal, or reject and resubmit with their proposed alternative changes. This process can take days to work through, in part because missions may have to review and consider changes and what impact they may have to spacecraft operations. In addition, the scheduling teams work normal business hours in different time zones, and this also can add to the time for a change proposal to work through the process to concurrence.



it can be expected that a vehicle of this complexity, coupled with avoiding risk to the crew, will make future Artemis flights similarly unpredictable.

- 2) **Schedule deconfliction:** schedule conflicts that are introduced by additional requirements can be time consuming to clear. While the Service Scheduling Software (SSS) includes both automated algorithms and constraint visualization to facilitate interactive schedule deconfliction, it can still be a significant effort to explore the best ways to fairly make cuts to existing missions in order to fit in new requirements.
- 3) **Overall system performance:** DSN systems are hosted on relatively old (2015) on-premise Linux and Sparc servers, and working with long multi-week contingency schedules was reported as slow by users.
- 4) **Negotiation process:** The DSN process of collaborative peer-to-peer negotiation did not lend itself to the rapid large scale changes imposed by Artemis launch date delays. For example, the DSN SSS software works by having a user exclusively “lock” a section of the schedule while they are making proposed changes, blocking other users from simultaneously changing the same time range. While serializing work in this way makes it easier for users to manage their proposed updates, it also acts as a bottleneck in the overall process when large numbers of updates are required.
- 5) **Trajectory updates:** Responding to updated trajectory files was sometimes confusing and delayed: missions provide numerous contingency trajectories pre-launch, which must be cleaned up and managed post launch as trajectories are updated with accurate tracking data. This process includes a number of manual steps.
- 6) **Situational awareness:** the need for schedulers to be aware of change proposals, or new conflicts, was implemented in SSS through a dashboard — but drilling down to get to the details requires explicit action on the part of each scheduler. From a process perspective, a better model that supports faster turnaround is to support filtered notifications per scheduler, that answer at a glance questions about what needs to be done next. This was one of the motivations for the DSN Intelligent Virtual Assistant (DIVA), as described in reference [7].

### 3. Improvements to DSN Scheduling

In this section we describe how the challenges and improvement suggestions noted above have been addressed.

#### 3.1. Performance

Improving the overall performance of the Service Scheduling Software (SSS) application has been approached from several different directions. As noted above, the servers on which the scheduling application runs are 2015 vintage compute nodes running in an Oracle Private Cloud Appliance (PCA) infrastructure. These servers are approaching end of life, and are in the process of being upgraded to current generation Nutanix hyper-converged infrastructure with redundant servers. This should provide a factor of several speedup in processor, I/O bandwidth, and storage performance. In addition, the Nutanix architecture will provide redundancy in case of hardware failures, and will make software version upgrades much more straightforward. The servers are in the process of deployment by the middle of the calendar year 2025.

In parallel with the server upgrade, an analysis has been conducted of opportunities within the software to tune the application and database. Several database tables have been identified with inefficient joins and are being restructured. In conjunction with migrating the database server from Oracle to Linux on Nutanix, this is expected to significantly speed up the application.

In addition to software changes, plans are proceeding to introduce a database schedule archive capability into the SSS application, to reduce the size of the active operational database. This requires running a second version of the software in read-only mode for viewing and reporting from historical data.

#### 3.2. Schedule Deconfliction Support

Resolving conflicts in the schedule can be a time-consuming process due to the complexity and interaction among all of the project scheduling requirements. In case of a launch delay, a typical scenario is to have a launch, and its associated high importance supporting tracks, “drop” onto a fully negotiated and baselined schedule. Needless to say, this can be disruptive to the missions already scheduled at that time. It is frequently the case that there is very little lead time to come up with a deconflicted schedule ready to support the launching mission, while also supporting as much as possible the missions that were already scheduled.

The scheduling development team at JPL has been working on algorithm updates focused on this specific scenario — which is very different from the initial schedule deconfliction phase for the whole week. Among the factors in play are how to propagate the impact of the launching mission, taking into account the priorities and preferences of the displaced tracks. There are also questions of how best to present options to the schedulers working on this problem, especially where there are alternative impacts based on different decisions to be made. Good progress is being made towards a software update that is expected to address a large part of this problem by providing multiple options for resolving the conflicts.

### *3.3. Ephemeris processing and viewperiods*

Pre-launch ephemeris files are used in scheduling to calculate when to place tracks within expected visibility times at all DSN antennas. After launch, updated ephemeris files are received that reflect any variations from the predicted values, for example due to burn times or other factors. Additionally, for launches that do not occur at the start of the launch window interval, the update will reflect the changes due to the actual time of launch. For Artemis 1, which typically had a 2 hour launch window, the actual launch could have been any time within that period, with resulting significant changes in visibility for the first few subsequent tracks.

As might be expected, delays within the window can have the effect of shifting the times of start and end of acquisition by some amount. More concerning is the case when an entire visibility period may occur only if launch occurs at a particular time in the window, but disappears if launch happens at other times. In this case, for example, a track scheduled on the presumption of launch at or near the open of the window, can shift and eventually disappear as the launch time moves towards the end of the window.

The DSN team is approaching this in two ways: the first is to obtain sample ephemeris files for a periodic set of times within each launch window, say spaced at 5 minutes. By analyzing the effect of the change in launch time, the team can be prepared for shifts in visibility that might affect tracking plans. In some cases, tracks can be scheduled for times that are in the intersection of all visibility periods for any launch time in the window, and thus be essentially guaranteed to be usable, no matter what time the launch occurs within the window. The second approach is to speed up the processing of ephemeris updates to the point where making a schedule change becomes practical and efficient, even in the immediate aftermath of launch.

### *3.4. Situational Awareness*

The DSN scheduling application, SSS, provides various dashboards and views of the schedule and scheduling process that are available to all users. However, it is based on a presumption that users are logged into the application and actively browsing for status and other information. It was recognized a few years ago that a more proactive mechanism would be very useful to the scheduling user community, and potentially reduce the cost of scheduling the DSN. The result has culminated in a new application called DIVA, for DSN Intelligent Virtual Assistant, that runs separately from SSS but monitors it for various situations. Among the capabilities provided by DIVA are:

- a mission-focused dashboard to display the status of schedule negotiation and any outstanding conflicts, with quick links to the full schedules
- notifications over Slack of key events, such as change proposals that need a response, or new antenna conflicts
- quick metrics summary of tracking hours by mission by week
- a mechanism to capture and submit requirements on a weekly basis

There is a separate paper on DIVA in this conference to which the reader is referred[7].

### *3.5. Real-time contingency scheduling*

Under extremely rapid turnaround requirements, the normal DSN negotiation process may not have time to function normally. This is due to extreme time pressure to come up with a feasible schedule, and the inherent delays due to the back and forth nature of negotiation — proposing, counter proposing, and obtaining mission concurrence for all schedule changes. To address this risk, the DSN has developed a special scheduling process that can be invoked only in these extreme (dire) circumstances, in which a designated “Real-time Schedule Arbiter” is appointed to direct and perform schedule changes, following which impacted missions can appeal to make or undo changes. The appeal process

can involve escalation if necessary. The goal of this process is to reach a conflict-free and feasible schedule in order to meet a very near-term deadline.

#### 4. Conclusions and Future Directions

As this is written, Artemis 2 is scheduled to launch in early 2026, and all of the improvements noted above are expected to be deployed and in use ahead of launch.

For Artemis 3 and future missions, there is another set of changes underway: DSN is adding S-band capabilities to additional antennas at each DSN complex, in anticipation of two vehicles, both with astronauts on board, at the Moon needing simultaneous DSN support. The expectation when tracking a crewed vehicle would normally be to dedicate *two* antennas, one as primary and one as a hot backup. When needed, the latter hot backup asset could assume the primary role in a matter of minutes. For two simultaneous crewed vehicles, a new capability is being developed to use a *shared* backup antenna, designated *standby* backup, ready to switch over to primary for either one of the crewed vehicles. Because this involves a change in frequency and telecom configuration, this is not just a scheduling change but also affects DSN processing end to end, driven by the way the primary and backup antennas are designated in the schedule.

Acknowledgements: This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

#### References

- [1] "Basics of Space Flight - Solar System Exploration: NASA Science," <https://solarsystem.nasa.gov/basics/chapter18-1/>, Apr. 2021.
- [2] Johnston, M. D., Tran, D., Arroyo, B., Call, J., and Mercado, M., "Request-Driven Schedule Automation for the Deep Space Network," *SpaceOps 2010*, Huntsville, AL, 2010.
- [3] Johnston, M. D., and Tran, D., "The DSN Scheduling Engine (DSE): Automated Scheduling of Activities for the NASA Deep Space Network," *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, Sapporo, Japan, 2010.
- [4] Johnston, M. D., and Tran, D., "Automated Scheduling for NASA's Deep Space Network," *6th International Workshop on Planning and Scheduling in Space (IWPS)*, Darmstadt, Germany, 2011.
- [5] Carruth, J., Johnston, M. D., Coffman, A., Wallace, M., Arroyo, B., and Malhotra, S., "A Collaborative Scheduling Environment for NASA's Deep Space Network," *SpaceOps 2010*, Huntsville, AL, 2010.
- [6] Dodd, S., "NASA Science Committee, Presentations from August 29-30, 2023 Meeting," <https://science.nasa.gov/science-committee/>, Aug. 2023.
- [7] Johnston, M. D., Mason, J., Lucena, G., Christman, P., Wong, A., Watson, D., and Peich, K., "Automating the Request Process for the Deep Space Network," *Proceedings SpaceOps 2025*, Montreal, Canada, 2025.