

Scheduling Resources on the Deep Space Network and the Mars Relay Network

Eve Pereira^a

^a *Jet Propulsion Laboratory, California Institute of Technology, USA, eve.pereira@jpl.nasa.gov*

Abstract

Efficient communication between Earth and spacecraft is essential for space exploration. The Deep Space Network (DSN) is the biggest communications network used by NASA (as well as various international and private industry partners) for data transmission directly to and from many spacecraft. The DSN is a limited resource and, with a growing user community, time allocations on antennas are coordinated via a collaborative effort by the DSN scheduling group.

Data transmissions from Mars lander missions entail an additional layer of complexity. Antennas on Mars landers are smaller, which causes lower data rates and longer transmission times. The creation of the Mars Relay Network (MRN) helped circumvent these problems by using Mars orbiter missions to act as relay interfaces; however, this added the need for coordinating the use of relay assets within the Mars community. Time allocations of relay overflights are coordinated via a collaborative effort by the MRN.

For both the DSN and the MRN, scarcity of the resources and the ever-evolving nature of requirements increase the difficulty of the scheduling task. The growing complexity of resource requirements has made past attempts at automating scheduling a challenge that has yet to be overcome.

The DSN antenna scheduling tools and processes and the MRN relay scheduling tools and processes are handled by two separate entities that have no functional overlap and are completely independent of each other. Experience with both systems, with 20+ years of operations use in two separate fields and completely independent of each other, clearly shows that a communal, human-involved approach to scheduling deep space assets significantly improves operational flexibility and accommodates the need for an egalitarian approach to prioritization over an automated and/or a direct user-to-asset scheduling approach.

Keywords: Deep Space Network, Mars Relay Network, DSN, MRN

1. Introduction

Transmitting data from spacecraft to Earth is a vital component of space exploration. The Deep Space Network (DSN) is the primary communications network between Earth and numerous deep-space and near-Earth spacecraft. With only 14 antennas supporting 50+ DSN users [1], the network is a limited resource and antenna allocation requires coordination that is handled via a collaborative effort by the DSN scheduling group.

Data transmissions from Mars lander missions add an additional layer of complexity to the process. Due to the smaller antenna size of landed assets and the consequent lower data rates, transmission from the surface of Mars directly to Earth is slow and has limited daily opportunities (when the landers face Earth). The creation of the Mars Relay Network (MRN) [2,3], which enlists orbiters to relay data from landers to Earth, helped solve these problems. Orbiters have bigger antennas than landers, which allows for larger data rates and faster transmissions. In addition, relay overflights remove the limitation of landers having to be in direct line of sight of antennas on Earth, and with orbiters passing the landers multiple times every day, many more opportunities are available to collect data from the landed assets. However, the cost of relay is the added component of scheduling the use of assets within the Mars community. Time allocations of relay overflights, which orbiter talks to which lander, navigational limitations, cross talk, and other factors need to be considered during coordination, and, after 20+ years of experience, the process has evolved into a successful, collaborative effort by the MRN.

For both the DSN and the MRN, scarcity of the resources (i.e. there are not enough antennas to provide comprehensive communications opportunities to all the spacecraft that want them) and the ever-evolving nature of requirements increase the difficulty of the scheduling task. Attempts at automating scheduling have been made [4,5,6]; however, the problem of limited resources and a large user community along with intricate details of requirements have proven to be a stretch for AI and, at this point, require human involved scheduling to assure fairness and adequate coverage for all users.

The many hurdles faced by DSN antenna scheduling and the MRN relay scheduling are similar but are handled by two separate entities that have no functional overlap and are completely independent of each other. Experience, as well as attempts to implement multiple approaches with both systems throughout each of their 20+ years of operations, clearly show that a communal, human-involved approach to scheduling deep space assets significantly improves operational flexibility and accommodates the need for an egalitarian approach to prioritization over an automated and/or a direct user-to-asset scheduling approach.

This paper will provide insight into the processes and procedures of the current approach to DSN scheduling (section 2) and MRN coordination (section 3). Section 4 will provide a comparison of these two approaches, highlighting the pros and cons of each. A discussion of possible alternative solutions will be presented in section 5.

1.1 Definitions

The following definitions will be used in reference to DSN scheduling:

- “Viewperiods” are intervals of time when a spacecraft is within view of and can be tracked from a DSN ground station.
- “Track” or “DSN track” refers to a time when a spacecraft is communicating with a DSN antenna.

The following definitions will be used in reference to MRN coordination:

- “Overflights” are intervals of time when a landed asset on Mars is within view of a spacecraft orbiting Mars.
- “Relay pass” or “relay session” refers to a time when a lander has established a radio link and is in communication with an orbiter.
- “Planning period” is a 2-week increment of MRN schedule that is being coordinated at the same time.
- “Tentatives” are potential overflight opportunities chosen by landers in the first round of reviews.
- “Proposals” are overflight opportunities that have been chosen by landers in the second round of negotiations.
- “Requests” are the ultimate passes chosen by lander teams.
- “Forward link” is the transmission of data from Earth to a Mars lander via relay through a Mars orbiter
- “Return link” is the transmission of data from a Mars lander to Earth via relay through a Mars orbiter

The following definitions will be used throughout this paper:

- “Lander” is any spacecraft that has landed on Mars, including both mobile (rovers) and stationary assets.
- “Orbiters” in this context are spacecraft orbiting around Mars.
- “Data” refers to binary information that is passed between Earth and a spacecraft either directly or via a relay pass.
- “Uplink” is sending data from Earth to a spacecraft via the DSN, or from an orbiter to a lander during Mars relay.
- “Downlink” means sending data from a spacecraft to Earth via the DSN, or from a lander to an orbiter during Mars relay. “

1.2 Background

The DSN was officially established in 1963 to provide independent, centrally operated communications support to deep-space missions. [7] The DSN provides tracking services, which include commanding, downlinking of science and spacecraft data, observation of spacecraft navigational data, and gathering of data for radio science experiments. More importantly, it allows for the monitoring of the health and safety of a spacecraft. From its early beginnings in 1963, the network has grown to entail 14 antennas in 3 station complexes in Goldstone (USA), Canberra (Australia), and Madrid (Spain). Each complex consists of one 70m antenna, and three or four 34m antennas.

Arguably, the most famous mission sets tracked on the DSN were the Apollo missions to the moon in the 1960ies, and the Voyager missions, which launched in 1977 and are tracking via the DSN to this day. Currently, the DSN supports around 40-50 users, including deep-space missions, near-Earth orbiters, and ground-based telemetry observations. These missions have a wide variety of objectives, such as orbiting a planet, landers, rovers, sample return, impact, etc., and a broad range of targets, including planets, asteroids, comets, moons, and more.

While during its early days the DSN was easily available to its small user community, the drastic increase in the number of DSN users has made it a highly oversubscribed asset and the coordination of asset allocation an increasingly complex task for the DSN schedulers. [8] Additional factors contributing to the high contention are viewperiod alignments between missions, overlapping mission critical events, antenna downtimes, and an increasing number of critical events that require increased antenna coverage. It is the schedulers’ responsibility to work through this plethora of conflicts and negotiate a schedule that is conflict free and meets mission requirements (Figure 1).

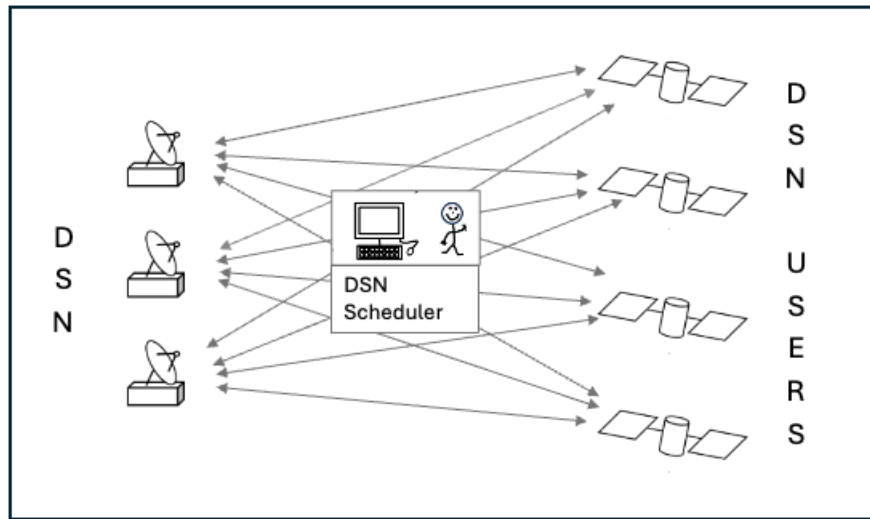


Figure 1: DSN scheduling

Though relay at Mars has been utilized since 1971, the current instance of the MRN, a highly successful collaboration between the National Aeronautics and Space Administration (NASA) and the European Space Agency (ESA), began evolving in 2004 with the twin rovers Spirit and Opportunity, and officially came into existence in 2010 with the development of the MaROS (Mars Relay Operations Service) tool in preparation for the Curiosity rover's 2011 post-landing operations [9]. The current participants in the MRN consist of two NASA landers: Curiosity and Perseverance, three NASA orbiters: Mars Odyssey (ODY), Mars Reconnaissance Orbiter (MRO), and the Mars Atmosphere and Volatile Evolution (MAVEN or MVN) mission, and two ESA orbiters: Mars Express (MEX) and the Trace Gas Orbiter (TGO). During a relay session, the orbiters function as an intermediate vehicle to transmit data from a lander to the DSN (return link) or from the DSN to a lander (forward link). Though both current landed assets do have the ability to communicate directly with the DSN, due to size and power limitations, their antennas and radio systems are substantially smaller than those of the orbiters, hence limiting the maximum data rate, which in turn significantly increases data transmission times. An additional drawback is the lack of line of sight from Mars to Earth for about 12 hours per day due to planetary rotation. Today's landers do have regular direct DSN passes but use these primarily for uplinking commands. Nearly all data returned to Earth from Mars is transmitted via relay orbiters.

The MRN consists of a smaller number of users than the DSN, nevertheless, allocation of relay assets still requires a coordinated and agreed-upon process that involves all relay partners. Like DSN negotiations, the MRN coordination process cannot be fully automated at this point and is reliant on hands-on involvement of the participants. Besides being "node-poor", i.e. having a limited number of available assets, factors such as orbiter availability (each orbiter has their own science objectives to fulfil), orbit distance (determines the frequency and duration of potential overflights), geometric feasibility (elevation angle of overflights), lander staff availability (lander planning cycles are dependent on the results of previous cycle data and thus have a much shorter turnaround time), etc. must be taken into consideration. A coordinated approach to managing the flow of Mars relay operations has been developed and documented [10]. It is the participants of the MRN's responsibility to navigate through these challenges, follow the MRN operations guidelines, and coordinate a conflict-free plan that meets mission needs (Figure 2).

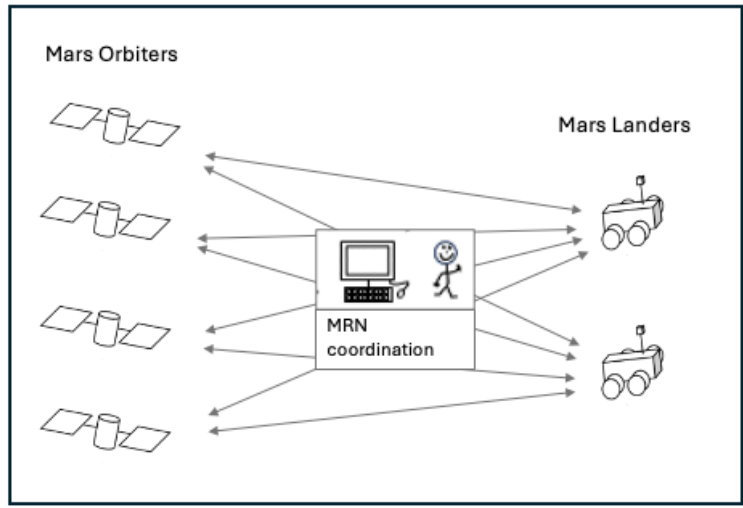


Figure 2: MRN coordination

Figure 3 shows the data flow to and from Mars, utilizing both DSN scheduling for conflict-free antenna allocation and MRN coordination for relay asset assignments.

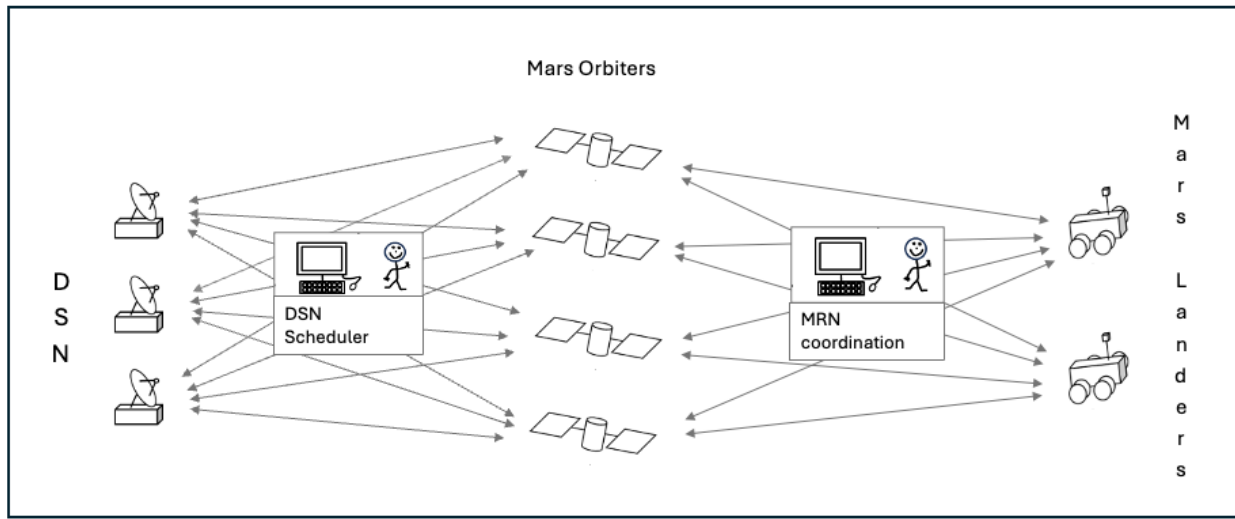


Figure 3: Forward and return link through DSN scheduling and MRN coordination.

2. DSN Scheduling

2.1 Participants

DSN scheduling is a combination of automated schedule generation and hands-on negotiations by various parties. Currently, between 40-50 deep-space and near-Earth missions from several organizations within NASA, as well as private industry, universities, and international partners use the DSN. In addition, ground based observations, monthly DSN maintenance, and antenna upgrade activities require scheduled blocks of time on the 14 operating DSN antennas.

The primary representatives of each mission are the schedulers, who have a deep understanding of requirements and negotiate on behalf of their missions (Figure 4); the mission planners, who provide DSN requirements; and the navigators who provide ephemeris files and navigational data for viewperiod generation.

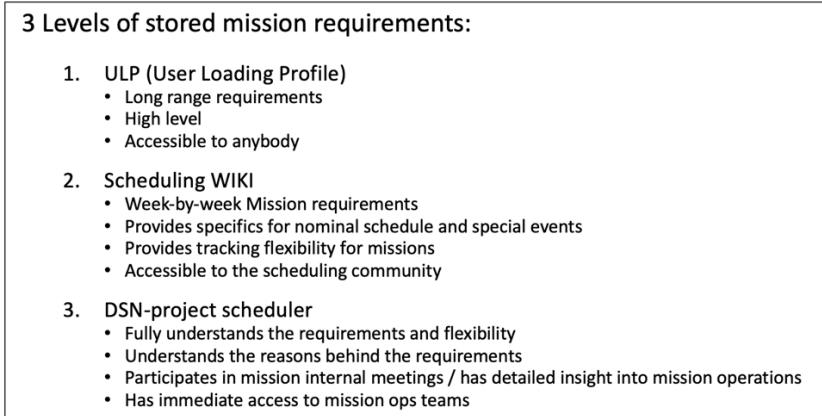


Figure 4: Three layers of DSN scheduling knowledge

Representatives of the DSN include the Mission Interface Manager (MIM) who coordinates technical details between the missions and the DSN, the Network Operations Project Engineer (NOPE) who provides guidance regarding antenna hardware and software, the Schedule Process Owner (SPO) who oversees and coordinates the continued flow of the scheduling process, the Builder of Proposals (BOP), who deconflicts the initial schedule, and the Operations Chief (OPS Chief), who manages real-time operations. In addition, the SPO office serves as a neutral party to arbitrate if conflicts cannot be resolved on the scheduler-level and need to be elevated to project management, provides supportability analysis, and coordinates the downtime planning process.

2.2 Process/Timeline

The DSN scheduling process is handled in three segments. Long-Range, Mid-Range, and Real-Time operations (Figure 5).

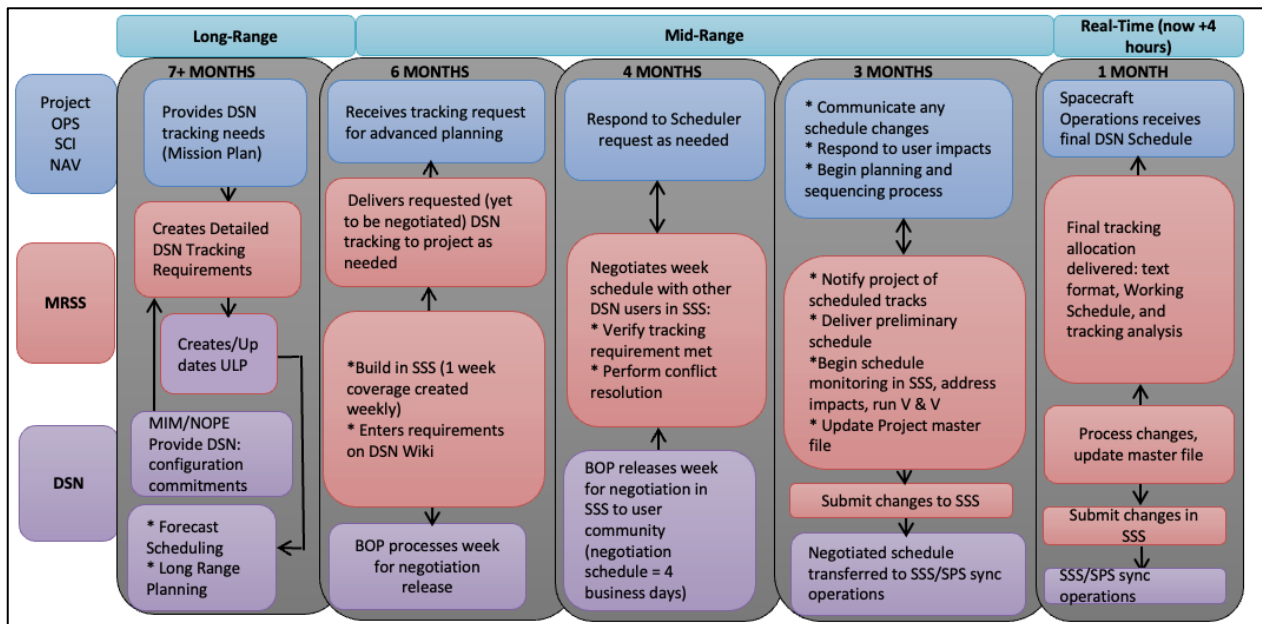


Figure 5: MRSS/Project/DSN Scheduling Process Timeline

The Long-Range planning process begins several years before a DSN track is executed. The mission planner relays mission objectives to the DSN MIM, who aids the mission in defining their high-level DSN tracking requirements. This is summarized in a User Loading Profile (ULP), a chart that lists monthly hours of tracking time

and antenna size required over the entire mission lifespan. The ULP is passed along to various DSN teams to begin long term preparations for tracking support. The ULP is also passed along to the DSN scheduling teams*.

At about 18 months before launch the scheduling team begins to attend mission planning and other relevant meetings to increase their understanding of the mission's objectives and understand their DSN tracking needs. The scheduler works closely with the mission planning team to define DSN requirements and flexibilities, such as track duration and frequency, antenna size, and support configurations during different mission phases and high-level events, potential tracking restrictions, and any other factors that may impact the ground team from communicating with the spacecraft.

At around 25 weeks before real-time tracking, the schedulers submit project tracking requirements to the DSN via the Service Scheduling Software tool (S³) weekly in one-week increments [11]. S³ ingests the inputs from the roughly 50 DSN users to generate the first version of the DSN master schedule. This version is highly contentious, containing on average 375 tracks in conflicts. Conflicts include, but are not limited to, two or more missions with overlaps in scheduled time on the same antenna, the scheduling of more equipment than available at the same DSN Complex, two or more tracks beginning their activity at the same time in the same complex, or tracks being placed outside of their mission's viewperiod. The number of conflicts also varies depending on the number of high-level activities occurring in a week.

This schedule is released to the Builder of Proposal (BOP), who attempts to minimize the number of conflicts through the application of strategies that utilize the designated flexibilities for each of these tracks. The BOP is highly knowledgeable of the more detailed, written and unwritten needs of every spacecraft, their flexibilities, and minimum tracking requirements. Once the BOP reduces the number of conflicts to a sufficient level (no more than 20), the schedule is released to the project schedulers for a negotiation period via S³.

The schedulers review, renegotiate, and optimize their mission's coverage, using their intimate knowledge of requirements and flexibilities of their missions. Though each scheduler is looking to maximize coverage for the project they represent, they also actively collaborate with their peers to generate the most efficient master schedule that meets all missions' needs. Negotiations are primarily conducted via S³ but also involve conversations between schedulers via the use of instant messengers, emails, phone calls, and old-fashioned hallway conversations. In an effort to optimize scheduling efforts, a review of the DSN scheduling process showed that there are three layers of knowledge for the scheduling task (Figure 4), with most of the detailed requirements knowledge residing with the schedulers.

During highly contentious weeks the scheduler and the mission planner often work closely together to develop solutions that maximize efficiency of the limited DSN tracking time available to the projects. These solutions may include the replanning, postponing, or cancelling of science activities to fit within the DSN schedule, especially during overly subscribed weeks when missions are required to reduce their tracking time to the bare minimum that is needed to maintain the health and safety of the spacecraft. Negotiations generally involve multiple iterations of the cycle until the schedule is conflict free (Figure 6).

* DSN scheduling is in rare cases done mission internal without employing a professional scheduling team; however, since most missions are scheduled by the MRSS team, the focus of this paper will be on missions scheduled by this team.

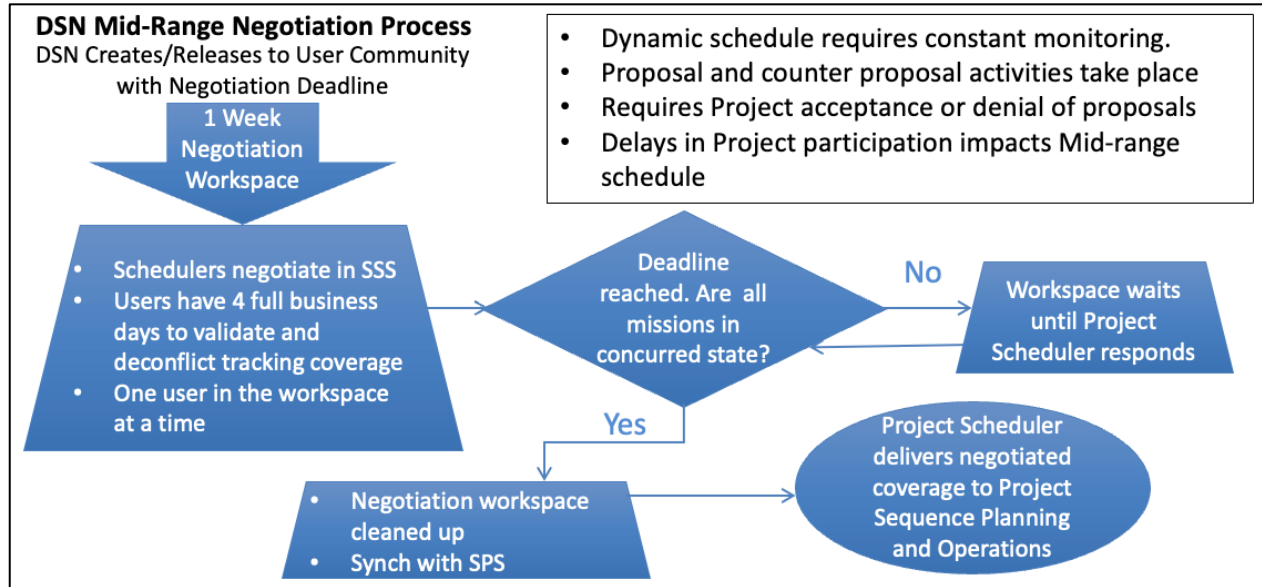


Figure 6: DSN Mid-Range process

Negotiations are completed at approx. 16-18 weeks before real-time. This conflict-free schedule remains available and open for changes in S^3 until a track is executed. Though discouraged, adjustments for various reasons (last-minute changes to plans of a mission, unexpected antenna downtimes, spacecraft emergencies, etc.) often are necessary and permitted until the beginning of track. During emergencies, stemming from either a spacecraft or an antenna, DSN passes can be adjusted mid-track.

Starting four hours before start of a track, the OPS Chief takes control over the schedule. He or she oversees the smooth continuations of operations, and in rare cases, should an unexpected situation arise, begins to coordinate and implement any changes while working directly with on-console/on-call mission operators. For example, if a spacecraft experiences an emergency or an antenna encounters technical issues in Real-Time, the OPS chief will step in and try to acquire track time on another station. The OPS chief will work with the on-call mission operators from active DSN users to evaluate priority and then work out the most efficient reassignments of DSN antennas.

2.3 Tools

Although each scheduling team utilizes their own internal tools to generate mission specific files that will be used in the command product generation process, S^3 is the central hub for DSN schedulers to negotiate DSN tracking requirements (Figure 7).

The S^3 functionality is strictly schedule centric, supporting tracks throughout the various steps of the process. As outlined above, interaction with the S^3 tool begins when schedulers submit their long-term spacecraft tracking requirements. S^3 retrieves applicable viewperiods previously generated from trajectory files submitted by the Navigation teams, and generates the initial master schedule, which is released into the scheduling process for deconflicting and optimization. S^3 is the repository for the DSN master schedule from the initial generation until after tracks have been executed and the schedule is archived. While schedulers communicate in person or through email, phone, messengers, email, etc., any official negotiations are conducted exclusively through S^3 . Negotiating through a centralized tool enforces the conformity to standardized formats of inputs and products, which is important as the finalized schedule will be utilized by DSN operators to actually control the antennas.

It is noteworthy to point out that because of the functionality of S^3 , the need for any regular in-person meetings between mission schedulers has been eliminated.



Figure 7: S³ view of the DSN master schedule for week 14-2025

3. MRN Coordination

3.1 Participants

Like DSN scheduling, Mars Relay coordination is a combination of automated schedule and product generation via hands-on negotiations by various parties. Currently, 7 Mars missions spanning multiple organizations within NASA and ESA use the MRN for daily operations.

The primary representatives of each mission are the relay planners for both orbiters and landers, who negotiate relay requirements on behalf of their missions, and who then generate the relay files that will then be incorporated into their respective mission’s spacecraft command products.

Representatives of the MRN include the MRN OPS team, who coordinates the relay process, verify proper documentation, and ensure the smooth flow of operations, and the MaROS support team, who verify the stability and robustness of MRN tools and pursue the continued improvement and development of new software. In additions, MRN management serves as a neutral party when mitigating relay opportunity conflicts and to arbitrate in the event of conflict between projects.

3.2 Process/Timeline

The Mars relay coordination process is handled in two parts: the strategic and the tactical planning processes. Additionally, the MRN office annually generates the “Mars Mid-Range Schedule”, which is a consolidated timeline that lists all relay related timelines and provides the relay community a guideline for operations due dates.

The Strategic Relay Coordination process begins up to four months before the first relay session of a two-week planning period. The process begins with the creation of several products. Navigation teams deliver orbiter trajectories and lander locations, which are then used to predict when the orbiter will be visible to the surface asset (overflights). The orbiter teams deliver files containing various planned onboard activities, such as times available for relay, their DSN schedule, and when personnel will be available to radiate commands to the orbiter. These products are used to generate a report of potential relay opportunities, which are available to all participants through the Mars Relay Operations Service (MaROS) tool. [12]

The next step in the process is pre-coordination, during which the lander projects identify potential candidate sessions (tentatives) for relay support from orbiters. They also consider other external constraints, such as power availability on the landers, schedule constraints for lander activities, the local times of day of the landers, and the data volume predictions. Using these considerations, the lander projects identify relay sessions that are most advantageous. Should the situation arise where tentative passes by one lander can affect the others’ (through co-location, crosstalk, schedule conflicts, etc.), the lander teams negotiate until they agree upon which orbiters and overflights are to be used by each lander. The process generally consists of a series of emails or messages sent via other messaging systems, but, if more coordination is required, an in-person meeting can be held.

Once the candidate relay sessions have been identified, the lander teams submit proposals to MaROS. After evaluation, the orbiter teams respond via MaROS with an acceptance, a denial, or a modification of the proposals. This process may occur in an iterative fashion. The final selections by the lander teams are then submitted to MaROS

as requests. Requests are either acknowledged or rejected by the orbiter teams, and a final cycle of back-and-forth negotiations ensues.

Negotiation between the lander projects and the orbiter projects conclude before the bi-weekly Short-Range Relay Coordination (SRRC) Coordination Meeting, which is coordinated by the MRN OPS team and during which all requests are reviewed, any outstanding conflicts are resolved, and an official relay schedule is signed off by all projects. The SRRC meeting officially marks the end of the strategic part of the relay planning process and the tactical phase begins.

The Tactical Planning Process manages short-term changes to the relay plan in the timeframe between the SRRC meeting and pass execution. This generally only entails minor modifications, such as a change in a selected return-link data rate but may also contain the occasional response to an emergency by either landers or orbiters. Negotiations similar to those used in the strategic process commence, and final agreement of modifications is submitted to MaROS. When a change has been implemented, the orbiter projects verify that the as-implemented state of any relay session is accurately reflected within MaROS.

Occasionally the standard process is augmented by additional coordination to support special events such as solar conjunction periods, emergencies of a spacecraft, the addition of a new relay partner [13], or any other unexpected event.

3.3 Tools

Though each spacecraft team utilizes their internal tools to generate mission specific files that will be used in the command product generation process, MaROS is the central hub for relay negotiations to be tracked and validated (Figure 8).

As outlined above, interaction with the MaROS tool begins in the early stages of relay negotiations when various spacecraft teams submitting inputs to generate a report of potential overflight opportunities for a planning period and continue through post-pass reporting. MaROS shows at any point the latest official status of negotiations of any overflight. The tool also provides automated detection of potential schedule conflicts and the validation of relay session parameters for each pass.

In addition to being the primary support tool for relay support negotiations, MaROS provides several other functionalities, such as calculation of latency predictions, computation of data volume predictions, transfer of forward and return link data between the lander Mission Operations Centre (MOC) to the orbiter's MOC, generation of geometric relationships (e.g., rise, set, elevation angles, etc.) between spacecraft, a three-dimensional visual representation of the MRN, and a collection of as-flown metrics from all participants in the MRN for each relay session.

Parameters	Latencies		Strategic				Post Pass	
	Geometry	Tentative	Proposal	Planned	Request	Implemented	Orbiter Scorecard	Lander Scorecard
Category/Plan		normal	normal	supported	normal	supported		
HallStartTime	2025-066T01:54:29.896	2025-066T01:54:46.000	2025-066T01:54:46.000	2025-066T01:54:46.000	2025-066T01:54:46.000	2025-066T01:54:46.000	2025-066T01:54:46.000	2025-066T01:54:48.000
HallDuration	0T00:17:06.378	0T00:17:03	0T00:17:03	0T00:17:03	0T00:17:03	0T00:17:03	0T00:17:32	0T00:19:59
HallEndTime	2025-066T02:11:36.274	2025-066T02:11:49.000	2025-066T02:11:49.000	2025-066T02:11:49.000	2025-066T02:11:49.000	2025-066T02:11:49.000		
LMT	1438T08:29:33.519	1438T08:29:49.192	1438T08:29:49.192	1438T08:29:49.192	1438T08:29:49.192	1438T08:29:49.192		
Submitter	Navigation	M20	M20	erro tgo	M20	European Relay Coordination Office	European Relay Coordination Office	m20 tel
SubmitTime	2025-076T14:44:05.000	2024-311T19:28:20.000	2025-043T01:05:45.000	2025-015T14:05:11.000	2025-066T01:05:48.000	2025-056T15:59:18.000	2025-067T02:01:55.000	2025-067T18:17:25.000
DataVolumeLRE				1034.5	1031.0	1031.0		1091.4
DataVolumeLRE_ID				431589	437755	450038		462533
Negotiated								
LinkType		both	both	both	both	both	return	both
Protocol		PROXIMITY_1	PROXIMITY_1	PROXIMITY_1	PROXIMITY_1	PROXIMITY_1	PROXIMITY_1	PROXIMITY_1
Non_Coherency		Ncoh_NoTurnarond	Ncoh_NoTurnarond	Ncoh_NoTurnarond	Ncoh_NoTurnarond	Ncoh_NoTurnarond	Ncoh_NoTurnarond	Ncoh_NoTurnarond
ForwardRate		32	32	32	32	32	32	32
ReturnRate		32	32	32	32	32	32	32
Doppler		NONE	NONE	PHASE	NONE	NONE	NONE	
G2_Inversion_TX		G2vectNotInvertd	G2vectNotInvertd	G2vectNotInvertd	G2vectNotInvertd	G2vectNotInvertd		
G2_Inversion_RX		G2vectNotInvertd	G2vectNotInvertd	G2vectNotInvertd	G2vectNotInvertd	G2vectNotInvertd		
Fwd_Link_Channel		0	0	0	0	0		
Fwd_Link_Carr_Mode		FRES	FRES	FRES	FRES	FRES		
Fwd_Link_Encoding		CONVOLUTIONAL	CONVOLUTIONAL	CONVOLUTIONAL	CONVOLUTIONAL	CONVOLUTIONAL		
Frame_Size		1016	1016	1016	1016	1016		
Window_Size		20	20	20	20	20		
Rtn_Link_Channel		0	0	0	0	0		

Figure 8: MaROS screen showing an M20-TGO overflight (Individual submitter names were removed)

4. Comparison of Scheduling Approaches

4.1 Similarities of DSN scheduling and MRN coordination

There are numerous similarities between DSN scheduling and the MRN coordination process, which all lead back to the basic challenge the community faces: how do you allocate tracks or passes in a heavily oversubscribed system in a manner that is effective, egalitarian, and affordable? Both networks, completely independent from each other and with no functional overlap, have over time organically grown into the same communal, human-centred negotiations processes that have been working for 20+ years. These processes have some basic commonalities:

- *Process flow*: Both networks begin the process with a very high-level schedule that the community works through in several iterative negotiation steps, during which requirements and allocations are narrowed down to eventually produce a conflict-free schedule that is acceptable for all parties.
- *System setup and tools*: Each of the two processes involves a common tool (S³ and MaROS) that serves as the central hub to the entire community. Though some of the negotiation might happen directly between some of the participants, the official latest status of the schedule will always be tracked in the tool. Both S³ and MaROS are central to the process from very early on – both tools start out with navigational data that indicates potential passes – until the end, when they retain post-pass tracking data for archival/historical purposes.
- *Human factor*: Another commonality is that highly knowledgeable representatives for each project are hands-on involved in the negotiation of DSN tracks and MRN passes throughout the process. These participants know their projects’ detailed requirements and flexibilities, and they use this knowledge to develop a highly efficient schedule that optimizes an exceedingly oversubscribed asset.
- *One-on-one support*: For the majority of support on both the DSN and the MRN, tracking is supported via one-on-one pairings. For most DSN tracks, a spacecraft is scheduled on one antenna at a time, thus monopolizing the asset. Similarly, Mars relay orbiters can only talk to one rover at a time and must do so in a way that avoids cross-talk.
- *One-on-one support exception*: An exception to the one-on-one pairings applies to both networks when two or more service users are in close proximity. The DSN supports MSPAs (Multiple Spacecraft Per Aperture) downlink when spacecraft are near enough to each other that their signal can be detected within the beamwidth reach of an antenna. Due to the need for proximity this feature can currently only be used

with the Mars missions, and, occasionally, for short periods when spacecraft that launched together are on a similar trajectory.

For Mars Relay, the UHF radios of orbiters and landers are limited to one-on-one connections; however, if two landers are close to each other, they can use the same geometric overflight of an orbiter by utilizing the “shared overflight” functionality, which allows the orbiter to switch mid-pass from one lander to another.

4.2 Differences of DSN scheduling and MRN coordination

Though the high-level approach to DSN scheduling and MRN coordination is very similar, the two networks are completely independent of each other, and, as would be expected, there are differences:

- *Operations size:* The most notable difference between DSN scheduling and MRN coordination is the size of the community and the number of negotiations. The DSN user community is far larger than that of the MRN. It provides communication services across 14 antennas to approx. 50 users with a wide variety of objectives and locations in space and on Earth. In contrast, the MRN currently provides relay services to 5 orbiters and 2 landers (though MaROS, the ground system that supports the MRN is scalable and adaptable and could be easily implemented at the moon or any planet).
- *Operations scope:* The DSN is a dedicated communications network, and their core business is to provide telecommunications support between various spacecraft and Earth. The MRN is composed of science orbiters with their own science objectives, often adjusting their science plans to accommodate relay passes. Though there are discussions of potentially sending dedicated relay orbiters in the future, at this point each Mars mission (orbiter or lander) is supporting various instruments and is working to achieve their own science goals.
- *Personnel:* DSN schedulers are full-time schedulers. They may represent multiple missions, but the primary objective of their work is to negotiate the DSN schedule for their customers. MRN relay planners negotiate relay passes; however, they also generate command products for their spacecraft and are often involved in multiple teams within mission operations.
- *3rd party involvement:* All DSN assets are controlled by a 3rd party, the DSN, whereas relay assets are controlled by individual missions. This has no impact on the actual data flow, but introduces additional challenges (antenna maintenance, etc.) that need to be resolved.
- *Automation-level:* All DSN negotiations are done via S³, which eliminates the need for any communitywide in-person meetings. Though some level of communication still happens outside of the tool, this is generally done on a one-on-one or small sub-group level and mostly involves members of the same scheduling team. DSN users have no direct interaction with each other and in most cases their only common denominator is the use of the DSN antennas, which are largely the same in what they can do and how they are set up for interaction with their customers.

The MRN coordination process on the other hand includes several meetings that require representation from all missions. This is primarily driven by the fact that there is direct communication between spacecraft within the MRN and technical aspects of this communication increase the complexity of the process - all the relay orbiters have different capabilities, are built distinctly unique, and have their individual science missions in addition to the relay efforts.

5. Evolution of the systems

Both the DSN and the MRN were not conceived as fully planned ideas that were then implemented, but rather both systems organically evolved over time based on the need of the community. This natural evolution brought forth developments that could not have been foreseen at the early beginning stages. The continuously expanding customer base of the DSN increased the complexity of the scheduling process. While early on reams and reams of printouts and coloured pencils were the central tools to negotiations, the development of S3 afforded efficiency and time-savings, and allowed the scheduling community to keep pace with the ever-increasing number of DSN customers.

A similar natural evolution was seen at the MRN – while not quite paper and coloured pencils, the early tools of the MRN were a number of reports for long- and short-range planning, that then were posted to a website which could be accessed by users. Any communication and negotiations were conducted via email, phone calls, or in person. The final versions of the reports were ultimately reviewed and discussed in the short-range relay planning meeting. Though this meeting exists to this day, the process leading up to and post-meeting have been modernized. The variety in Mars

project interfaces (every orbiter and lander behaves differently and has a unique ground data system) drove up operations costs when new missions were introduced. The development of the MaROS tool provided a centralized system with standardized formats that allow for easy interactions between missions, no matter what their internal system encompasses.

Standardization of interfaces help streamline processes, provide increased efficiency, and ease the introduction of new participants. On the DSN side, tracking requirements are generally built and submitted directly in S³. Another, less commonly used option, is to submit requirements and flexibilities by uploading a standardized XML file to the tool. If exact event windows for tracks are available, these can also be submitted to the system in one of two standardized formats. The MRN specifies and documents common MaROS file formats that are being used by all participants via SISs (System Interface Specification).

6. Alternative Solutions

The most obvious solution to simplify both these scheduling processes is to turn a scarcity of resources into an abundance of resources: build more antennas for the DSN and send more relay orbiters for the MRN. Though planning for more antennas and future relay assets is ongoing, the actual implementation of either is a process that takes years of preparation and is exceedingly costly, thus not an immediate alternative option.

For an alternative to the current approach, ideally, both the DSN scheduling and the MRN coordination functions would be fully automated; however, the combination of limited resources, oversubscribed systems, and complex requirements have made a fully AI-based solution unachievable for now. Developers from the DSN and the MRN engage in ongoing efforts to improve and expand on their tools (S³ and MaROS, respectively). These efforts have been successful but are gradual and limited in scope. A big limitation to accelerating AI efforts is a lack of funding, although the bigger obstacles to fully automating both the DSN scheduling and the MRN coordination processes are the sheer scope and details of these tasks, and the difficulty – and sometimes reluctance – of system users to change their established tools and processes.

Automating the DSN schedule creation based on levels 1 and 2 scheduling knowledge (Figure 4) is possible to some extent. One of the initial steps in the current process is the automated generation of the master schedule via S³, though this is still dependent on the manual submission of requirements by the schedulers, and the outcome is highly contentious and requires human intervention to clear the conflicts. Automation of level 3 is not a feasible option at this point – there is simply too much detail and flexibility that cannot be fully captured by an automated system. Currently, AI is not an alternative solution, but rather a series of small steps towards the automation of existing processes.

A different solution to the current community-based approach would be direct project-to-antenna and lander-to-orbiter scheduling. In this scenario, each mission would submit any requests for DSN antenna time directly to the DSN, and each Mars lander would send their requests directly to the orbiter teams. It would then be the responsibility of the DSN and the orbiter teams to either accept or reject these requests. Presumably, there would be a process for the acceptance or rejection of passes. This may be a pre-determined priority system (e.g. manned spacecraft before unmanned, or rover before stationary lander), a “first-come-first-serve” basis, a combination of both, or any other schema that can be formulated.

The biggest advantages to this direct scheduling approach are cost and time savings. For DSN scheduling, the roles of the BOP, and the schedulers would be fully eliminated. In addition, the role of the SPO would be reduced from overseeing and ensuring the smooth execution of the scheduling process to simply accepting or rejecting tracking requests based on the pre-defined adopted schema, or, taking this idea one step further, the SPO would be eliminated, and projects would submit their requests directly to the stations.

Similarly, for the MRN, the iterative planning cycles (tentatives, proposal, request) would be eliminated and landers would just submit simple requests to the individual orbiter missions, and in turn eliminate the need for MRN involvement. In both cases, this reduction in staff would provide substantial savings. Additionally, the functionality of software could be reduced, or the software could be eliminated altogether – either way, providing savings. Another advantage of the direct scheduling approach is the lack of a defined timeline. The elimination of several iterative negotiations cycles minimizes the number of deadlines, thus providing greater flexibility.

The biggest disadvantages of the direct scheduling approach are the inefficiency of the process, the lack of coordination, and the potential unequal and undesirable outcome. Depending on the schema adopted for acceptance or rejection of DSN tracks or relay passes, the outcome can be inconsistent and vary greatly between weeks and/or projects. Missions may be continuously below their minimum requirements, they may have passes at unfavourable

times (local night shift), they may have violations of gap- or other timing requirements, or they encounter a multitude of other scenarios produced by an inefficient schedule – in short, an unfortunate subset of users might be “starved out” of their requirements, while missions that are deemed a higher priority could have an overallocation of resources. Additionally, missions wouldn’t have any advance input to schedule changes which could require missions to rework their command products to prevent data loss. Continuous readjustments and below ideal schedules significantly complicate spacecraft operations and increase cost.

For the MRN, past experience has proven that this solution is less than ideal. The direct user-to-provider approach was the initial way the MRN was scheduled. Lander teams requested support directly from the orbiters and had to comply with whatever data formats were required by the orbiter projects. Additionally, they needed to build a specific interface to interact with each orbiter directly. Not only was this costly but also required an immense coordination effort on each project. The introduction of MaROS abstracted these mission-unique interfaces away and phenomenally changed the entire network. Landers are now able to request relay services from any orbiter via the single portal of MaROS using a single set of file formats.

7. Conclusions

Both the Mars Relay Network and the DSN today are considered great successes, despite their immutable limitations; however, with an ever-growing number of users, both networks will need to keep exploring further options for automation, increased efficiency, and optimization. Currently, the system of human-involved, communal negotiations is the best solution because it provides functioning answers to the problem of oversubscription of resources. The intimate knowledge that DSN schedulers and MRN coordination personnel have of the requirements, needs, and flexibilities of their spacecraft cannot be fully recreated via automation. The process of iterative, communal negotiations and collaboration provides efficient, egalitarian, and optimized allocation of resources.

20 years of operations experience by these two separate entities that have no functional overlap, yet that notably evolved to similar solutions, clearly show that in the current environment a human-involved-over-automated and communal-over-individual approach to scheduling deep space assets significantly reduces cost and effort over a direct user to asset scheduling approach.

Though nearly full automation of both the DSN and MRN processes could ultimately lead to highly efficient schedules at decreased cost, it is not a realistic solution now and won’t be until the problem of scarcity of resources has been resolved. Simply stated, we need more antennas for the DSN and more relay orbiters for the MRN!

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 (80NM0018D0004).

Many thanks to Amanda Kniepkamp, Ricky Cors, Adam Spanier, and Rylee Mitchell for unlimited access to their DSN scheduling knowledge. A special thank you to Roy Gladden, who patiently guided me through the process of writing this paper, reviewed and provided feedback, and was overall ridiculously helpful.

References

- [1] JPL: Deep Space Network NOW: <https://eyes.nasa.gov/apps/dsn-now/dsn.html>
- [2] “A Martian Telecommunications Network: UHF Relay Support of the Mars Exploration Rovers by the Mars Global Surveyor, Mars Odyssey, and Mars Express Orbiters”, C. Edwards et al, International Astronautic Congress, IAC-04-M.5.07, Oct 2004, Vancouver, Canada.
- [3] NASA: Mars Relay Network, <https://science.nasa.gov/mars/mars-relay-network/>
- [4] “Scheduling NASA’s Deep Space Network: Priorities, Preferences, and Optimization”, M. Johnston, ICAPS 2020, October 2020, retrieved from https://icaps20subpages.icaps-conference.org/wp-content/uploads/2020/10/SPARK-2020_paper_1.pdf

- [5] “Prioritization and Oversubscribed Scheduling for NASA’s Deep Space Network”, Shouraboura et al, ICAPS 2016 – SPARK Workshop, London, UK
- [6] “A Demand Access Paradigm for NASA’s Deep Space Network”, Dhamani et al, retrieved from: <https://ai.jpl.nasa.gov/public/documents/papers/Dhamani-IWPSS2021-paper-18.pdf>
- [7] NASA: History of the Deep Space Network, <https://www.nasa.gov/history/history-of-the-deep-space-network/>
- [8] NASA Office of Inspector General, “Audit of NASA’s Deep Space Network”, 12 Jul 2023, IG-23-016 (A-22-01-00-MSD), retrieved from <https://oig.nasa.gov/wp-content/uploads/2023/12/ig-23-016.pdf>
- [9] “The History of the Use of Orbiter Relay at Mars, 1972 to 2024”, R. Gladden, et al, SpaceOps 2025, May 2025, Montreal, Canada
- [10] “Lessons Learned from the Mars Relay Network: Considerations for Future Relay Networks”, R. Gladden et al, 2024 IEEE Aerospace Conference, March 2024, Big Sky, MT, USA.
- [11] “Automating Mid- and Long-Range Scheduling for NASA’s Deep Space Network”, Johnston et al, SpaceOps 2012, June 2012, Stockholm, Sweden
- [12] “Mars Relay Operations Service (MaROS): A Present Service Preparing for the Future”, R. Gladden, SpaceOps 2014, May 2014, Pasadena, CA, USA
- [13] “Preparing the Mars Relay Network for the Arrival of the Perseverance Rover at Mars”, Gladden et al, 2022 IEEE Aerospace Conference, Aug 2022, Big Sky, MT, USA