SpaceOps-2025, ID # 344

Preparing the Mars Relay Operations Service for the Challenges Ahead

Brandon Sauer*, Michael P Newcomb[b]

[a] *Jet Propulsion Laboratory, California Institute of Technology, USA, brandon.sauer@jpl.nasa.gov*
[b] *Jet Propulsion Laboratory, California Institute of Technology, USA, michael.p.newcomb@jpl.nasa.gov*
* Corresponding Author

Abstract

The Mars Relay Network (MRN) will need to adapt to the evolving needs of current and future missions that rely on it for transmitting data to and from Earth. Research was conducted with the goal of identifying the factors driving adaptation and to assess how the Mars Relay Operations Service (MaROS) – a ground system service used for planning relay – can support the MRN in addressing future challenges. The research involved conducting interviews with various subject matter experts involved in Mars telecommunications, relay networking, and technology. It also included reviewing published materials on future projections for a Mars telecommunications network detailing capabilities and operational concepts of future missions and architectures. Our research indicated challenges centered around the theme of scalability facing both the network's physical infrastructure and its operational personnel that if not addressed, will have a negative impact on network efficiency and ultimately, Mars mission performance. This paper will provide an overview of those challenges, followed by recommendations for mitigation through addition of new capabilities to the MaROS software platform.

These recommendations include equipping MaROS with capabilities that support orbiter-to-orbiter relay planning, automating scheduling to meet mission and network constraints, updating to universally recognized interface standards, and more. Through implementation of these recommendations, a foundation can be established to address the upcoming challenges and accommodate the demands of a growing network and accompanying operations teams. If these recommendations and others are added to the suite of features MaROS provides, it will be better-positioned to address the exigent factors of current and future Mars missions.

---

1. Introduction

As Mars missions continue to evolve to capture higher volumes of data, the need for efficient telecommunication with Earth becomes increasingly critical for ensuring the success of those missions. The Mars Relay Network (MRN)] is an international collaboration of orbiter mission "providers" that relay science and telemetry data from landed mission "users" at Mars to with Earth. The MRN is responsible for nearly all of the data exchanged between Earth and Mars. A pivotal tool to the MRN is MaROS (Mars Relay Operations Service), a ground system service developed and operated by the Jet Propulsion Laboratory (JPL) which facilitates the planning, scheduling, and negotiation details for all relay interactions [1-4].

This paper investigates the future capabilities and configurations of the MRN, with a focus on how new relay link technologies, evolving operational dynamics, and a potential influx of new relay customers could impact the system. After identifying the potential challenges Mars mission operators will face, we examined the current capability of the

MaROS system to identify potential short-comings. By conducting primary research interviews and reviewing existing literature on Mars relay services, future planned missions, and projections of telecommunications infrastructure, our research aims to identify potential risks and opportunities to ensure the continued successful operation of the MRN from the perspective of maintaining a capable and efficient ground service.

One key finding of the research is the identification of a significant risk: there is a need to address scalability challenges ahead of new technologies and participants joining the MRN in the coming years, which will require changes to the network from both technological and personnel operations processes. These changes could necessitate substantial updates to MaROS in order to maintain its effectiveness and address these new operational needs. This research underscores the importance of forward-thinking, proactive strategies to ensure the adaptability and long-term success of the MRN.

2. Literary Research Overview

A broad review of existing literary material was conducted focusing on relay network architectures, proposed Mars missions, and projections of telecommunications technologies and paradigms. The scope of the research offered insight into various ways the MRN could be expected to evolve over the next couple decades and how emerging technologies could present challenges to future MRN missions.

One of the primary findings in our research was the prevalent prediction that future Mars missions will have highly increased data volume requirements. Missions with high-resolution imaging, data-intensive instruments, and eventually, human exploration, warrant the need for magnitudes more data to be transmitted to Earth with less latency than the current relay network architecture can support [5]. It's worth noting that while much of the existing research details the efficiencies of new protocols, compression methodologies, and bandwidth technologies for both the proximity and deep-space links used in the network, it was determined that MaROS planning and scheduling activities are largely unaffected by such details of the physical link layer.

The NASA Moon-to-Mars initiative implies that communication standards, link technologies, and architectures that will be implemented as part of the current emphasis on lunar explanation will eventually influence the future of Mars missions [6]. It is expected that the Delay/Disruption Tolerant Networking (DTN) paradigm will be introduced at Mars in the future [7], resulting in MRN participants adopting new data storage and communication protocols, on-board automated routing capabilities, and other mechanisms that differ significantly from the current relay process, and especially the ground systems such as MaROS. As is the case at the moon, there is also an initiative to involve commercial industry in the future implementation of the MRN [8].

Additionally, architectures predict the introduction of dedicated relay orbiters at Mars – as opposed to current relay support provided by science orbiters – that will result in higher data throughput, coverage, and availability capabilities. These dedicated relay orbiters may act as constellations, supporting cross-links between multiple orbiters for increased coverage and availability. Eventually, dedicated relay orbiters at higher orbit elevations will act as "trunk links" capable of collecting data from all other spacecraft at Mars and, with near-constant view of Earth and coverage from deep-space ground stations, relay vast amounts of data as soon as possible [9].

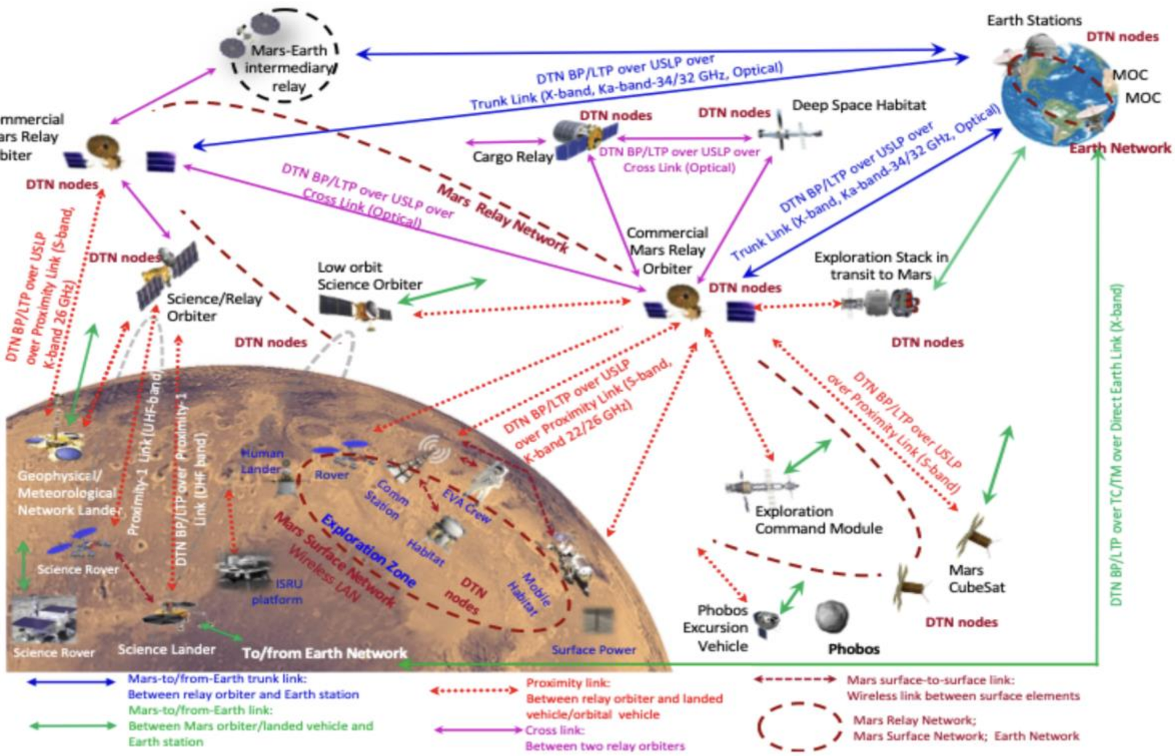

Figure 1: A long-term future Mars communication architecture from the IOAG Future Mars Communications Architecture [8], depicting orbiter-to-orbiter cross-links, multiple Mars-Earth trunk-links, and DTN-enabled nodes.

The literature review provided a broader perspective for both conducting and reconciling the findings from the expert interviews and helped frame the challenges identified by the research community.

3. Expert Interview Research Methodology

Our research aimed to identify emerging link technologies and network configurations that would serve the Mars community and understand the impacts to the operation of the MRN through interviews with a wide variety of experts. The research sought to define the challenges facing the MRN in the next five years and beyond, focusing on the operational and network factors that will cause potential disruption in relay operations of future missions if not addressed.

In order to gather the data needed to respond to the research goals, specialists within the field were sought out and we then conducted qualitative interviews with 18 subject matter experts (SMEs). These experts represented a wide range of areas critical to Mars relay operations, including telecommunication, technology development, DTN, and spaceflight operations for Mars orbiters. The interviews were designed to explore the challenges facing Mars relay based our findings from literary review, expanding on both current and future technological and operational needs. Given the range of expertise and knowledge, the experts did not require a great deal of prompting to provide insights.

Each interview lasted between 30 and 45 minutes, and the initial questions posed to the experts that led to lengthy discussion included:

- "In terms of link capabilities, what challenges do you see based on current technologies being used?"

- "What do you see as some of the primary challenges facing Mars relay in the near future (within the next five to ten years)?"

These open-ended questions were crafted to gain diverse information from these individuals regarding the limitations of existing technologies and hardware, as well as to gather expert opinions on the challenges that may arise as Mars missions' needs continue to evolve parallel to future MRN infrastructure. The qualitative nature of the interviews allowed for rich, in-depth responses, which provided valuable perspectives on both the immediate and long-term needs of the MRN.

Ethical guidelines were followed throughout the study to ensure that participants' rights and confidentiality were respected. All participants were made aware of the research objectives and the voluntary nature of their participation. Their responses as statements were collected, and participants were assured that their anonymized responses would be used solely for the purposes of this research project [10].

The data from the expert interviews was captured and reviewed using a thematic analysis, also known as the "Atomic Research Nugget," [11] approach. This approach allows for an insight to be derived from a direct or summary quote from a participant and then associating it to a commonly found or emerging theme in the research. Issues discussed in the interviews, such as limitations of the current network and anticipated future challenges, were categorized into themes, which were later corroborated against our literary research results to ensure a more comprehensive understanding of the challenges facing the MRN in the future.

While the interviews provided valuable insights into the future challenges facing the MRN, several limitations must be acknowledged. First, the expert interviews focused on a relatively small sample of 18 individuals, meaning the findings may not fully represent the broader range of expert opinion in the field. Additionally, the review of papers and reports encompassed a wide range of sources, but the dynamic nature of space mission timelines, advances in spacecraft hardware, and the momentum behind telecommunication technologies may result in recent developments not being captured. And lastly, it must be acknowledged the future-looking nature of this qualitative research required at least some amount of speculation and assumption from our SMEs and ourselves in order to make our best guess at what the future may hold.

This approach, combining expert interviews with a comprehensive review of existing research, enabled an in-depth exploration of the challenges facing the MRN. The insights gleaned from this overall approach will inform ways in which MaROS can adapt and mitigate these challenges.

4. Identifying Future Challenges

By taking the 18 different interviews [12] and cross-referencing them with the landscape of available research, we were able to discern high-level patterns regarding the challenges facing the MRN in the next decade or more. **Scalability** emerged as a central theme throughout the research findings, with two distinct partitions to this theme identified: **Scalability in the Infrastructure** and **Scalability in Operations**. Both are essential to ensuring the MRN can meet the demands of future Mars missions. The following challenges are ordered by the number of times each was brought up by different SMEs during interviews.

Scalability in the Infrastructure refers to the physical network infrastructure's ability to handle an expanding number of missions that will need to utilize relay and to accommodate advancements in telecommunications technology.

1. **Evolution of Link Capabilities (11 mentions)**

The evolution of link capabilities was brought up the most by experts. Future spacecraft will utilize newer technologies, such as frequency changes for both the direct-with-Earth link and proximity links and adoption of new radio protocols. Directional proximity links may have implications on various relay planning procedures as opposed to current omni-directional links, as might the implied data rates, but it was determined that the MaROS planning and scheduling process is largely agnostic to changes in the physical link layer.

2. Preparing for Delay-Tolerant Networking (DTN) (6 mentions)

Many experts agreed that delay tolerant networking is the future of the MRN, but the path of how we transition to a DTN-enabled network is still developing. DTN is an entirely new communications paradigm, and there will undoubtedly be an impact to the operations concepts for the entire network. This study focused considered how MaROS can support the transition of the MRN to a fully DTN-enabled network, including the inevitable future of a hybrid network with some legacy nodes mixed with newer DTN-enabled nodes.

3. Entropy of Existing Relay Nodes (6 mentions)

Over the next decade, some of the older MRN orbiters are expected to degrade, with a high probability that at least two relay providers will be inoperable or less efficient for relay by the early 2030s. This reduction in relay providers, combined with scarcity of plans for replenishment of new providers will put additional strain on the remaining orbiters to meet increasing demand, thereby making relay planning and scheduling more complex and competitive.

4. Increased Data Rate and Coverage Needs (5 mentions)

Future missions, including data-intensive imaging robotic missions and eventual manned missions, will need to be able to transmit data at rates and volumes magnitudes larger than the network can support today. This increased data throughput will also come with more expansive coverage needs, where communication will need to be as synchronous as possible. Long-term future missions may have the need for “demand-access” relay communication, with the expectation that they can transmit data via relay at any time and have an orbiter available to capture this data and relay it to Earth almost immediately.

5. New Missions Joining the Network (3 mentions)

Though an exact timeline across the many national space agencies and now private industry is impossible to predict, it is safe to say that there are more players in the game with interest in sending missions to Mars than ever before. The MRN will likely continue to welcome new participants, when possible, but all of these new participants will undoubtedly put additional stress on the network infrastructure as they compete for telecommunications resources.

6. Ground Stations will be More Oversubscribed (2 mentions)

The proximity link between spacecraft at Mars is only one leg of the relay. Deep-space ground stations, such as the Deep Space Network (DSN), are already considered an oversubscribed resource that are imperative for communicating with deep-space missions. With the expected influx of lunar and other deep-space missions, this resource will become even more impacted. Reduced Earth communication opportunities means that scheduling time for orbiters downlinking relay data will be more competitive and thus difficult to obtain.

Scalability in Operations refers to the increasing number of teams of people that operate the missions that comprise the MRN and their interactions with one another to coordinate and enable relay. MaROS, as the centralized interface for these missions to interact with one another, will be under increased stress to support these teams' coordination on the ground. SMEs identified several key areas where scalability in operations was critical for the continued facilitation of relay.

1. **Downlink Constraints on Planning (8 mentions)**

   As orbiting relay providers became inoperable or less performant and deep-space ground stations become less available, the planning process will in turn experience more constraints in relay planning and scheduling. Fewer opportunities available for downlinking data could significantly complicate the scheduling process for relay user teams, likely reducing the amount and timeliness of data that a mission can return to Earth. The implications of this finding could be quite severe, reducing each mission's ability to operate on Mars.

2. **Need for More Flexible Scheduling (7 mentions)**

   Though some orbiters have made strides to support more flexible scheduling, the current baseline process for relay planning is multiple weeks long. Future missions will likely need this process to be more agile, allowing for more relay opportunities and flexible scheduling not beholden to the weeks-long planning process of today. This flexibility will reduce the resources each mission devotes to relay planning, but the path to updating the operations process to be more flexible remains unclear, thus far.

3. **Growth of Planning Process Membership (6 mentions)**

   Each new mission joining the MRN, as either a provider or user, comes with the overhead of more support staff focusing on planning operations. This is a whole new set of people who must learn how to conform to the planning process, and consequently, must learn how to interact with MaROS – sometimes through the implementation and maintenance of specific tools. From the MaROS-perspective, increasing the number of users we must support will likely result in increased costs for operations and maintenance.

4. **Adapting to New Partners and Their Needs (4 mentions)**

   The planning process needs to evolve to better support a growing community of commercial and government organizations. These organizations may seek to establish their own approach to relay planning and operations if the current one is deemed too inefficient or cumbersome for their mission operation needs. MaROS should continue to evolve its capabilities to be adaptable to new partners' processes and the collaboration necessary to provide ongoing value to new organizations in the future.

5. **Manual Analysis of Post-Pass Telemetry and Anomalies (3 mentions)**

   When a relay pass does not perform as anticipated, mission operators must scour telemetry data for the landers and orbiters to determine the cause of the anomaly. Each mission has their own nuances that have to be considered and manual detection of these anomalies for different mission combinations is time and resource intensive. In a larger network with more distinct nodes and more mission interaction combinations, maintaining a manual process like this for anomaly detection and analysis will require even more human intervention.

6. **Increasing Lines of Communication (1 mention)**

   As new missions join the MRN with their own operations staff, there will be more people that need to communicate with one another during the relay planning process than ever before. This problem, though simple, is somewhat unique to the nature of relay, as by definition, it requires working and communicating with others. In the case of the international collaboration present in the MRN, it is expected that this communication may be highly asynchronous, with operators in different time-zones, and potentially speaking differently languages. Though possible to manage through conventional email and internal communication tools, this could prove to be a significant communication challenge to ensure that teams are aware of parallel relay activities occurring.

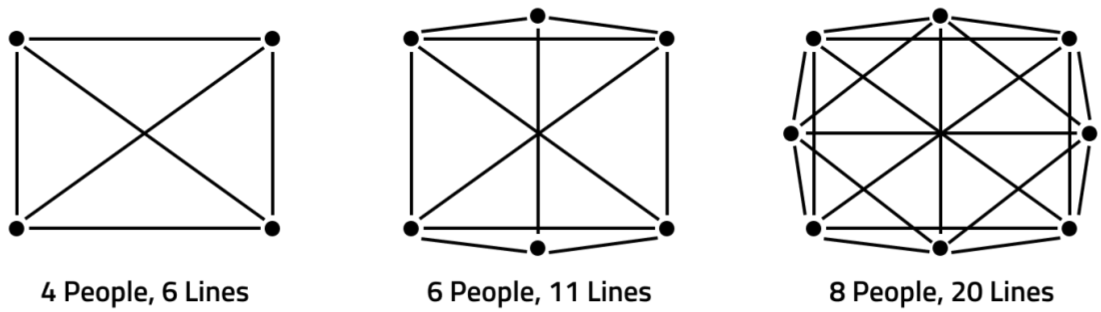

Figure 2: Lines of communication visualization, showing number of lines growing as team size expands

5. Reconciling Results with MaROS Capabilities

After consolidating the results of our research into a primary set of challenges, we set out to find potential enhancements to existing MaROS capabilities or new features that could help lessen the impacts identified. Whenever possible, we preferred cross-cutting recommendations that would address multiple challenges simultaneously and tried to only consider single-issue solutions for challenges that could not be otherwise addressed or if the proposed recommendation was somewhat simple ("low-hanging fruit"). We also determined that some challenges that were identified separately as Scalability in the Network and Scalability in Operations were relevant to the same solutions for similar reasons and thus could be roughly combined for the purposes of validating the recommendations (e.g. "New Missions Joining the Network" and "Growth of Planning Process Membership" were ultimately similar enough to be mitigated by the same recommended solutions). After this consideration, we detailed a set of three larger, cross-cutting recommendations for new features that could be added to MaROS to mitigate multiple challenges, along with three single-item recommendations that could help to address some of the challenges specifically or with less effort.

Recommendation	Challenges Addressed
Supporting Orbiter-to-Orbiter Relay Add the ability to plan communications between two orbiters in MaROS, as opposed to the current capability of only lander-orbiter relay.	● New Missions Joining the Network ● Growth of Planning Process Membership ● Ground Stations will be More Oversubscribed ● Preparing for the Future of DTN
Automated Scheduling Suggestions Ingest geometric view periods along with mission-specific and network-wide constraints to programmatically generate a suggested schedule for the entire network that maximizes data volume and timeliness while minimizing scheduling conflicts.	● New Missions Joining the Network ● Growth of Planning Process Membership ● Downlink Constraints on Planning ● Entropy of Existing Nodes ● Lines of Communication Increase ● Adapting to New Partners and their Needs

	• Preparing for the Future of DTN
Standardizing MaROS Interfaces Conform to a more universally-recognized set of file formats – such as CCSDS standards – for users to input/export planning data to/from MaROS.	• New Missions Joining the Network • Growth of Planning Process Membership • Adapting to New Partners and their Needs
Centralized Communication Platform Integrate a built-in communication platform within MaROS for centralized location for negotiation and discussion of relay planning.	• Increased Lines of Communication
Automated Anomaly Detection and Improved Post-Pass Analysis Tooling Automatically ingest post-pass telemetry to detect anomalies in relay performance to explain reductions in returned data volume and notify interested parties.	• Manual Analysis Post-Pass Telemetry and Anomalies
Improved Event Message Service Allow for more self-service subscription to listen for MaROS events to further improve mission automation.	• Need for More Flexibility

Figure 3: Table summarizing the MaROS improvement suggestions along with the challenges they could each help mitigate.

5.1 Supporting Orbiter-to-Orbiter Relay

While MaROS is currently geared toward coordinating relay between landers and orbiters, adding support to schedule relay between orbiters will enable future missions that expect orbiter-to-orbiter relay as part of their operations concept and will pave the way for multi-hop networking that can benefit all relay users.

Future proposed network architectures plan to rely heavily on orbiter-to-orbiter cross-links as part of their operations concept. In more advanced architectures, lander and orbiter missions alike will rely on dedicated high-orbit relay providers that act as a “trunk link”, providing significantly higher throughput for communication with Earth than what is achieved through current science orbiters. Proactively providing support for these use cases in MaROS will enable future missions to leverage existing multi-mission services to standardize relay planning activities.

As more relay user missions arrive at Mars, there will also be significantly more lunar and deep-space missions expected to utilize deep-space ground stations for Earth communications, reducing the opportunities for Mars relay providers to downlink relay data. The trunk-link concept is expected to serve as a major benefit for all Mars missions constrained by reduced ground station availability, and thus enabling this support in MaROS will be pivotal for supporting the trunk-link architecture. Adding support for orbiters to transfer science or relay data to one another could also allow the network to further take advantage of limited ground station opportunities and decrease latency within the downlink-constrained network.

DTN will introduce a new paradigm for the interfaces used to transfer data between nodes, with orbiter-to-orbiter and multi-hop network transfers included as a standard part of routing data through the network. We believe that there are many lessons we can learn from planning orbiter-to-orbiter cross-links and multi-hop transfers in the current manual fashion to help pave the way for DTN-automated interactions.

The process for supporting orbiter-to-orbiter relay planning in MaROS would be extensive, as it involves expanding on nearly all of the current system's capabilities. MaROS would first need the ability to calculate the times where each orbiter spacecraft has a line of sight with one another, called "view periods". As of the time of writing, development is underway to provide orbiter-to-orbiter view periods from the subsystem that calculates lander-to-orbiter view periods, MPX (Metric Predicts Executive), with the intent to eventually ingest these into MaROS whenever new orbiter predicted ephemerides are published. MaROS itself would then need to be updated to ingest these view periods and allow orbiter operators to request relay interaction with other orbiters during a view period. As part of the negotiation process, the other orbiter's operators would need to be able to acknowledge and agree to those requests.

MaROS calculates various scheduling conflicts based on different constraints that would need to be modified to support orbiter-to-orbiter relay, and new types of conflicts specific to orbiter-to-orbiter relay may need to be introduced. Updates would also need to be made for the sub-system that models data volume prediction for each relay pass – the Relay Telecom Predictor (RTP) – to model orbiter-to-orbiter communication. Another sub-system, "Where's My Relay Data" (WMRD) – a dashboard that provides insight into the relay data-flow via specific monitoring data – would also need to be updated. Language throughout the system would need to be updated to not exclusively refer to relay service users as "landers".

5.2 Automated Scheduling Suggestions

Currently, each lander relay planning team is responsible for generating their own schedules for each orbiter, specifying those schedules into MaROS planning products, and submitting them into MaROS. Because these schedules are planned by each mission independently and asynchronously, and due to the nuances that each mission must consider to meet their science, relay, and operations needs, this process requires extensive communication between all participants in the network to ensure everyone is satisfied. This problem becomes exponentially more complex and expensive as the network continues to grow.

To alleviate this scalability challenge, we propose the implementation of an automated relay scheduling algorithm that consumes mission-specific and network-wide constraints, then generates potential relay schedules for all MRN missions that optimize for maximum return data volume and minimize latency while simultaneously minimizing conflicts throughout the network. These schedules would be presented to MRN user mission operators to evaluate, iterate on, and approve – leaving the user mission operators in the driver's seat, but reducing a lot of the overhead that goes into selecting each relay pass for their mission and negotiating with other missions.

Each relay user mission in the network requires a communication planning staff of several people, increasing the total number of people needed for coordination and submission of data products throughout the network. Additionally, each mission team is currently responsible for developing and maintaining their own software tools for scheduling and requesting relay (for relay users) or automation for acknowledging and processing relay requests (for relay providers). Automated scheduling suggestions will reduce communications planning teams' need for creating and maintaining external tools to interface with MaROS and greatly reduce the barrier to entry for new missions joining the network. We also believe that automating the scheduling process can help mitigate the need for direct negotiation of schedules across mission operators.

Assuming relay provider nodes age out of the network over the coming years and relay support becomes even more of a commodity, relay service users will have to compete and negotiate for relay opportunities. Especially if the relay orbiters continue to double as science missions, relay service user missions will be severely constrained for data volume. The network will need to function as efficiently as possible to meet the needs of each mission involved. An optimized schedule generated by MaROS can help to utilize the network to its fullest potential, allowing missions to take full advantage of sparse relay opportunities with minimal negative impact on other missions.

As the commercial partners and other organizations join the network to provide relay services at Mars, future partners will search for methods of saving time and money through means of automation. To ensure that the MRN can attract providers to replenish the network, modernization and optimization of scheduling must be implemented.

DTN will present a completely new paradigm for relay scheduling, as it in theory allows for a more "hands-off" approach to routing data through a network, provided there is sufficient infrastructure available to support it. With each node having awareness of available contact windows and the priority needs of customers, scheduling of data transfer can happen almost entirely between the network nodes themselves with much less human intervention. Investing in the development of automation-assisted scheduling for the MRN today will provide not only a technological segue for the eventual introduction of DTN, but also a cultural one. Current relay service users are accustomed to having fine-grain control of every aspect of relay planning for their mission, which will no longer be the case when a DTN-enabled network is automatically routing data. Participants in the MRN will experience a smoother transition to allowing an automated process to select relay passes for them.

To implement an automated scheduling algorithm in MaROS, we would first need to gather a complete list of constraints from all existing relay missions, considering constraints that may exist in future network configurations based on our research and other network-wide constraints. We would then need to find or develop a constraint language that can codify these constraints against MaROS data.

Latency prediction, along with data volume predictions (which is already broken into a separate service, RTP), are key performance indicators for relay pass selection. Presently, the MaROS latency prediction logic is tightly coupled with the submission of relay pass requests by the landers. In order for an algorithm to generate a schedule that meets relay users' latency constraints, MaROS would need the ability to calculate "hypothetical" latencies for subsections of each view period.

With the constraint language and updated latency prediction models in place, we would then implement the scheduling algorithm as a sub-system within the MaROS ecosystem. The algorithm would ideally generate numerous schedule suggestions ranked by data volume returned and minimized latencies across a planning period and present these to MaROS users in a clear way. These schedules would be reviewed by each relay user mission team, at which point they can choose to accept a schedule as is or make iterative adjustments as needed.

5.3 Standardizing MaROS Interfaces

MaROS supports over a dozen unique file interface formats that have been organically introduced and modified throughout its fifteen years in operation. These files take the form of a variety of XML formats and are used by mission operators for inputting planning, configuration, and post-pass data into MaROS and, to a lesser extent, exporting data from the system. Generation of these files is typically automated by the mission teams through custom tools, but many of these formats leave room for human error which has caused its fair share of difficult-to-trace issues throughout the years. Changes to these file formats, though intentionally uncommon and backwards-compatible whenever possible, create a disruption across most network missions to update their tools to adhere to the new format.

These file formats were developed in-house out of necessity, as there were no universal standards that could achieve what was needed when these interfaces were introduced. Now, however, the Consultative Committee for Space Data Systems (CCSDS) has developed a vetted and highly generalized suite of universal service management standards that could likely be used by MaROS for ingesting and exporting most if not all data products. Adopting the more universally recognized CCSDS standard formats would result in a more stable and versatile interface for new missions joining the network, and less maintenance for the MaROS team to update and maintain their own custom formats.

As new missions join the network, operators must learn the nuances of interacting with MaROS, including its unique file interface. While the MaROS team strives to keep documentation for these interfaces clear and up-to-date, there is still frequent confusion amongst users on the proper usage of some features. Furthermore, in situations where the MaROS file interfaces require a change, each mission must evaluate whether they can support such a change, which may result in an overall expensive update for the entire network. Utilizing a universally recognized set of standards would result in a more stable interface for new missions to adopt. The CCSDS standards are thoroughly documented and reviewed, and may be used by mission teams for other operational processes outside of MaROS, reducing the MaROS learning curve.

CCSDS standards are publicly available, and with new organizations and companies expected to join the network in the coming years, they could view these new standards as a more modern and simple approach for distributing relay planning data to MaROS. Conversely, if MaROS fails to support such standards, it could result in a lack of interest in using MaROS as the center for relay planning data due to its cumbersome standards.

Outside of this research task, some brief analysis was done to discuss how the CCSDS Service Management specs can be adapted for use in MaROS. A single MaROS file format was "translated" into the CCSDS "SMURF" specification, with positive results, while other file formats were tentatively mapped to CCSDS formats that could potentially support their functionality, but further comparison is needed.

New standards may also present opportunities to accept new types of data into the system, such as data volume prediction and relay data-flow monitoring that are currently handled internally through the RTP and WMRD sub-systems, respectively. Once new file formats are thoroughly mapped, MaROS can initially support these formats via the implementation of an internal converter to translate the new file formats into existing internal MaROS data structures. To keep supporting legacy missions, MaROS would be required to continue supporting the existing file formats for the foreseeable future.

5.4 Centralized Communication Platform

While we expect that automated scheduling can assist with the increased lines of communication that come with adding new missions and organizations to the network, there will always need to be teams of humans in the loop. Up to this point, much of the negotiation between lander teams for relay planning has been handled outside of MaROS, notably through use of internal 3rd-party tools for instant communication. However, new organizations and companies, with different time zones, operational procedures, and cybersecurity requirements will warrant a more sophisticated negotiation communication platform.

A MaROS-hosted chat capability, accessible to all participants in the network, would be a reasonable short-term solution to address this challenge. This integrated communication platform could allow for contextual linking capabilities to MaROS data items (e.g. mentioning an “overflight ID” could automatically link to the details for that pass) to enhance the negotiation communication experience. If selected for implementation, we would plan to research existing platform tools that could be integrated into MaROS and any tweaks we may need to make it compatible with MaROS context.

5.5 Automated Anomaly Detection and Improved Post-Pass Analysis Tooling

After a relay pass has completed, relay service user teams analyze the results to ensure that the amount of data expected to be returned was indeed received on the ground at the expected time. Significant discrepancies between predicted and actual data volume or latency are flagged by teams for anomaly investigation. These discrepancies can be caused by unexpected local terrain occlusions, spacecraft hardware issues, interference from other spacecraft in the network, or a host of other issues that require a great deal of manual attention to track down.

The recently introduced "Where's My Relay Data" (WMRD) sub-system in MaROS provides users with the ability to manually analyze post-pass telemetry and other monitoring metrics in a dashboard. We believe it would be possible to expand on WMRD's capabilities to implement some level of automated anomaly detection and analysis – potentially with the help of AI models – to notify relay teams of significant discrepancies in expected pass execution and potential causes. This capability is expected to save significant resources used for manual analysis currently and improve the overall efficiency of the network.

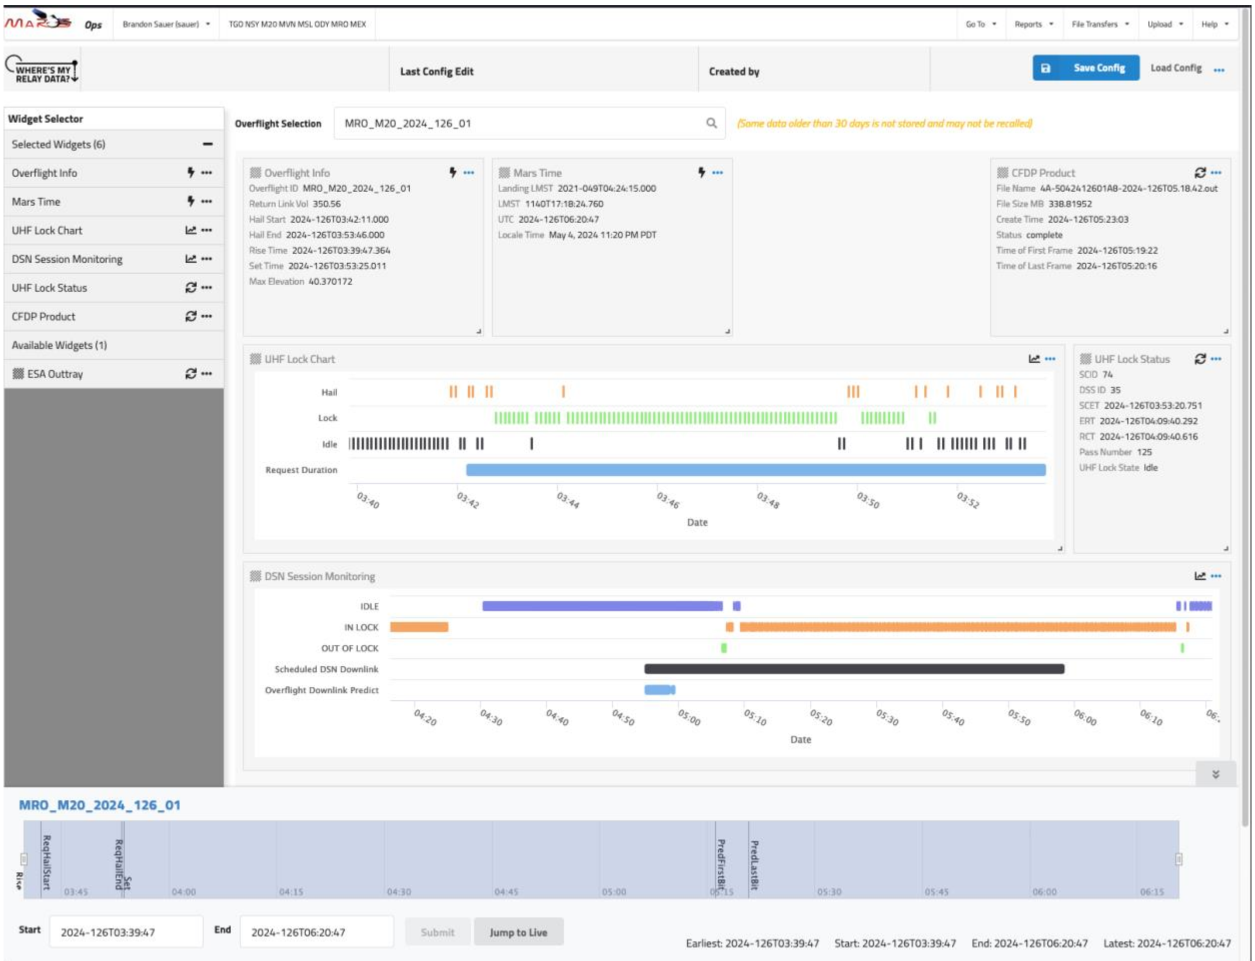

Figure 4: A sample screenshot of the Where's My Relay Data Dashboard in MaROS

5.6 Improved Event Message Service

The current relay planning process is somewhat rigid, requiring plans to be submitted weeks in advance of when the relay passes will take place at Mars. Though MaROS supports tactical changes, it can be challenging for orbiter teams to make rapid, manual changes to their existing plans. However, relay service users of the future may have an increased need for short-notice changes to the strategic plan, especially as we continue to progress toward human exploration on Mars.

Utilizing an internal event message bus, MaROS recently implemented a mechanism to push newly published files to ESA for their orbiters. This mechanism, paired with a great deal of automation efforts on ESA's side, allows them to automatically support tactical changes made before the predicted Last Nominal Uplink Time (LNUT) with little-to-no human intervention. Such automation has made ESA-provided relay opportunities much more flexible to change and should be encouraged for existing and future orbiters as well, but would currently require a great deal of custom integration for each mission. If MaROS could broaden its current event message bus system – used primarily internally thus far – to create a standard way to push published files to subscribers, similar automation would be more accessible for current and future missions. Additionally, the message bus should be expanded to be self-service, allowing users to set up their own message queue and subscribe to specific subsets of event types. Improving this process would help pave the way for future missions to adopt a demand-access scheduling architecture.

6. Conclusion

Looking ahead to the 2030s and beyond, the inverse relationship between an increasing number of spacecraft needing relay services and a decreasing number of spacecraft providing relay service poses various challenges for the Mars Relay Network to overcome. Through our research, we found that the overall expected increase in participants in the MRN will present various challenges that the existing process and systems will likely struggle to keep up with. If these challenges remain unaddressed, we could see significant negative impacts to the amount of data returned from Mars.

To address these challenges, moving toward automation and standardization will be important. Combining this with streamlined operations will not only reduce the number of personnel required per mission but also simplify the onboarding process for new missions joining the MRN. By creating more standardized, automated workflows, future missions will be able to more easily integrate into the network, allowing for greater efficiency and scalability. This approach will be crucial for ensuring that the MRN can meet the increasing demands going forward, providing a key value proposition for future missions looking to leverage its capabilities.

While we don’t have direct control over the growth and impending architecture of the network, we can advocate for these suggested improvements to the MaROS system and related ground services to prepare for these challenges ahead. As the MaROS development team, it will be our goal to acquire the resources necessary to continue to investigate and implement these changes over the coming years to best support the interests of the MRN.

7. 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).

8. References

1. “Mars Relay Operations Service (MaROS): Rationale and Approach”, R. Gladden, SpaceOps 2010, Apr 2010, Huntsville, Alabama, USA.
2. “Mars Relay Operations Service (MaROS): Managing Strategic and Tactical Relay for the Evolving Mars Network”, D. Allard, R. Gladden, IEEE Aerospace Conference, Mar 2012, Big Sky, Montana, USA.
3. “Mars Relay Operations Service (MaROS): A Present Service Preparing for the Future”, R. Gladden, SpaceOps 2014, Apr 2014, Pasadena, California, USA
4. Mars Relay Network, https://science.nasa.gov/mars/mars-relay-network/, 2025
5. “Expanding the Horizons of Mars Science: A Plan for a Sustainable Science Program at Mars”, NASA Mars Exploration Program (MEP), 2024

6. "2024 Moon to Mars Architecture Concept Review: Architecture-Driven Technology Gaps" https://www.nasa.gov/wp-content/uploads/2024/12/acr24-architecture-technology-gaps.pdf?emrc=363045
7. "Delay- and Disruption-Tolerant Networks (DTNs): A Tutorial (Version 3.2)", Forrest Warthman, 2015
8. "Exploring Mars Together: Commercial Services Studies", Richard Davis, Steve Matousek, Mars Exploration Program, December 4th, 2024
9. "The Future Mars Communications Architecture". Interagency Operations Advisory Group (IOAG), n.d.
10. Kaiser K. Protecting respondent confidentiality in qualitative research. Qual Health Res. 2009 Nov;19(11):1632-41.
11. User Interviews. "Atomic Research Nuggets." *User Interviews*, 2025, www.userinterviews.com/ux-research-field-guide-chapter/atomic-research-nuggets.
12. "Anonymous Interviews with Mars Relay Network Subject Matter Experts." Interview. Conducted by Michael Newcomb and Brandon Sauer, May 24, 2024 to June 14, 2024.

9. Author Biographies

***Brandon Sauer** received a B.S. in Computer Science from Cal Poly Pomona. After spending much of his 10-year career at JPL as a software engineer in IT and in mission operations, he is currently the task lead for the Mars Relay Operations Service (MaROS) and other related relay software services.*

***Mike Newcomb** received a M.S. in Communications Design from Pratt Institute. Joining JPL after working for digital consulting agencies in New York for seven years, winning multiple industry awards for design innovation. He is a Senior Lead UI Designer in the Human Centered Design group, and has worked as a designer on the Mars Relay Operations Service and Planetary Data System at JPL.*