Technical Support Package (TSP) for

The Mars Relay Network (MRN) Participation Guide

Version 1.0 (Initial Release), 25 August 2025

Questions on the content of this document or to begin discussions on participation in the MRN should be submitted to imrcwg_admin@jpl.nasa.gov.

The contents of this document were assembled at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004), with the participation of the European Space Agency.

Specific contributors to this document include Neil Chamberlain, Brandon Sauer, Eve Pereira, and Roy Gladden from the Jet Propulsion Laboratory, California Institute of Technology; and Silvia Sangiorgi, Peter Schmitz, and Olivier Reboud from the European Space Agency.

Table of Contents

1 Introduction

This technical support package (TSP) is a companion to the Mars Relay Network (MRN) Participation Guide. It contains additional detail and context regarding the architecture and operation of the MRN, from both flight and ground perspectives. Included herein are the following sections:

- **Section 2:** A description of how the Proximity-1 specification has been tailored for use in ultra-high frequency (UHF) transceivers as presently used in the MRN.
- **Section 3:** A description of the supported protocol modes and other supported enterprise specific modes as used in the current MRN, as typically documented in an interface description document (IDD). Also included is a description of how Proximity-1 relay sessions are typically implemented.
- **Section 4:** A summary of the historical context behind the UHF transceivers as presently used in the MRN.
- **Section 5:** A summary of the expected contents of a radio compatibility test plan.
- **Section 6:** A description of the Mars Relay Operations Service (MaROS) and related systems.
- **Section 7:** A description of the file formats used to interface with MaROS.
- **Section 8**: A description of the processes used to coordinate relay activities in the current MRN.
- **Section 9**: A description of the ESA Relay Coordination Office (ERCO), including how it interoperates with MaROS.

2 Enterprise-Specific Tailoring of the Proximity-1 Specification

The Proximity-1 specification [1] [2] [3] [4] was developed, standardized, and documented through the Consultative Committee for Space Data Systems (CCSDS), and adapted for use at Mars for proximity communications by NASA and ESA. As such, some of the CCSDS requirements have been tailored for use at Mars, some features are not supported, and some capabilities that are not described in the specifications have been added to many of the UHF[1] transceivers in use at Mars, specifically within the Electra family of software-defined radios. Supported Proximity-1 capabilities are summarized in Table 2.1. Explanations for the non-compliances and tailoring are provided below the table.

In Table 2.1, two columns are presented. The first column, marked "Minimal", suggests a minimal set of capabilities that are thought to be acceptable for a new relay service provider to implement should the Proximity-1 specification be adopted. The second column, marked "Preferred", indicates the full list of capabilities that are presently implemented in the UHF transceivers at Mars [5], representing capabilities that are equally desired of new relay service providers who support the Proximity-1 specification.

Three indicators are used in the table. A green check indicates that the given capability is fully supported. A yellow check indicates that the given capability is supported, but with features that deviate from or are additive to the Proximity-1 specification. A red X indicates that the capability is not supported.

Table 2.1. List of Proximity-1 Capabilities

✓ (supported), ✓ (tailored), X (not supported).

Proximity-1 Specification	Minimal	Preferred
Modes:		
Full-Duplex	✓	✓
Half-Duplex	X	X
Simplex	✓	✓
Sequence Controlled (Reliable Quality of Service, or QoS)	✓	✓
Expedited Mode (Unreliable QoS)	X	✓
Bitstream (Unreliable)	✓	✓
Standby	✓	✓
Hailing	✓ (tailored)	✓ (tailored)
Coherent Turnaround (Transponder)	X	✓
Timing Services	X	✓ (tailored)
Frequency Band:		
Return-Link (390-405 MHz), with specified Channels 0-7	✓	✓ (tailored)
Forward-Link (435-450 MHz), with specified Channels 0-7	✓	✓ (tailored)
Modulation:		
CW (i.e., No Modulation)	X	✓
FSK	X	X
PSK Residual Carrier	✓	✓

[1] It is noted that the frequencies utilized at Mars for proximity communications are currently implemented in the UHF band. Though this is a feature of the existing Mars Relay Network, it is expected that additional frequency bands will be implemented in the future. Much of the discussion in this section is independent of the frequency band.

BPSK Suppressed Carrier	X	✓
QPSK	X	X
Line Coding:		
Manchester (Bi-Phase)	✓	✓
Non-Return-to-Zero (NRZ)	X	✓
Return-Link Coding:		
Bypass (i.e., no coding)	✓	✓
Convolutional (7,1/2)	✓	✓
Concatenated RS + Convolutional	X	X
LDPC (2048, 1024)	X	✓
Forward-Link Coding:		
Bypass (i.e., no coding)	✓	✓
Convolutional (7,1/2)	✓	✓
Concatenated RS + Convolutional	X	X
LDPC (2048, 1024)	X	X
Symbol Rates:		
Return-Link: 1 kbps – 4096 kbps	✓	✓
Forward-Link: 1 kbps – 4096 kbps	✓	✓

The remainder of this section describes the rationale for deviations from the Proximity-1 specification as implemented in the UHF transceivers as presently used in the MRN (i.e., where red Xs and yellow checkmarks are seen in Table 2.1 above).

2.1 Half-Duplex Mode

In half-duplex mode, a single frequency is used for both return-link and forward-link communications. This necessitates data flow in only one direction at a time and consequently reduces the overall throughput of data (i.e., it is like a "walkie-talkie" radio). Half-duplex mode can be attractive if two-way communication is required where frequency channels are highly constrained (e.g., channels are generally unavailable, channel isolation is an issue, or communication is limited by interference). Since the Electra family of radios, as presently used in the MRN, is equipped with a full complement of forward-link and return-link channels, the radios are constructed with an effective diplexer, and interference is generally not an issue in Mars UHF relay, half-duplex mode was not implemented in the software and firmware for two-way communications.

However, the Electra family of radios is implemented with the necessary hardware to operate in half-duplex (called the *half-duplex overlay*, or HDO). HDO is used when operating in bitstream/simplex mode to bypass the diplexer and thereby reduce circuit losses. Operationally, HDO switching is exercised during a bit error rate (BER) test, which is a non-Proximity-1 capability of NASA proximity radios (see Section 2.7.4).

2.2 Hailing

The default Proximity-1 hailing channel is Channel 1 (i.e., forward-link at 435.6 MHz, return-link at 404.4MHz). The Proximity-1 specification allows enterprise-specific hailing channel frequencies. NASA and ESA use Channel 0 (i.e., forward-link at 437.1 MHz, return-link at 401.585625 MHz) for the hailing channel. For proximity link radio equipment that supports only a single channel (such as in the Mars 2001 Odyssey's C/TT-505 radio), the hailing channel is the same as the working channel. If the proximity link radio equipment supports multiple channels (such as in the Electra family of radios), the specification requires the hailing channel to be distinct from the working channel. However, for operational simplicity, relay sessions in the current MRN are often operated on the same channel as the hailing channel.

Presently, the MRN operates with an "orbiter hailing" paradigm, where full-duplex relay sessions are always initiated by the relay service providers. This is an artifact of the way in which the network was built up over several decades. Though it is technically possible to invert this paradigm to allow relay sessions to be initiated by the relay service users, significant effort would be required to do so within the current MRN.

2.3 Timing Services

The Proximity-1 protocol specifies two timing services: 1) time-tagging of transfer frames and 2) transferring time to a remote asset. The Electra family of radios supports time-tagging but does not support the transfer of time to a remote asset. It is possible to correlate time tags and infer the delay imposed by the channel if time tags are available from both the relay service provider and relay service user and the relationship between time tags and respective clocks (as well as internal delays) is known for each transceiver, though this correlation is not a capability practiced in the current MRN.

2.4 Frequency Channels

The Proximity-1 specification for frequency channels and those implemented in the Electra family of radios are shown in Table 2.2, with differences italicized for Channels 2 and 3.

Table 2.2. Proximity-1 Frequencies As-Specified vs As-Implemented

	Proximity-1 Specification		**As Implemented**	
Channel	**Forward-Link**	**Return-Link**	**Forward-Link**	**Return-Link**
0	437.1	401.585625	437.1	401.585625
1	435.6	404.4	435.6	404.4
2	*439.2*	*397.5*	*437.1*	*391.0*
3	*444.6*	*393.9*	*442.5*	*392.0*
4	435 – 450	390 – 405	443.0	393.0
5	435 – 450	390 – 405	437.1	388.0
6	435 – 450	390 – 405	437.1	405.0
7	435 – 450	390 – 405	437.1	406.0

2.5 Modulation

The Proximity-1 specification permits a simple radio to receive a frequency shift keying (FSK) modulated carrier, but requires it to transmit using phase shift keying modulation (PSK). The advantage of FSK is that it allows non-coherent demodulation (i.e., it does not require a phase lock loop, PLL) at the cost of a 3 dB loss in signal-to-noise ratio. The Electra family of radios does not support FSK modulation because of the associated performance degradation and the widespread availability of PLLs.

Quadrature phase shift keying (QPSK) is a bandwidth-efficient technique for modulating two PSK bitstreams in a single channel. Since bandwidth in the proximity channels is typically not an issue (because the available spectrum is presently sparsely used at Mars), the additional complexity associated with supporting QPSK was unjustified at the time of implementation.

2.6 Channel Coding

Channel coding (or forward error correction) enables a link to operate at lower signal thresholds. Coded links typically operate at higher symbol rates (compared to uncoded links) but provide higher data rates for a given signal-to-noise ratio. The Proximity-1 specifies four coding options:

- Bypass (i.e., no coding)
- Convolutional Coding
- Concatenated Reed-Solomon with Convolutional Coding
- Low Density Parity Check (LDPC) Coding

The Proximity-1 specification calls for the G2 vector (one of two polynomials involved in the encoding of the bitstream) to be inverted in the convolutional code encoder. This effectively inverts every other symbol, which affords a degree of randomization in the coded symbol stream, ensuring a nominal density of symbol transitions for symbol synchronization when the modulation is BPSK with non-return-to-zero (NRZ) line coding.

In the Electra family of radios, the G2 vector is not inverted owing to the use of a commercial Viterbi decoder. A scrambler is employed to guarantee sufficient symbol transitions. Relay service users with radios that do not employ a scrambler to randomize the symbol stream (and instead rely on G2 inversion) should also ensure they use differential encoding to encode the bitstream, as the decoding process can result in sign ambiguity. The Proximity-1 standard specifies three options for scrambling: 1) bypass, meaning no scrambling, 2) Consultative Committee for International Telephony and Telegraphy (CCITT), which supports lossless data compression [6], and 3) Intelsat Earth Station Standards (IESS) [7]. Since neither of these scrambling algorithms are specified by CCSDS, scrambling is not considered required.

For historical development reasons, the UHF transceivers presently used in the MRN currently implement a non-compliant version of the Proximity-1 LDPC code specification [8], which includes a different randomizer and codeword synchronization marker. The Electra family of radios onboard the current MRN orbiters are equipped with firmware that implements the Proximity-1 compliant version of LDPC decoding on the return-link, but this version has not been tested nor exercised operationally.

The current MRN orbiters do not support LDPC coding on the forward-link, due to the associated complexity of the decoder required in the relay service user's radios. Also, the needed forward-link data throughput is operationally much smaller than that needed for the return-link.

Concatenated Reed Solomon encoding with Convolutional codes are optional according to the Proximity-1 specification and are therefore not required. As an alternative to implementing concatenated coding in the physical layer, relay service users may consider implementing a version of such coding in the data link layer, where data in transfer frames is first encoded using Reed Solomon encoding before being convolutionally encoded in the physical Layer.

2.7 Added Capabilities

Additional capabilities of the Electra family of radios are not specified in the Proximity-1 standard are listed in Table 2.3. The columns and marking within this table are consistent with that of Table 2.1.

Table 2.1. Additional Capabilities of NASA Proximity Link Radios

Capability	Minimal	Preferred
Adaptive Data Rates	X	✓
Open-Loop Recording	X	✓
Radiometric Samples (1-way, 2-way)	X	✓
Bit Error Rate (BER) Test	X	✓
Flight (Remote) SW/FW Load	X	✓
Forward / Return Band Swap	X	✓
Heartbeat Telemetry	✓	✓
Extended Telemetry	X	✓

2.7.1 Adaptive Data Rates (ADR)

Proximity radios with the ability to adapt their data rates do so by adjusting symbol rates to optimize the symbol signal-to-noise ratio (SSNR) of the return-link, as bounded by the maximum and minimum rates supported by the radio. For the Electra family of radios implemented onboard the current MRN orbiters, the ADR-supported data rates are from 32 kbps to 2048 kbps, which correspond to symbol rates of 64 ksps to 4096 ksps, with intermediate rates specified in powers of 2. Radios that do not support ADR are characterized as fixed data rate (FDR) equipment. The Electra family of radios support both ADR and FDR. While it is theoretically possible to adapt data rates on the forward-link, this is not done in practice because a low, fixed data rate is normally adequate for most forward-link needs and relay session control frame exchanges.

Presently, an ADR session is controlled by the relay service provider (i.e., the orbiter) proximity radio, which senses the received (i.e., return-link) SSNR and then passes the telemetry to the radio software, where decisions are made about changing data (or symbol) rates, depending on the estimated SSNR value relative to several defined thresholds.

NASA's enterprise-specific ADR control software defines three levels of threshold: a high threshold above which an increase of one data rate can occur, a low threshold below which a decrease of one data rate can occur, and a "deep" threshold below which a decrease of two data rates can occur. Associated with each of these thresholds are persistence thresholds (i.e., the number of SSNR samples) and associated wait durations that must occur before a rate change is allowed.

When a rate change is directed (either to increase or to decrease the data rate), the relay service provider (i.e., the orbiter) issues a communication change request (CCR) on the forward-link, directing the relay service user radio to change to the specified return-link data rate. If the link stalls or times out due to the non-responsiveness of the relay service user, the relay service provider will rehail the relay service user after a specified period referred to as the *rehail gap duration*.

These ADR settings are listed in Table 3.1.

2.7.2 Open-Loop Recording

The Electra family of radios can record a signal in open-loop mode. In an open-loop recording, the relay service provider (i.e., the orbiter) radio directly samples and digitizes the intermediate frequency waveform received and amplified from the antenna without the use of phase-lock loops and demodulation.

The received signal could, for example, be a relay service user's bitstream transmission (i.e., a signal transmitted in unreliable mode) or the radio could simply record the electromagnetic spectrum around a designated frequency. This latter approach is useful to detect electromagnetic interference (EMI) that may be emitted by the relay service provider spacecraft (in which case the recording is called a *sniff*). In either case, extraction of user return-link data or sniff data from the open-loop recording is performed by specialized ground software.

The typical maximum sampling rate of the relay service provider radios is 128 kHz[2], which restricts receivable data rates to 16 kbps or below, or alternatively restricts the bandwidth of a sniff spectrum to $\pm$64 kHz. For coded data transmissions, a data rate of 8 kbps has historically been chosen, with convolutional coding and residual carrier modulation, which results in a first-null bandwidth of 32 kHz. The open-loop sampling frequency must be at least twice the frequency of the coded signal to prevent aliasing.

Open-loop recordings produce data onboard the relay service provider spacecraft at a rate of approximately 250 Mbit/minute, which can result in very large recordings, depending on their duration. Long open-loop recordings are used during critical events (e.g., such as when a new Mars lander performs

[2] NASA's Mars Reconnaissance Orbiter (MRO) can support a sampling rate of 150 kHz.

its entry, descent, and landing, or EDL, at Mars) where link conditions are known to be sporadic and standard, full-duplex communication cannot be relied on. Such events often require special coordination with relay service provider operators to accommodate the resulting very large recording.

Because it is a processor-intensive activity, the Electra family of radios cannot perform an open-loop recording while the radio is performing other functions.

2.7.3 Radiometric Sampling

The Electra family of radios is capable of sampling the phase and power of the received carrier waveform. The collected samples, when packaged with time codes, are referred to as *radiometric* or *phase and power* data. These radiometric data, along with additional information (such as spacecraft ephemeris data), can be used to calculate a spacecraft's position. As such, the collection and processing of radiometric data products is considered a navigation service.

The rate at which radiometric samples are collected is programmable but is generally set to approximately 1 Hz. Radiometric samples can be collected in either one-way mode or two-way mode, the latter requiring the relay service user radio to respond with a coherent turn-around (i.e., in a transponding mode). The collection and generation of a radiometric data product may typically be requested for each relay session (see Table 3.1).

2.7.4 Bit Error Rate (BER) Test

A bit error rate (BER) test is a self-check test that can be used to characterize the performance of a radio receiver. In the Electra family of radios, a small fraction of transmitted power can be fed back (or looped back) to the receiver through the half-duplex switch. At the same time, external noise and EMI, if present, are received by the radio antenna. A known sequence of bits is transmitted and any errors in this sequence are detected in the received bit stream.

A BER test can be useful in assessing the impact of spacecraft EMI on the receipt of a relay service user's data, but only if a baseline of BER tests is established in a controlled, EMI-free environment. In tandem with sniff measurements (as in Section 2.7.2), BER tests can serve as a diagnostic for assessing EMI that may otherwise be difficult to discern during normal relay operations.

2.7.5 Remote Software / Firmware Load

Software-defined radios (SDRs) can have their software and firmware updated in the field (i.e., after installation and launch). In today's MRN, updates may range from simple bug-fixes to the implementation of all new capabilities. The processes of updating software/firmware is different for each participant in the MRN, albeit with common elements. Typically, special spacecraft sequencing is necessary in both instances.

For the Electra family of radios, the process starts by transferring the updated software/firmware images to the host spacecraft's memory. The images are then routed to the radio's non-volatile memory where they are unpacked and installed in the radio. A CRC checksum of the binary image in the radio's EEPROM is performed to indicate if the load was successful prior to installation.

After installation, radio functionality is usually verified extensively, for example by running a BER test and then performing several test relay sessions (i.e., relay sessions planned and executed on a flight timeline, but where the risk of anomalous (or no) data return is accepted).

2.7.6 Forward-Link / Return-Link Band Swap

Consistent with the Proximity-1 specification, full-duplex forward-link transmissions are radiated at frequencies between 435 and 450 MHz and return-link transmissions are radiated at frequencies between 390 and 405 MHz. This provides 15 MHz of bandwidth in each instance. The Electra family of radios can swap their transmit and receive bands using their half-duplex switching circuitry. This allows the radios to either transmit or receive (but not at the same time, as discussed in Section 2.1) anywhere in the 390 to 450 MHz frequency band.

This capability is potentially useful for orbiter-to-obiter crosslink occultation experiments, as well as for reducing circuit losses when operating in simplex mode. The Mars 2001 Odyssey orbiter and the Mars Reconnaissance Orbiter (MRO) performed a series of crosslink experiments in 2007 [9]. In these experiments, Mars Odyssey was transmitting at 437.1 MHz while MRO was receiving in open-loop mode at the same frequency. These experiments were chosen to coincide with ongoing relay sessions with the Mars Exploration Rovers (MERs), Spirit and Opportunity. Doppler shift was extracted from the open-loop recordings using specialized ground tools. Doppler shifts due to the spacecraft trajectories were removed from the data to produce a Doppler residual from which characteristics of the Martian atmosphere and ionosphere (such as electron density profiles) can be inferred. Today, similar experiments continue between ESA's Mars Express and the ExoMars Trace Gas Orbiter [10].

2.7.7 Telemetry

The Electra family of radios produces two types of non-user (or engineering) telemetry: heartbeat telemetry and extended telemetry (also referred to as *monitor data*).

Heartbeat telemetry provides a minimal set of telemetry variables, such as time codes, system voltages, sensor temperatures, and last command history. It is automatically polled by the spacecraft command and data handling (C&DH) system as part of its aliveness monitoring of the radio and is available during most nominal operating modes as well as during quiescent periods (i.e., when the radio is powered on but is otherwise not configured for operation). The timecodes in heartbeat telemetry are necessary to correlate the Electra internal time to the host spacecraft's event time.

Extended telemetry comprises a much richer set of engineering variables and includes heartbeat samples (for a total of ~180 telemetry channels). It is only available when the radio is in its standby mode or during an active relay session. Extended telemetry is not automatically polled by the spacecraft's C&DH system; it must be commanded to be output from the radio (normally at the end of a relay session).

The sampling interval for both types of engineering telemetry is determined by the heartbeat sampling interval, which is typically one or two seconds, though possibly longer in quiescent modes.

3 Typical Contents of an Interface Definition Document (IDD)

This section provides a generalized outline of the contents for a relay service provider interface definition document (IDD). The writing of an IDD is normally the responsibility of the new relay service provider, but collaboration with likely relay service users can be useful during its preparation.

3.1 Typical Contents

The cover material for an IDD should contain the following:

- The front page, indicating the date of the document's release, any document identifiers, version information, and the name of the document preparer.
- All approving signatures, which typically includes the managers of the relay service provider and acknowledgements from the management of known relay service users.
- A distribution list, indicating those who are designated recipients of the signed document.
- A change log for successive releases of the document.
- A table of contents.

The introduction for the IDD should contain the following:

- A description of the document's scope, including a list of relay service users who are expected to utilize the document.
- An outline of what is in the document.
- A list of any relevant reference documents that are useful to help readers understand the document's content or to describe other interface information that is documented elsewhere. Typically, information sourced from other authoritative documents is not repeated in the IDD, except where deviations or tailoring occurs (such as suggested in Section 2).
- Pointers to where relay service users may locate relay service provider trajectories, and the frequency at which these will be updated.

The document should include:

- A description of the physical and data link layers (see also Section 4).
- A description of the hailing and link establishment process (see also Section 3.7).
- A description of enterprise-specific or non-Proximity-1 modes that are supported (e.g., radiometric modes, adaptive data rates; see also Section 2.7).
- A description of the relay service provider data products and their file naming conventions.

The document should describe how the relay services can be accessed and utilized during operations, which includes:

- A list of use cases anticipated to be exercised during operations (as a subset of all the modes supported by the relay service provider, as above).
- Nomenclature for how relay opportunities will be identified during relay planning and coordination activities (see Section 6.3.2 for a discussion on overflight identifiers).
- A description of how the relay service provider's relay antenna will be pointed during a relay session.
- A description of any constraints on how relay sessions may be scheduled.
- The list of all parameters that must be negotiated between relay service users and the relay service provider (as suggested in Section 3.2).
- The list of all parameters that are considered fixed for all relay sessions (as suggested in Section 3.3). These parameters should help describe how relay services will be provided.
- A description of how the relay planning processes function for the relay service provider, which is expected to conform to established processes as much as possible (see Section 8).
- A description of how forward-link data products will be handled, including any limitations that are applied to the size and number of products (see Section 3.4). This should include a description of the end-to-end data flow and all data interfaces.
- A description of how incomplete return-link data products are handled (see Section 3.5). This should include a description of the end-to-end data flow and all data interfaces.
- A description of what data will be reported as part of the accountability process (see Section 3.6)

3.2 Negotiated Relay Session Parameters

The following table lists a series of parameters that are currently common across the MRN orbiters. These parameters are specified by the relay service user operators when requesting a relay session and confirmed by the relay service provider operators as part of the Strategic and Tactical Planning Processes (see Section 8.1).

Table 3.1: Examples of Negotiated Relay Session Parameters from the Existing MRN

Parameter	Description
Spacecraft ID	Decimal representation of the Spacecraft ID to use in the hailing sequence (typically as specified by the Space Assigned Numbers Authority, or SANA).
Hail Start Time	Time (in spacecraft ephemeris time, SCET) when the orbiter should start hailing the relay service user for the session.
Hail Duration	Duration of the relay session.
Relay Link Type	Indicator for if the relay session should be configured as duplex (i.e., data should be exchanged in both directions), simplex return (i.e., data should be exchanged only in the return-link direction), or simplex forward (i.e., data should be exchanged only in the forward-link direction).
Relay Protocol	Proximity-1 mode (reliable) or Raw Data mode (unreliable).

Parameter	Description
Coherency	Specified as coherent or non-coherent for the relay session. This is typically specified as non-coherent unless radiometric samples are collected.
Forward-Link Data Rate	The data rate to use in the forward-link direction (in kbps).
Return-Link Data Rate	The data rate to use in the return-link direction (in kbps). If an ADR algorithm is to be used, this is typically the data rate at which the relay session should start after the hailing cycle concludes.
Radiometric Mode	Indicator to specify if a radiometric product is to be generated during the relay session.
G2 Inversion	Inversion of the G2 polynomial in the convolution encoder, if applicable, typically specified as inverted or not inverted.
Forward-Link Channel	Frequency to use for the forward-link after the hailing cycle concludes, typically specified as a Proximity-1 frequency channel (see also Section 2.4).
Forward-Link Carrier Mode	Indicator to specify the carrier mode to use on the forward-link after the hailing cycle concludes, typically specified as residual or suppressed.
Forward-Link Encoding	Indicator to specify the encoding to use on the forward-link after the hailing cycle concludes, typically specified as bypass (i.e., no coding) or convolutional.
Frame Size	Indicator to specify the forward-link frame size (in bytes).
Window Size	Indicator to specify the number of frames that may be transmitted on the forward-link that may go unacknowledged before a retransmission of those frames occurs.
Return-Link Channel	Frequency to use for the return-link after the hailing cycle concludes, typically specified as a Proximity-1 channel (see Section 2.4).
Return-Link Carrier Mode	Indicator to specify the carrier mode to use on the return-link after the hailing cycle concludes, typically specified as residual or suppressed.
Return-Link Decoding	Indicator to specify the decoding to use on the return-link after the hailing cycle concludes, typically specified as bypass (i.e., no coding), convolutional, or LDPC.
Sampling Divisor	Indicator for the radiometric sampling divisor to use if radiometric data is to be collected during the relay session (per “Radiometric Mode” above). This value determines the rate at which radiometric samples are collected.
Carrier Loss Timeout	The duration that the carrier signal may be lost before the radio should attempt to rehail.
Rehail Gap Duration	The duration between rehail attempts, should a response not be received.
	ADR-Specific Parameters
Enable Flag	Specifies whether an ADR algorithm should be used on the return-link during the relay session.
Maximum Return-Link Data Rate	Indicates the highest return-link data rate (in kbps) that should be attempted when ADR is enabled.
Minimum Return-Link Data Rate	Indicates the lowest return-link data rate (in kbps) that should be allowed when ADR is enabled.
SNR Sampling	Indicates the sampling period (in seconds) for estimating the SNR when ADR is enabled.
SNR Step	The step (in dB) that should be used when determining the number of rate changes (1 or 2, usually 1) when the “High SNR Threshold" is exceeded.
High SNR Threshold	The threshold (in dB) above which an increase of at least one data rate should occur.

Parameter	Description
High Persistence Threshold	The number of consecutive high SNR samples that should occur before a data rate increase should occur.
High Wait Duration	The time (in seconds) after a "High Persistence Threshold" data rate increase occurs before another data rate change is permitted.
Low SNR Threshold	The threshold (in dB) below which a decrease of one data rate should occur.
Low Persistence Threshold	The number of consecutive low SNR samples that should occur before a data rate decrease of one data rate should occur.
Low Wait Duration	The time (in seconds) after a "Low Persistence Threshold" data rate decrease occurs before another data rate change is permitted.
Deep SNR Threshold	The threshold (in dB) below which a decrease of two data rates should occur.
Deep Persistence Threshold	The number of consecutive deep SNR samples that should occur before a data rate decrease of two data rates should occur.
Deep Wait Duration	The time (in seconds) after a "Deep Persistence Threshold" data rate decrease occurs before another data rate change is permitted.

Future Consideration: Self-Negotiated Relay Sessions

Relay sessions in today's MRN are selected and scheduled by the operators of the Mars rovers, who choose the time of each relay session and designate the radio parameters to use during the relay sessions. As conditions at the site of the Mars rovers may be different between relay sessions, the responsibility is left to the rover operators to choose a configuration that best meets their needs. Most of the time, relay sessions are configured with ADR turned on using a default set of parameters. However, at times it may be more advantageous to schedule a relay session as an unreliable session or with fixed data rates. It is left to the rover teams to select a correct configuration for the relay session, and the orbiter operators typically implement what has been requested even if it does not match expectations.

In the future, it is anticipated that relay sessions might be self-configured, similar to how modern-day cellphones negotiate and configure their connection with a nearby cellphone tower. However, so long as the available relay services remain over-constrained, it is expected that the relatively hands-on method described here of selecting relay session parameters will remain necessary to ensure a correct configuration is selected that best meets the needs of the relay service users.

3.3 Fixed Relay Session Parameter

The following table lists a parameter that is common across the MRN orbiters. This parameter is specified by the relay service user operators and is generally fixed for all users for all relay sessions, unless special arrangements are made to modify it. Adjusting this parameter is not typically part of the negotiation process. Indicating this parameter and others like it in an IDD is useful to indicate features of the relay services that may need to be known to properly utilize these services.

Table 3.2: Example of a Fixed Relay Session Parameter

Parameter	Value
No More Data Flag	A flag that indicates if the relay session may be terminated by the relay service user when it is determined there is no more data to send, typically specified as "off" but may be set as "on". This feature is not presently supported by the Electra family of radios.

3.4 Forward-Link Data Operations

Forward-link data is handled by the relay service providers as distinct files, the contents of which remain unknown to the providers.

Each relay service provider specifies the maximum number of forward-link files that may be transferred during a single relay session and the maximum size of those files, both individually and collectively. These restrictions are a function of the nature of the relay service provider's (i.e., the orbiter's) file handling system, the available onboard storage, how this storage space is shared among all the relay service users, the persistence of the data in this storage after a relay session concludes (regardless of the success of the data transfer during the relay session), etc.

The end-to-end topology for how forward-link data is transferred via the relay service provider would typically be expected to include MaROS as a portal from which forward-link data can be received by the relay service provider operators from the operators of all relay service users (see Section 9 for an exception). The flow of such data through the ground system of the relay service operator should be understood, including how the forward-link data is transmitted to the relay service provider spacecraft via ground stations and how it is handled onboard the relay service provider spacecraft prior to being transmitted to the relay service user spacecraft.

These details should be documented in the relay service provider's IDD. In principle, a general description is sufficient, as the relay service users usually do not need to know the specific details of this internal topology.

3.5 Return-Link Data Operations

Return-link file handling generally differs by deep space tracking network and the relay service provider. Presently, return-link data is handled differently by each of the relay service providers, with some packaging data as CFDP products [11] and others transferring the data as Advanced Orbiter Systems (AOS) frames [12].

Each relay service provider specifies the maximum amount of data that can be stored onboard the orbiter, which typically has implications for data persistence, retransmission, etc. They may also control this on a per-user basis or aggregate the data in a single data repository onboard the orbiter, in which case it may be possible for data stored for one relay service user to impinge upon data stored for another.

The end-to-end topology for how return-link data is transferred via the relay service provider should be understood, including how the data is handled onboard the relay service provider spacecraft after it is received from the relay service user spacecraft, how it is downlinked from the relay service provider spacecraft to ground stations, and how it is handled on the ground prior to being transferred to the operators of the relay service user spacecraft. This topology may include MaROS as a portal through which return-link data can be transferred by the relay service provider operators to the operators of all relay service users.

These details should be documented in the relay service provider's IDD. In principle, a general description is sufficient, as the relay service users usually do not need not know the specific details of this internal topology.

3.6 Relay Accountability

The operators of both relay service provider and relay service user spacecraft are expected to report on the performance of each relay session to MaROS from data extracted from their respective radio's engineering telemetry. This is normally done within 24 hours of a relay session on a best-efforts basis. There are two such reports that are typically submitted: an Overflight Scorecard (or just Scorecard) which provides a summary of performance data for the relay session, and an Overflight Performance Assessment File (OPAF) which reports values that vary as a function of time through the relay session. See Sections 7.3, 7.4, and 8.4.

Generally, it is left to each operator to decide what data to report. However, for reference, a minimal set of Scorecard and OPAF items are described here.

Summary data published in a Scorecard may include:

- The actual start time and end time of the relay session, as executed onboard the spacecraft. This may differ from what was negotiated during the Strategic and Tactical Planning Processes (see Section 8.1).
- The period(s) of time during which the relay session link is considered "closed", which should reflect the number of times that a rehail had to occur throughout the relay session.
- The total amount of return-link data transferred during the relay session. This value is typically compared to the predicted amount of data that was expected to be transferred, a key indicator of the success of the relay session.
- The total amount of forward-link data transferred during the relay session. This value is typically compared to the predicted amount of data that was expected to be transferred, as represented by the forward-link data submitted to the relay service providers.

- The data rates that were used during the relay session. For the return-link data rate during ADR relay sessions, this is typically the starting data rate.

- Any other key mode settings. Typically, a report of the relay session parameters, as suggested in Section 3.2, would be reported to indicate the as-executed configuration of the relay session.

Time-dependent data published in an OPAF may include:

- The actual return-link data rate achieved during the relay session, particularly for ADR relay sessions.

- The number of successfully received return-link bits.

- The number of successfully transmitted forward-link bits.

- An indicator of the carrier lock.

- A report of the measured signal-to-noise ratio.

This data can reflect the overall health of the relay session, and may point to issues with EMI, terrain reflectivity, etc.

Additional telemetry and session variables, such as a count of retransmission requests, a count of the number of out-of-sequence frames that were detected, radio temperatures, radio voltages, etc., may also be published.

3.7 Proximity-1 Relay Session Description

The process of establishing a reliable or expedited mode Proximity-1 relay session is referred to *hailing*. Hailing is performed in both full-duplex and half-duplex relay sessions; it is not performed in simplex (i.e., bitstream) relay sessions. The initiator of the hailing process is referred to as the *caller* or *local asset*, whereas the receiver of the hail is referred to as the *responder* or *remote asset*. The Proximity-1 specification allows for either the orbiter or the lander to be the caller or responder, but by convention in the MRN, orbiters have always been the caller in relay sessions to date.

Prior to the initiation of the relay session, the responder sequences its radio to be powered on and “listening” for a hail signal (referred to as *waiting for hail*) on the caller’s forward-link channel, previously agreed to during the Strategic and Tactical Planning Processes as described in Section 8.1. While waiting for the hail, the responder is in receive-only mode and is not transmitting. The default Proximity-1 hailing channel is Channel 1 (435.6 MHz), but the specification also allows the use of enterprise-specific hailing channels; the MRN presently uses Channel 0 (401.585625 MHz) as the default hailing channel. Hailing is normally performed at a low data rate (nominally 8 kbps) using unencoded, residual carrier modulation. However, compatibly-configured transceivers can hail on other channels and/or use other modulation and coding schemes (including, for example, suppressed carrier modulation).

Hailing is initiated by the caller radiating a carrier signal to enable carrier lock, then by issuing an idle bit stream to achieve symbol lock. This is followed, without response from the responder, by the caller radiating hail directives in a supervisory protocol data unit (PDU) or transfer frame, followed by idle bits to allow the directives to be received and processed. The supervisory PDU or transfer frame header contains the spacecraft ID (SCID) of the remote asset being hailed. If the SCID in the transfer frame does not match the SCID of the remote asset, then the frame is marked as invalid, and the remote asset should not respond. Similarly, if the remote asset does not receive the hail (e.g., because it is not in line-of-sight or because the transmission is corrupted), then the hail sequence repeats after a specified interval and continues until the hailing process either times out or the remote asset communicates (on the return-link frequency) that a valid transfer frame was received.

Once a valid supervisory transfer frame is received and its reception is communicated to the caller, both assets reconfigure their radios to enable data services for the main session. The directives communicated to the responder comprise the following major categories (see [3], Annex B for a complete description):

- Transmitter Parameters
- Receiver Parameters
- Extension Parameters
- Control Parameters

The Transmitter and Receiver Parameters contain the frequency channel, data rate, modulation, and data encoding, available in the Standard Proximity-1 Table, for the respective radios. Additional radio parameters (such as suppressed carrier modulation and data rates above 512 kbps) are available in the Extended Proximity-1 Table. Control Parameters set duplexity (full-duplex, half-duplex, or simplex) and allow the remote asset to terminate the relay session through the *Remote No More Data* command[3]. Control Parameters are also used to set the number of time samples that are collected during timing service sessions. Once the respective radios are properly configured for the main session, the process of carrier and idle transmission is repeated to establish the link at the new parameters, and the main session begins.

During a relay session, it is possible that the link may drop for a variety of reasons (e.g., such as due to terrain occlusion). In such instances, the caller will reinitiate the hailing process after a prescribed pause. This process repeats until the link is reestablished or the session ends. Relay user radios that support adaptive data rates will be commanded to change the return-link rate through a communications change directive issued by the relay provider radio. The session remains active throughout this process, unless the link drops during a rate change, in which case the radios rehail as previously described.

Once a relay session is terminated, the return-link data received from the relay service user and associated engineering telemetry are made ready to be transmitted back to Earth. During the data collection process, relay session data is packetized (or framed) and stored in the relay spacecraft's onboard memory. In principle, transmission to Earth is possible as soon as the session concludes, but the spacecraft's attitude, antenna pointing, prioritization of data products to be downlinked, and the availability of a deep space antenna on Earth to receive the data may delay the downlink. Additionally, the downlink process can become interrupted (for example by a fault in the deep space antenna), requiring the retransmission of missing data. MaROS return-link latency predictions are based on the expectation of a nominal downlink. When a downlink is partially received in the Earth ground system, relay service user operators are usually

[3] The UHF transceivers in the MRN do not implement session termination using the no more data command. Instead, relay sessions run for the planned duration and, if the relay user has no more data to send (which is rare), fill bits are sent on the return-link until the session otherwise concludes, or the relay user can stop transmitting and remaining receptive to a hail.

furnished with a partial product until retransmission is completed, at which time the complete data is delivered to the relay service user operators. Retransmitted data products can sometimes take days to complete.

In special circumstances, such as during the EDL of a Mars lander, the delay associated with collecting all relay data from a session prior to downlinking it is often deemed unacceptable. In such circumstances, *bent-pipe* operation is typically used to return the data as it is received. For a relay service provider to provision bent-pipe relay, it is necessary to be able to simultaneously point the proximity link antenna towards the relay service user spacecraft and the direct-to-Earth (DTE) antenna to Earth, as well as have the necessary software to flow the data to the orbiter's downlink system in realtime or near-realtime. Currently in the MRN, only the Mars Reconnaissance Orbiter (MRO) and the Mars 2001 Odyssey orbiter have a bent-pipe capability. Because this capability requires the reconfiguration of the spacecraft and special planning, it is only used during critical events where there is an urgency to receive the relay data.

4 Historical Context for the Design of the MRN's Proximity Radios

This section provides a summary of the relay configurations as presently implemented in the MRN.

4.1 Radio Hardware and Software

NASA-sponsored UHF radios for Mars proximity communications are derived from portable radios originally developed for terrestrial military communications. As such, these radios are relatively compact, lightweight, power efficient, and reliable. Early proximity radios, such as the C/TT-505 [13], were hardware defined. They offered a limited selection of frequency channels, data rates, and modulation and encoding schemes, and could not be updated in the field. Later proximity radios, such as the Electra family of radios [14] [15] are software defined, supporting updates to their firmware and software that enable them to be reconfigured in the field to perform new functions. The Electra family of radios also implements the full set of Proximity-1 extensions, facilitating (for example) higher data rates and additional frequency channels for frequency agility[4]. Frequency agility can be helpful (for example) to avoid EMI from other sources (including the spacecraft itself). Historically, UHF proximity radios typically utilize solid state power amplifiers (SSPAs) that generate between 5-10 watts of radio frequency (RF) power, in line with their military precursors.

4.2 Antennas

NASA-sponsored UHF proximity link communication antennas have historically been low-gain, omni-directional antennas. Omni-directional antennas have the advantage of performing relatively well over a wide range of aspect angles and don't need to be pointed or steered. Adding a gimbal assembly to steer an antenna (or adding an electronic scanning capability) adds significant complexity in terms of mass, power, and operations; and so therefore have been avoided on Mars-bound spacecraft.

The current orbiters and rovers in the MRN carry one or more circularly polarized UHF antennas that support the full Proximity-1 frequency band of 390 to 450 MHz. Since a radio is typically connected to one antenna at a time during a relay session, a diplexer is required to separate return-link and forward-link channels in full-duplex mode (a diplexer is not required in half-duplex mode or in simplex mode).

NASA's Mars orbiters, landers, and rovers have historically employed either circularly polarized or linearly polarized UHF antennas. A common choice for a circularly polarized antenna is the helical antenna, which is usually either a single helix spiral or two interlaced spirals fed with a 90 degrees phase difference, referred to as a *quadrifilar helix*. The orbiters in the MRN radiate with right-handed circular polarization. The use of a linearly polarized antenna on the lander results in roughly half as much power received by the lander, compared to a circularly polarized lander antenna (the same is true of the return-link).

[4] Frequency agility in this context means the ability to select a fixed frequency for the return-link or the forward-link for a given relay session. It should not be confused with an ability to adjust the working frequencies of a relay session dynamically in response to local conditions.

Additionally, a common implementation for a linearly polarized antenna is a simple monopole, which has a null in the zenith direction. Despite this limitation, early NASA rovers (Sojourner and the Mars Exploration Rovers) all used monopole antennas whereas more recent landed missions, including the Curiosity and Perseverance rovers, carried helix antennas.

4.3 Propagation Channel

The communication channel between radios comprises the antennas and the intervening propagation channel. This channel determines how much signal power is radiated in a particular direction, as well as the frequency response. The relative geometry between two radios plays a large part in the coupling from one antenna to another and the overall link performance. To first order, the received coupled power (assuming both antennas are in free-space) is predicted by the Friis equation [16]. This depends on the gain of the antennas, the transmitted power, and the path loss, which scales 20 dB per decade increase in the slant range (R). This is true at any frequency, and at UHF (around 400 MHz) the path loss scales as

$$PL(400\text{ MHz}) = 62 + 20 * \log_{10}(R(\text{km})) \quad (\text{dB})$$

whereas at X-band (around 8 GHz), the path loss scales as

$$PL(8\text{ GHz}) = 89 + 20 * \log_{10}(R(\text{km})) \quad (\text{dB})$$

4.3.1 Antenna Designs

Omni-directional antennas tend to have the same low gain (i.e., roughly a few dB) independent of frequency, therefore an X-band link using omni-directional antennas would have a disadvantage of 27 dB in coupling compared to UHF. This same reasoning could argue for a lower frequency than UHF, but lower radio frequencies necessitate the implementation of larger antennas, making accommodation on a spacecraft more of a challenge. The choice of UHF for Mars proximity communications therefore represents a compromise between efficient link performance and the size of the antennas that could reasonably be accommodated on the spacecraft.

4.3.2 Overflight Durations

Overflight durations depend on orbital mechanics and the time that a spacecraft can be oriented to accommodate relay activities. Historically, these have ranged from about 10 minutes for MRO, which is in a circular orbit with an altitude of about 250 km, to 30 minutes for MAVEN, which is in an elliptical orbit with a periapsis of about 180 km and an apoapsis of about 4500 km. Though MAVEN has view periods that are longer than 30 minutes, it is limited to 30-minute relay sessions due to limited available power. Relay session durations for overflights that occur at low elevation angles (from the viewpoint of a lander) typically have shorter durations due to terrain obscuration.

4.3.3 Data Rates

Generally, it is the relay user's responsibility to calculate the link budgets for relay sessions, typically in an attempt to estimate the total amount of data that may be transferred during a relay session. This is done using relay provider ephemeris data and other information typically documented in an IDD (see Section 3). A link budget based on the free-space Friis equation is normally sufficient for most planning and prediction purposes. The amount of data that can be returned during a relay session will depend on several factors previously discussed, as well as the data rates(s) that can be supported during the relay session.

Proximity radios that can adapt their data rates as a function of local conditions (i.e., ADR, see Section 2.7.1) return significantly more data than radios that have a fixed data rate for the duration of the relay session. In a fixed data rate relay session, the data rate is typically selected so that the link closes at the longest slant range and with the most unfavorable antenna pointing. Radios that support the use of an adaptive data rate algorithm adjust their data rate automatically to optimize data throughput for a given time during the relay session.

For example, NASA's 2001 Mars Odyssey orbiter and ESA's ExoMars Trace Gas Orbiter (TGO) both operate in circular orbits that provide a view of landed spacecraft for about 15 minutes during an overflight. Odyssey supports fixed data rates up to 256 kbps but usually runs at lower data rates to maximize data return over the relay session. Thus, a 15-minute relay session at 256 kbps will return approximately 230 Mbit of data. In comparison, TGO supports adaptive data rates up to 2048 kbps, and can therefore return over 1 Gbit during a 15-minute overflight. These example data volumes depend on the modulation and encoding schemes (which are different for the respective orbiter radios) as well as data rate(s) and overflight geometry. MaROS supports the prediction of return-link data volumes and associated latencies (see Section 6.5).

4.3.4 Terrain Scattering

Terrain affects UHF signal propagation in two principal ways. For lander-orbiter elevation angles below the terrain horizon, the UHF signal is effectively blocked, preventing the link from closing. Overflight predictions can take such blockage into account by computing a terrain mask as a function of azimuth angle relative to the lander, or more simply by assuming the link will not close below a fixed elevation angle corresponding to the maximum elevation of any in-view terrain.

For lander-orbiter elevation angles above the terrain horizon, the link can become impacted by multipath reflections and (to a lesser extent) by diffraction. Multipath reflections are prevalent below approximately 20 degrees above the terrain horizon. Such reflections can be additive or subtractive; it is the subtractive reflections that cause destructive interference with corresponding transient fading of the UHF signal. Without accurate terrain and scattering models, it is difficult to predict when such fades will occur. However, because multipath fades are fleeting and occur at lower elevation angles where data rates tend to be lower, their overall impact on data volume throughput is relatively low. Relay planners can therefore choose to ignore multipath fades or accommodate their impact through appropriate margining.

4.3.5 Electromagnetic Interference (EMI)

Another factor that can impact the amount of data returned during a relay session is EMI generated by the spacecraft. EMI can arise from several sources, such as the actuators of instruments installed on the spacecraft or switching power supplies. Because of the proximity of such sources to the UHF antenna, coupling factors between EMI sources and the UHF antenna can be significantly larger than coupling between the user and provider UHF antennas, with the result that low power EMI emissions can add significantly to received noise power. EMI is typically a more significant problem on the return-link (i.e., due to the relay service provider spacecraft) because the data rates are usually higher than those of the forward-link and therefore the thresholds[5] are higher.

New participants in the MRN should proactively mitigate potential EMI issues prior to launch through careful design and verification during integration and test. EMI issues that persist or arise during flight operations can sometimes be mitigated through operational work arounds and other methods, though this is not ideal. For example, a noisy instrument can be powered down or put in a quiescent mode during a relay session to reduce EMI, but at the cost of not being able to collect instrument data. In addition, relay sessions can be operated at lower data rates to compensate for EMI, but at the cost of reducing relay data throughput. Radios with frequency agility can also select channels that minimize the impact of EMI. Software-defined radios can also be configured to implement digital filters to notch or otherwise limit received EMI. While these various approaches can be used to mitigate EMI to some degree, it is generally preferable to start with a well-designed and EMI-free spacecraft from the outset.

[5] A threshold is the minimum signal power that must be received to sustain the link by limiting bit errors. Low data rates (which the forward-link typically operates at) have low thresholds and therefore larger SNR margin over ambient noise.

5 Generic Radio Compatibility Test Plan

This section provides a generalized outline of the contents of a test plan used to confirm radio compatibility between participants in the MRN. The writing of a test plan is normally the responsibility of the new entrant, whether a relay service provider or a relay service user, but it naturally requires collaboration with participant test teams during its preparation. A test plan is distinct from a test procedure in that it does not go into the specifics of configuring equipment, the methodologies for recording test data, and the criteria for pass/fail assessments. A test plan should speak only in general terms of pass/fail criteria and should provide a list of all applicable test procedures.

The cover material for a test plan should contain the following:

- The front page, indicating the date of the document's release, any document identifiers, version information, and the name of the document preparer.
- All approving signatures, which typically includes the managers of the new entrant and acknowledgements from the management of other related parties.
- A distribution list, indicating those who are designated recipients of the signed document.
- A change log for successive releases of the document.
- A table of contents.

The introduction of the test plan should contain the following:

- A description of the document's scope, including the list of test participants who are expected to utilize the document.
- An outline of what is in the document.
- A list of any relevant reference documents that are useful to help readers understand the document's content or to describe other useful information that may be documented elsewhere. Typically, information sourced from other authoritative documents is not repeated in the test plan, except where deviations or tailoring occurs.

The document should include a description of the test objectives, including:

- An overview of the test schedule.
- An overview of the test venue.
- A summary of what the testing intends to accomplish, which may be characterized as a set of minimal goals for the testing (sometimes referred to as an *incompressible test list*), a larger set of baseline goals, and an extra set of any secondary goals (sometimes referred to as *stretch goals*).
- An overview of major test configurations.
- References to any necessary preceding tests and/or follow on tests.

The document should include a description of any major test preparation milestones, including:

- The generation of relay service requests (typically via MaROS), reflecting the tests to be exercised.
- The retrieval of these relay service requests from MaROS and the generation of the needed command sequences.
- The generation of any relay service user forward-link test files.
- The generation of any relay service user return-link test files.
- A description of verification activities that confirm that appropriate command sequences and relay service user data can be loaded into the appropriate testbeds.
- Logistics for the transport of test equipment to the test venue.
- A description of any test readiness review that will be planned.

The document should include a description of how the test venue will be configured, including:

- A description of the test venue and facilities that will accommodate the testing.
- A description of any applicable Missions Operations Center (MOC) interfaces and interactions.
- A high-level description of major hardware, software, and network components, with associated interfaces, including a system-level block diagram.
- A description of the applicable end-to-end test interfaces and how they will be exercised in the testing.
- A description of the specific hardware and software configurations of the testbeds.
- A description of how communications will occur between test participants, especially if all test participants are not co-located during the testing.
- A description of real-time telemetry flows and monitors.
- A description of the test timeline and the reference time (e.g., UTC).
- A summary of applicable test constraints and any considerations where the manner in which the testing occurs is significantly different than how it would execute in flight operations (referred to as *test-as-you-fly exceptions*).

The document should include an outline of how the tests will be executed, including:

- The test matrix, with indicators of which tests are considered minimal, baseline, and secondary.
- A list and summary description of all applicable test procedures.
- An outline of organizational responsibilities for all test participants.
- A schedule and nominal timeline for the execution of the tests, which should include time to review test results and to perform any applicable end-to-end test functions. The schedule should also include extra time for troubleshooting and retesting, as needed.

- A list of applicable constraints or idiosyncrasies that may impact the execution of the tests.
- A description of specific checks that will verify the test results, such as any analysis of engineering telemetry and transferred data products.
- A description of the pass or fail criteria for each test, specified as one of 1) pass, where the test goal is significantly met and relevant functionalities are mostly verified, 2) partial pass, where the test goal is partially met and some relevant functionalities are verified, or 3) fail, where the test goal is not significantly met and/or relevant functionalities are mostly not verified.
- A description of the anomaly reporting process in the event of test failures and how those failures will be dispositioned during testing.
- A schedule of any meetings or briefings that will be used to review the test plans and the results.

6 MaROS and Related Systems

In addition to MaROS, various other sub-services and systems are utilized to facilitate or enhance the relay coordination activities between the relay service users and the relay service providers, as described here. This section provides a high-level description of these services along with some information on how they are configured for new MRN participants.

6.1 MaROS Ongoing Development

MaROS and related systems are considered "living" software services and are expected to continue to evolve over time to remain current with new software technologies, cybersecurity requirements, and the operational needs of the MRN. Other changes may be needed to add new capabilities or to modernize its infrastructure.

Significant changes, including testing and release dates, are communicated to the MRN participants at IMRCWG meetings or via email. Typically, one major or minor MaROS release is organized each year, with a user acceptance test (UAT) period on the MaROS Testbed environment (MaROS-TB) lasting one month prior to the official installation of an update to the MaROS Operations environment (MaROS-Ops). These releases may include updates to any part of the MaROS ecosystem, its external interfaces, and in the attending documentation. Backwards compatibility at the interfaces is always attempted though it is not always possible.

Downtimes associated with these installations typically last at most six hours and are communicated well in advance. For critical issues, the team may implement and release patches to the services or clients at any time, which may be communicated with significantly less notice than planned releases.

The MaROS operations team assists all MRN participants as needed during these updates.

6.2 Gaining Access to MaROS and Related Systems

MaROS and related systems (other than the public-facing "Eyes on the MRN", see Section 6.7) are all secured behind firewall of NASA's Jet Propulsion Laboratory (JPL). Three layers of security are applied to control access to MaROS, related systems, and its database:

1. New MRN participants must specify one or more Internet Protocol (IP) addresses from which relay services will be accessed, typically at their home institutions. The MaROS operations team then coordinates with cybersecurity personnel to verify the addresses and add them to a whitelist, which enables access to MaROS from those addresses. All MRN participants bear the responsibility of maintaining their own secure operations environment with machines that have fixed IPs.
2. All MaROS users must have valid Lightweight Directory Access Protocol (LDAP) credentials to authenticate with the service. Again, the MaROS operations team assists with on-boarding

personnel to gain this LDAP access. For automated access to MaROS, an "application user" can be configured to communicate with MaROS via the use of a rotating token scheme for authentication.

3. Certain operators of each spacecraft represented within MaROS are identified as "mission administrators". These can determine which other users are associated with that spacecraft and manage user-level permissions to publish the various types of data to the MaROS database.[6] These mission administrators bear the responsibility of ensuring that only authorized personnel are granted access to MaROS and related services from their operations environment, including managing the removal of permissions for personnel who no longer require them.

The Metrics Predict Executive (MPX, see Section 6.4), the Relay Telecom Predictor (RTP, see Section 6.5), and the Where's My Relay Data (WMRD) service (see Section 6.6) can generally only be accessed via MaROS.

All participants in the MRN are expected to gain direct access to SPS (see Section 6.4) for the purpose of publishing their spacecraft ephemeris data. SPS is the primary repository for spacecraft ephemeris predictions and other data needed to interface with NASA's Deep Space Network (DSN). As such, it is required of participating missions to gain access to this service, which is gained separately from access to MaROS. The onboarding information mentioned in step 2 above for gaining LDAP access is applicable.

The MaROS operations team stands ready to aid new participants in gaining access to these systems and to configure these services for use. Contact information for the MaROS operations team will be provided to new participants after a formal agreement to join the MRN has been established. In most cases, two months of lead time is required to configure access to these services for new organizations and users.

Access to MaROS-TB is provided concurrently with access to MaROS-Ops.

6.3 MaROS

6.3.1 MaROS Description

MaROS serves as the center point of the relay coordination ecosystem [17], as discussed in the MRN Participation Guide. Refer to that guide for an overview of the MaROS ecosystem, for a general description of the native capabilities of MaROS, and for a summary of the methods that may be used to interface with MaROS. Refer to Section 7 herein for a description of the file formats used when publishing data to or extracting data from MaROS. Refer to Section 8 herein for a description of the relay coordination processes where MaROS is utilized.

---

[6] The act of publishing data into the database is restricted via application-level user permissions. Nearly all relay planning data published into the MaROS database is not considered sensitive, and thus can be viewed and downloaded by any other MaROS user. However, forward- and return-link data is generally considered proprietary and access to this data is also restricted at the application-level.

6.3.2 Configuring MaROS

Due to the generic and data-driven nature of the MaROS ecosystem, there is a small amount of configuration that must occur to set up each new spacecraft (either as a relay service user or a relay service provider) in MaROS and related systems. New MRN participants must work with the MaROS operations team to define the following:

- A 3-letter alphanumeric identifier for the new spacecraft (e.g., “MRO”), used by MaROS when automatically generating unique identifiers for each relay opportunity, which is used to associate all the data pertinent to those opportunities. These “overflight IDs” include both partners in a relay opportunity (relay service provider listed first), the year, day of year, and overflight number on that day of year, and looks like “MRO_M20_2025_043_02”. This overflight ID is intended to be easily recognized and interpreted by the human operators of the MRN.
- A SCID for the new spacecraft, typically as specified in the SANA registry. This SCID is primarily used by MPX to retrieve spacecraft ephemeris predictions from SPS when calculating view periods.
- The set of “linked” spacecraft with which the new spacecraft will perform relay activities.
- One or more “mission administrators” for the new spacecraft, specific operations personnel with access to MaROS (as in Section 6.2) that will manage the new participant’s MaROS users, define spacecraft-specific parameters, and configure other settings that describe how the new spacecraft behaves in the network.

Once the spacecraft has been added to MaROS using this information, the assigned mission administrators can then perform the following actions:

- Define and publish a Spacecraft Configuration File (SCF, see Section 7.11) that contains the parameters that can or must be defined for each relay session. As a relay service provider, the definition of these parameters is critical for users to request relay services from the provider spacecraft. As a relay service user, relay service requests must contain parameters defined by the relay service provider, but relay service user-specific parameters may also be defined. These parameters may be updated at any time via the publication of a new SCF to MaROS or via the MaROS graphical user interface (GUI). The SCF should also contain various system parameters and other settings that MaROS needs to calculate latencies, identify planning conflicts, handle forward-link data, and more. For a relay service provider, the SCF should also contain information about the onboard storage for the spacecraft, which is used when calculating return-link latencies.
- Manage the users that should have access to publish data to or extract data from MaROS on behalf of the new spacecraft. This function is only available via the MaROS GUI.
- Manage email notifications for these users. A variety of email notifications are generated by MaROS when certain service events occur as part of the relay coordination processes (see Section 8). A user may also update their own notification preferences as well. This function is only available via the MaROS GUI.

Once these configurations are in place, the new spacecraft is ready to perform nearly all relay coordination functions within MaROS.

6.4 SPS and MPX

6.4.1 SPS Description

The Service Preparation Subsystem (SPS) [18] is a service operated by NASA's Deep Space Network (DSN). It acts as a repository for ephemeris predict files, among many other responsibilities specific to the DSN. Data derived from these files, which are typically delivered to SPS in the form of a Spacecraft and Planet Kernel (SPK) file [19] or an Orbit Ephemeris Message (OEM) file [20], are utilized throughout the MaROS ecosystem to determine the geometric positioning of each spacecraft in reference to one another. In addition, these ephemerides are required for interfacing with the DSN for direct-with-Earth (DWE) communications, as well as for participating in the Multimission Automated Deepspace Conjunction Assessment Process (MADCAP) collision detection system [21], as applicable.

Frequent delivery of updated ephemeris predict files is highly encouraged of all participants in the MRN to ensure that the latest geometry is used for relay coordination and operations. This includes orbiter ephemeris data as well as surface position data, especially in the case of a Mars rover or other mobile platform.

6.4.2 MPX Description

The Metrics Predict Executive (MPX) is a service affiliated with SPS that functions in part as a sub-service to MaROS. It automatically ingests SPKs that are newly published to SPS and subsequently calculates all the unique view periods between each pair of relay service user and relay service provider spacecraft. Upon completion of these calculations, these view periods are then used to update the MaROS database with the latest geometry information to ensure that relay services are requested while the spacecraft are actually in view of one another. As a standalone service, MPX can also calculate detailed view period information with user-selected time steps for use in other analyses.

As it is operationally important to always use the latest published ephemeris, participants in the MRN may use the "Location or Attitude File" (LOAF) Visualizer, accessed via the MaROS GUI, to verify which ephemeris file is being used for a specific period. Operationally, new ephemeris files are "overlaid" on old ones in MPX to facilitate the continuity of available data. An example screenshot of LOAF is shown in Figure 6.1.

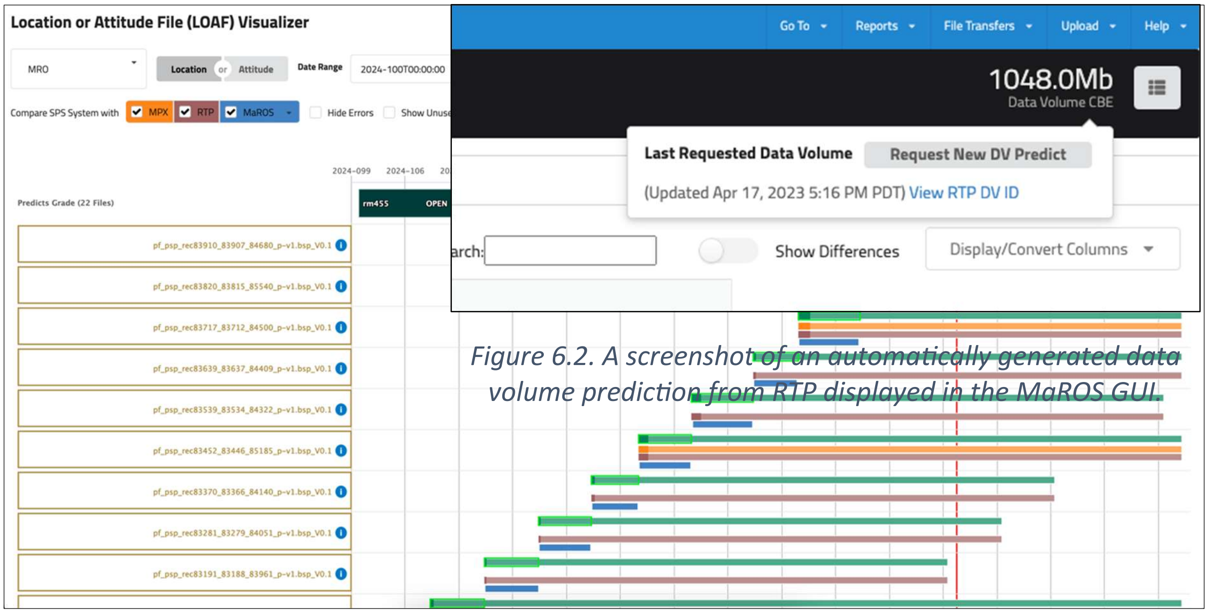

Figure 6.2. A screenshot of an automatically generated data volume prediction from RTP displayed in the MaROS GUI.

Figure 6.1. A screenshot of the Location or Attitude File (LOAF) Visualizer in the MaROS GUI, which shows the published orbiter SPKs used by MPX to calculate View Periods for MaROS and RTP.

6.4.3 Configuring SPS and MPX

For MaROS to automatically update its view period information when new ephemeris information is made available, new MRN participants must first configure the new spacecraft in SPS and MPX, where ephemeris information must be published. Once a valid repository is configured in SPS for the new spacecraft, the MaROS operations team will work with the MPX operations team to add the spacecraft to its list of spacecraft for which view periods are calculated. Once this is done, MaROS will automatically ingest and update view periods with every new ephemeris delivery and generate new return-link data volume predictions via RTP, if appropriate.

The use of SPS and MPX as sub-services of MaROS is required.

6.5 RTP

6.5.1 RTP Description

The Relay Telecom Predictor (RTP) is a service that supports the computation of predicted return-link data volumes that might be returned during any given relay session. The RTP framework runs models provided by the operators of the relay service user spacecraft. RTP acts as an “on demand” service, and the generation of a new data volume prediction can be triggered at any time, though usually after some state or condition has changed.

Typically, updated data volume predictions are automatically initiated by MaROS upon the delivery of new ephemeris data as well as when relay session request information is updated. The results are then made visible within the MaROS GUI. It may also be manually initiated during planning efforts as part of "what if" scenarios by users (see Figure 6.2). This "on demand" approach helps to ensure that the "best" data volume predictions are automatically tracked and available.

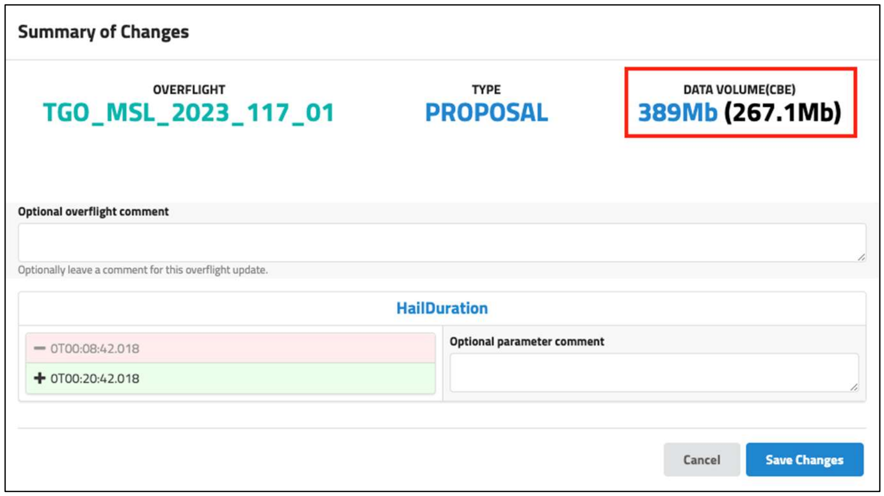

Figure 6.2. A screenshot of a manually initiated data volume prediction from RTP that was calculated before proposed relay session configuration changes were made to a relay service request in the MaROS GUI.

RTP can run data volume models generically, with models supplied by system users in the form of a containerized image that can be executed in a secure cloud environment. This architecture enables RTP to support any variety of model complexity. New relay service users are encouraged to assess whether existing data volume models can be used for their purposes or if they would like to provide their own. Additionally, relay service providers should coordinate with relay service users to confirm that spacecraft-specific properties reflected in the models are accurate (i.e., antenna patterns, attitude vector definitions, etc.).

6.5.2 Configuring RTP

To utilize the data volume prediction services available with RTP, relay service users must provide a model in the form of a containerized program that is compatible with the RTP system. At their core, these models need only accept input in a standardized JavaScript Object Notation (JSON) format provided by RTP, and likewise output data volume information in a JSON format expected by RTP.

RTP identifies to the model the set of data files (ephemeris files, spacecraft attitude files, etc.) that it determines to be most relevant to the indicated pair of spacecraft and time window. Alternatively, users can specify exactly which geometry files to utilize. RTP connects to SPS to provide SPK files to the models and is also capable of pulling files from a MaROS-based repository.

Optionally, users can provide a pre-calculation model for more computation-heavy calculations (such as Link Power and Signal-to-Noise Ratio (LPSNR)) that can be fed into data volume models for much faster computation on a per-overflight basis.

The use of RTP as a sub-service of MaROS is optional, as data volume estimates can be supplied manually via relay session requests.

6.6 WMRD

6.6.1 WMRD Description

"Where's My Relay Data" (WMRD) is a dashboard (see Figure 6.3) within the MaROS GUI that provides real-time updates during the transfer of return-link data at various checkpoints in the return-link data flow. This is presented as a series of widgets in the MaROS GUI that provides telemetry and monitoring information for each overflight in MaROS that relates to the flow of return-link data through the various flight and ground data systems.

Participation by relay service users and providers in submitting WMRD monitoring data is optional, but it is highly encouraged to facilitate transparency and to assist MRN users in identifying the return status of their data.

6.6.2 Configuring WMRD

With a widget-based design paradigm for the WMRD Dashboard, new widgets can be added to support new spacecraft should the relay dataflow for that spacecraft differ from existing relay dataflows. While some data in WMRD can be populated from generic data collection mechanisms maintained by the MaROS operations team, WMRD supports self-reporting of real-time monitoring data via a MaROS representational state transfer (ReST) application programming interface (API) endpoint. Self-reported data can be delivered in a wide variety of formats, with a data transformation layer within WMRD that parses the data into the format needed by the WMRD Dashboard.

The MaROS operations team stands ready to work with new MRN participants to develop new widgets and the underlying pipelines needed to display data in the WMRD dashboard.

The use of WMRD as a sub-service of MaROS is optional.

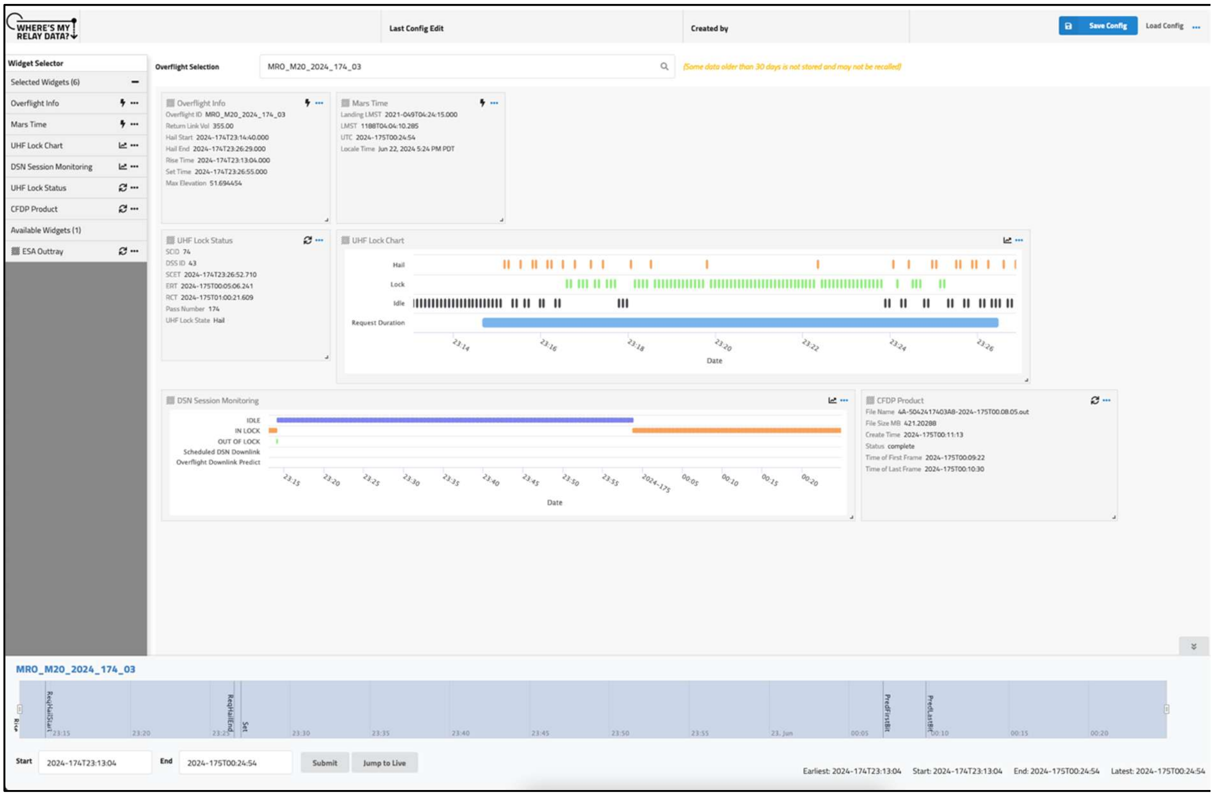

Figure 6.3. A screenshot of the "Where's My Relay Data" Dashboard within the MaROS GUI, displaying monitoring status of an overflight.

6.7 Eyes on the MRN

6.7.1 Eyes on the MRN Description

The "Eyes on the MRN" (referred to as the *MRN Visualizer*) is an interactive, 3-dimensional visualization of overflight information that is made available from MaROS and MPX. The tool allows users to view overflight information from a "birds-eye" view of Mars (see Figure 6.4), or from the perspective of the spacecraft involved (see Figure 6.5). This visualizer is available publicly at https://eyes.nasa.gov/mrn.

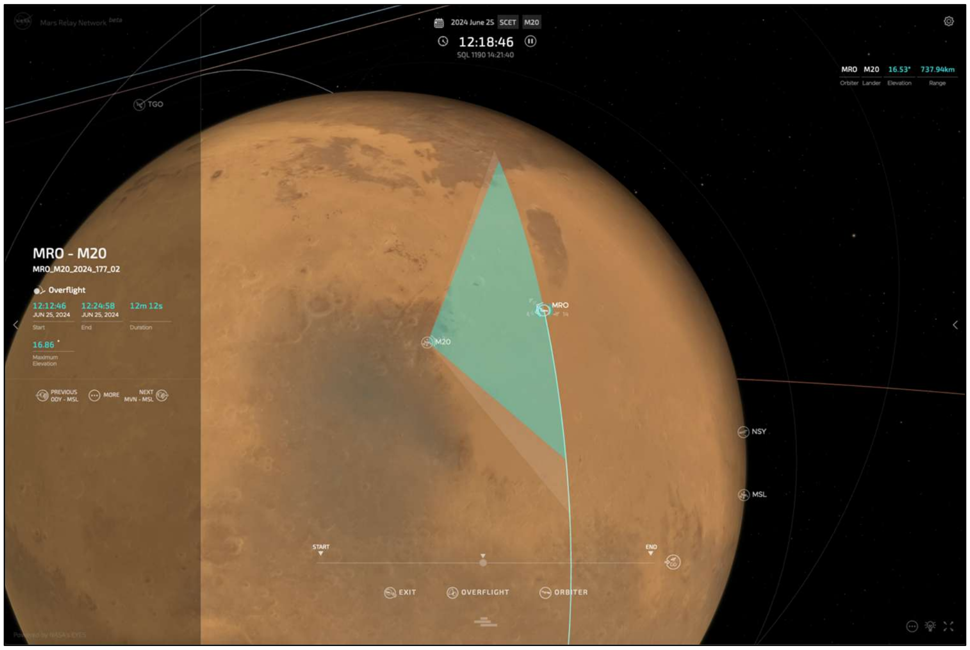

Figure 6.4. A screenshot of the "Eyes on the MRN" 3D visualizer showing an overflight between the Perseverance Rover (M20) and Mars Reconnaissance Orbiter (MRO).

Figure 6.5: A screenshot of the "Eyes on the MRN" 3D visualizer showing an overflight between the Perseverance Rover (M20) and the Mars Atmosphere and Volatile Evolution (MAVEN) orbiter (MVN) from the perspective of M20.

6.7.2 Configuring the "Eyes on the MRN" Visualizer

Similar to MPX, for a new spacecraft to be added to the MRN Visualizer, an internal configuration will need to be updated by the Eyes operations team to include the new spacecraft. Additionally, a 3D model of the spacecraft can be added for an accurate (though static) depiction of the spacecraft to appear in the tool. Predicted and as-flown spacecraft attitudes and trajectories are typically extracted from SPS and/or the Navigation and Ancillary Information Facility (NAIF) repository [22]. For landed missions, a high-resolution terrain map of the surrounding area allows the tool to display local terrain in relation to the lander's view of the orbiter during an overflight.

The MaROS operations team stands ready to work with new MRN participants to facilitate the inclusion of new spacecraft in the visualizer.

The inclusion of new MRN participants in the "Eyes on the MRN" Visualizer is considered optional but is preferred.

7 MaROS Interface File Formats

All MRN participants utilize XML-based file formats to interface with MaROS for both automated and manual interaction within the relay planning processes (see Section 8). This section includes a summary of these file formats and how they are used. Detailed specifications for these file formats will be provided to new participants after a formal agreement to join the MRN has been established.

7.1 Orbiter Request File (ORF)

The Orbiter Request File (ORF) is used by relay service users to request relay sessions from relay service providers. The ORF consists of one or more relay service requests, indicating both the time of the request and any parameters that are required, as defined by the relay service providers. It may also contain additional parameters specific only to the relay service user, such as those needed for tracking and implementation purposes. See Section 3.2 for a list of "negotiated" parameters that may be included in an ORF. An individual relay service request may also be submitted directly via the MaROS GUI. More typically in operations, however, many requests are submitted as a batch via an ORF upload via the MaROS command-line interface (CLI).

7.2 Overflight Acknowledgement File (OAF)

The complement to the ORF, the Overflight Acknowledgement File (OAF) is used by relay service providers to acknowledge requests received from users. During the Strategic and Tactical Processes (see Section 8.1) where the negotiation for relay services occurs, it is expected that the relay service provider will return the same set of parameters indicated in received relay service requests. Disagreements are flagged by MaROS as conflicts, suggesting the need to resolve the conflicts through negotiation until an agreeable plan is achieved. The OAF may also contain additional parameters specific to the relay service provider, as well as provider-specified forward- and return-link latency predictions for the given relay sessions. An individual relay service request may also be acknowledged directly via the MaROS GUI. More typically in operations, however, many requests are acknowledged as a batch via an OAF upload via the MaROS CLI.

7.3 Overflight Scorecard

The Overflight Scorecard (or Scorecard) is published by both the relay service provider and the relay service user after a given relay session has occurred as part of the Relay Accountability Process (see Section 8.4). It reports key metrics indicating the relay session's technical performance (see Section 3.6). As with the ORF and OAF, the parameters included in the Scorecard are typically defined by the relay service provider and echo those parameters used in the ORF and OAF to negotiate for relay services; this approach allows for direct accountability of how the relay session actually occurred compared to what was planned to occur. Under normal operational circumstances, there should be no deviations between what was planned to occur and what actually occurred. Key data included in the Scorecard includes the total data volume transferred during the relay session in both the forward- and return-link direction, periods when the relay session link is considered "closed", and the achieved data rate for the relay session, as applicable.

Typically, in operations, a Scorecard file is published via the MaROS CLI, one file per relay session, though the file format supports submitting data for multiple relay sessions simultaneously. In addition, the MaROS GUI supports the upload of these files.

7.4 Overflight Performance Assessment File (OPAF)

The Overflight Performance Assessment File (OPAF) is also published by both participants after a relay session has occurred as part of the Relay Accountability Process (see Section 8.4), typically in tandem with an Overflight Scorecard. It reports time-dependent lists of relay performance data to convey the spacecraft's telemetry and other ancillary data (see Section 3.6). As with the Scorecard, this data is used for reporting the technical performance of the relay session and typically includes a time-ordered listing of session metrics, such as the instantaneous data rate achieved during the session, the received signal level, and the measured signal-to-noise ratio. The contents of the OPAF can be tailored based on the needs of the relevant operators and are negotiated between the relay partners to ensure the most meaningful data are reported by each participant.

All MRN participants can easily access customizable reports in graphical format for each overflight on the MaROS tool where certain data from ORFs, OAFs, and Scorecards, and all data from OPAFs may be plotted. See Figure 7.1 for an example of how this data can be represented in graphical format via the MaROS GUI.

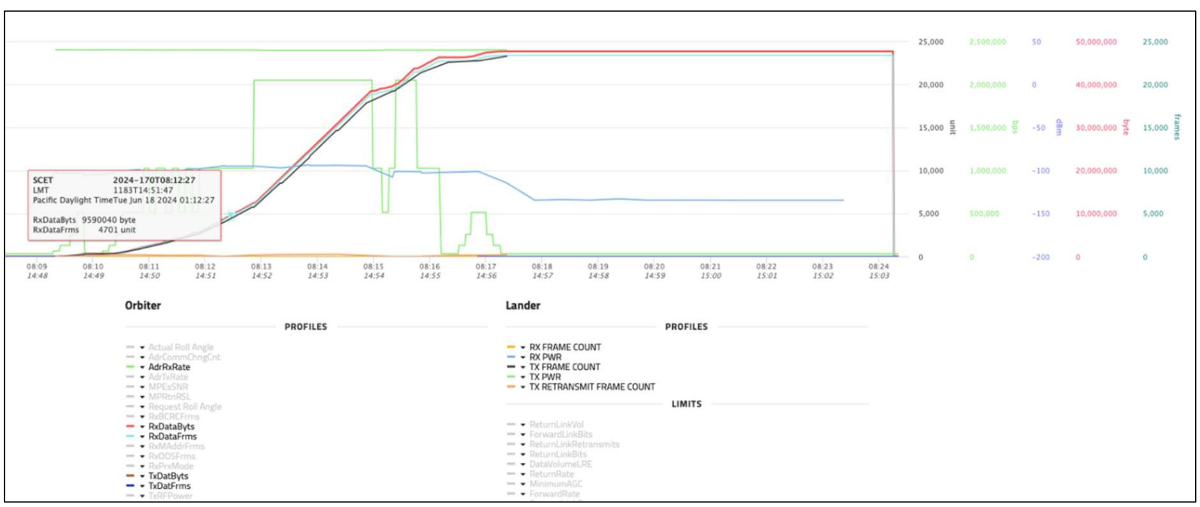

Figure 7.1. An example of a relay report shown in graphical format in the MaROS GUI.

Typically, in operations, an OPAF is published via the MaROS CLI, one file per relay session, though the file format supports submitting data for multiple relay sessions simultaneously. In addition, the MaROS GUI supports the upload of these files.

7.5 Orbiter Sequence of Events (OSOE) File

The Orbiter Sequence of Events (OSOE) file contains a time-dependent list of relay service provider events, as derived from the intended spacecraft commanding during the relay planning periods. Typically, it includes information regarding uplink and downlink windows from deep space tracking antennas, DTE data rates, and periods of time during which relay support is not available to relay service users. The contents of the OSOE are primarily used by MaROS to help calculate forward- and return-link latencies.

Typically, in operations, an OSOE file spans a designated period of time and is published via the MaROS CLI, though the upload of these files is also supported via the MaROS GUI.

7.6 Generic Sequence of Events (GSOE) File

The Generic Sequence of Events (GSOE) File is similar to the OSOE. It also contains a time-dependent list of events that may communicate any type of event across all participants of the network. It includes events defined by the relay service provider or the relay service user, such as periods of time during which the relay service provider spacecraft cannot perform spacecraft pointing in support of relay activities, designators for what spacecraft attitude the relay service provider spacecraft is scheduled to be in as a function of time, or indicators of when the lander teams may need certain types of relay support. By design, the GSOE is intended to be a means to communicate unexpected or spacecraft-unique data to help facilitate relay coordination activities.

Typically, in operations, a GSOE file spans a designated period of time and is published via the MaROS CLI, though the upload of these files is also supported via the MaROS GUI.

7.7 Contact Schedule File (CSF)

The Contact Schedule File (CSF) contains a time-dependent list of when a person or function is available from the relay service provider operations team to support direct-from-Earth (DFE) uplink activities in support of the Forward-Link Process (see Section 8.2). This data is primarily used to help calculate forward-link latencies but may also communicate when people are available for in-person support during tactical operations (as in the Tactical Planning Process, see Section 8.1).

Typically, in operations, a CSF file spans a designated period of time and is published via the MaROS CLI, though the upload of these files is also supported via the MaROS GUI.

7.8 Forward-Link Trigger (FLT)

The Forward-Link Trigger (FLT) file defines a list of a relay service user's forward-link data files, and issues instructions to MaROS and the relay service providers regarding how those data files should be handled in the Forward-Link Process (see Section 8.2). From a process viewpoint, the relay service users first submit their forward-link data as files to a MaROS-hosted file repository and then submit an FLT, which may be submitted as a file to MaROS (or constructed directly in the MaROS GUI) to initiate the process of transferring the forward-link files to the relay service provider operators. The FLT may also contain additional data or keywords to help direct the transfer of the files.

Typically, in operations, FLT data is constructed and published directly via the MaROS GUI, though MaROS supports the upload of these files via the MaROS CLI.

7.9 File Transfer Trigger and Messaging File (FTF)

The File Transfer Trigger and Messaging File (FTF) supports file exchanges between all the operators of spacecraft participating in the MRN by providing a means to submit meta-data related to those files. In the forward-link process, the FTF can functionally duplicate the FLT file, and it can also be used to transfer non-forward-link files or return-link data products. The FTF enables a more generalized means of exchanging data between the operators of the spacecraft in the MRN.

Typically, in operations, an FTF is constructed and published directly via the MaROS GUI, though MaROS supports the upload of these files via the MaROS CLI.

7.10 Overflight Summary File (OSF)

The Overflight Summary File (OSF) is the primary means of exporting relay session data from MaROS. It is an export-only file and lists overflights (each designated by a unique MaROS-generated overflight ID) with all or a selection of the parameters that outline geometric overflight information, relay session request and acknowledgement information, conflict information, latency information, etc. This file is received from MaROS by both the relay service users and the relay service providers during the Strategic and Tactical Planning Processes (see Section 8.1).

Typically, in operations, OSFs are downloaded from MaROS via the MaROS CLI, though MaROS supports the download of these files via the MaROS GUI.

7.11 System Configuration File (SCF)

The System Configuration File (SCF) contains a time-dependent list of parameters that enable users of MaROS to configure it to best reflect the nature and behavior of a spacecraft. Among other functions, the SCF is principally used by the relay service providers to specify the parameters that must be specified by relay service users when requesting relay services, as in Section 3.2.

Typically, in operations, SCF data is constructed and published directly via the MaROS GUI, though MaROS supports the upload of these files via the MaROS CLI.

> *Legacy Consideration: The Lander Orbit Propagation and Timing Geometry (LOPTG) File*
>
> In addition to the formats described above, the Lander Orbit Propagation and Timing Geometry (LOPTG) file format is a legacy file format that can be used to import relay view periods into MaROS in the absence of overflights determined by MPX. Before MPX came online as a MaROS-affiliated system in 2012, this file was the only means by which overflight information was published into the MaROS database. Each lander and rover team generated their own LOPTGs to identify relay opportunities, which necessitated an additional interface between the relay service providers and the relay service users to exchange orbiter ephemeris data. With the introduction of MPX, which centralizes and automates view period calculations, LOPTGs are no longer used operationally, though they may still be used for testing purposes.

8 Steps in the Relay Coordination Processes

This section describes the steps in each of the relay coordination processes as described in the MRN Participation Guide. Applicable file formats are indicated in the right-most column of the tables below (see also Section 7).

It should be noted that the MRN continues to evolve for a variety of reasons, which sometimes necessitates changes to these processes. For example, old spacecraft may be retired or lost, new spacecraft may bring new capabilities, and new ground system technologies may become available. When changes must occur, it is the goal of the IMRCWG to perturb existing implementations as little as possible, attempting to maintain heritage systems in-place and to reduce costs of adaptation. However, this is not always possible, and MRN participants need to remain flexible.

Throughout, MaROS is referenced as the system via which relay coordination activities occur. Despite the central role that MaROS plays, MaROS itself does not enforce the relay coordination processes as described below, instead remaining flexible to support a wide variety of possible processes.

8.1 The Strategic and Tactical Planning Processes

8.1.1 The Strategic Planning Process

Table 8.1 indicates typical steps in the Strategic Planning Process. All steps are relative to the Short-Range Relay Coordination (SRRC) Meeting, which occurs less than a week before the start of the two-week period being planned, which is often referred to as the *planning period.*[7] All data exchanges are via MaROS unless otherwise noted.[8] For every data publication, MaROS generates a notification that it sends to subscribed participants.

[7] The relay service provider operators define the boundaries of the planning periods for their spacecraft. Efforts are made network-wide to make the boundaries for each provider be as similar as possible to boundaries from other providers, though operational realities may make that impractical. Having similar planning period boundaries across the network enables the relay service user operators to effectively request relay services across the breadth of the MRN in a consolidated manner.

[8] ESA has implemented an internal service called the ESA Relay Coordination Office (ERCO), which acts as a portal for all ESA-MaROS interactions and manages data transformations between the MaROS-based file formats and those needed by ESA mission planning activities. ERCO is hosted at the European Space Operations Centre (ESOC) in Darmstadt, Germany. See Section 9.

Table 8.1: Steps in the Strategic Planning Process

Epoch	Event	Applicable File Formats
SRRC-28 days	The relay service provider operators generate and deliver initial data products, as needed: orbit ephemerides (delivered to SPS), lighttime information, deep space tracking times, anticipated spacecraft events that would hinder relay services (referred to as *non-relay periods*), realtime operator scheduled times, planning period boundaries, etc.	SPK or OEM, OSOE, GSOE, CSF
	Overflights are automatically identified by MPX. MaROS automatically calculates preliminary latencies and identifies possible planning conflicts between potential relay sessions due to geometry or orbiter-specified non-relay periods.	
	The relay service user operators extract this data from MaROS and use it to identify possible relay opportunities.	OSF
SRRC-21 days	The relay service user operators meet for a “pre-coordination” meeting, during which they negotiate to determine how to allocate the available overflights. This is an optional preliminary step intended to avoid conflicts later in the planning process.	
SRRC-14 days	The relay service user operators submit proposals for relay services to the relay service providers via MaROS. This is an optional preliminary step for those relay service providers that accept all relay service requests.	ORF
SRRC-7 days	The relay service provider operators extract these proposals from MaROS and evaluate their ability to support them. This is an optional preliminary step for those relay service providers that accept all relay service requests.	OSF
SRRC-7 days	The relay service provider operators submit their plans for providing relay services as proposed by the relay service user operators.[9] This is an optional preliminary step for those relay service providers that accept all relay service requests.	OAF

[9] Each relay service provider operator has their own criteria for determining if they can support relay services as proposed. MRO, for example, evaluates these proposals in the context of their science plans, whereas TGO accepts all proposals and performs their science planning around them.

Epoch	Event	Applicable File Formats
from SRRC-7 days to SRRC-1 day[10]	The relay service user operators extract this data from MaROS and evaluate it to confirm the willingness of the relay service provider operators to support the relay services as identified. This is an optional preliminary step for those relay service providers that accept all relay service requests.	OSF
SRRC-1 day (afternoon)	The relay service user operators submit their final selections for relay services as formal requests.	ORF
SRRC-1 day	The relay service provider operators extract these requests from MaROS and evaluate their ability to support them.	OSF
SRRC-1 day (afternoon)	The relay service provider operators acknowledge the requests as to be implemented. In rare cases, relay service providers may reject the requests as unsupportable, and iteration between the relay service user operators and the relay service provider operators may be necessary. The relay service providers may wait until after the SRRC meeting before submitting these acknowledgements.	OAF
SRRC-1 day (end of day)	A data lock in MaROS is applied to the planning periods to prohibit the submission of additional or updated relay service requests and comprehensive reports are generated of the relay service requests included in the designated planning period.	OSF
SRRC meeting	The SRRC meeting is held during which the MRN participants review the relay service request reports, make adjustments, if required (during or shortly after the SRRC meeting); and approve the relay plans for the designated planning periods as the final plan for relay services to be implemented for the upcoming planning period. The SRRC meeting is considered the end of the strategic planning process, as all negotiations are concluded.	
SRRC+0 or +1 day	The relay service provider and relay service user teams begin implementation of all command products needed to execute the agreed-upon plan for the planning period. .	

[10] This window of time can be used to iterate between the operators of the relay service users and the relay service providers to come to agreement on what relay services are to be supported. In practice, iteration is rare.

Epoch	Event	Applicable File Formats
SRRC+7 days	The planning period goes active and the included agreed upon relay services commence.	

8.1.2 The Tactical Planning Process

Table 8.2 indicates typical steps in the Tactical Planning Process, which is functionally very similar to the Strategic Planning Process. All steps are relative to the Last Nominal Uplink Time (LNUT), a forward-link latency calculated by MaROS or otherwise provided by the relay service providers. The LNUT represents the last opportunity for a tactical change request to be made with any assurance that the change can be accommodated by the relay service provider operators. All data exchanges are via MaROS unless otherwise noted. For every data publication, MaROS generates a notification that it sends to subscribed participants.

Table 8.2: Steps in the Tactical Planning Process for a Relay Service User[11]

Epoch	Event	Applicable File Formats
LNUT-hours or days	A relay service user operator identifies a need for a tactical change and considers the feasibility of implementing it given any known restrictions by the applicable relay service provider operators, especially the proximity of the current time to the LNUT.	
LNUT-hours or days	The relay service user operator submits a tactical proposal for updated relay services to the relay service provider operator via MaROS. This is typically accompanied by a separate email to appropriate spacecraft management.	ORF
LNUT-hours or days	The relay service provider operator extracts the proposal from MaROS and evaluates their ability to support it.	OSF
LNUT-hours or days	The relay service provider operators submit their plan for providing relay services as proposed.	OAF

[11] The process for a relay service provider to implement a tactical change is fundamentally the same as shown here.

Epoch	Event	Applicable File Formats
LNUT-hours or days	The relay service user operator extracts this data from MaROS and evaluates it to confirm the willingness of the relay service provider operators to support the tactical change. Iteration of the tactical proposal with the relay service provider operators may be necessary until an agreement is met or the tactical change is abandoned.	OSF
LNUT-hours or days	If proceeding, the relay service user operator submits their final request for the tactical change as a formal tactical request.	ORF
LNUT-hours or days	The relay service provider operator extracts the tactical request from MaROS and evaluates their ability to support it.	OSF
LNUT-hours or days	The relay service provider operator acknowledges the tactical request as to be implemented.	OAF
No later than LNUT	The relay service provider and relay service user operators begin implementation of all command products needed to execute the agreed-upon tactical change, as agreed, and ensure that the needed command products are transferred to the applicable spacecraft before the LNUT.	
Time of relay session	The updated plan for the relay session executes.	

During the negotiation processes described here, MaROS supports the specification of various types and categories of relay service requests and acknowledgements, which are used to support the various permutations of how relay services are scheduled across the network, as described next.

8.1.3 Relay Service Request Types

Relay service requests, as submitted to MaROS via an ORF, are always specified with a designated request type. These request types are strictly defined by MaROS and are used to aid in the workflow of relay service negotiations. They may be of the following types:

- Tentative: A tentative request, sometimes called a *strategic tentative request*, is one published to MaROS as a placeholder entry during the Strategic Planning Process, as used by the relay service user operators. The publication of a tentative request does not trigger any notifications from MaROS to the relay service provider operators and does not have a corresponding acknowledgement type, as per Section 8.1.5. This request type is most useful when trying to schedule relay services across multiple orbiters where various conflicting options exist or when allocating relay sessions across multiple users who share similar view periods with the relay service provider spacecraft.

- Proposal: A proposal, sometimes called a *strategic proposal*, is one published to MaROS during the Strategic Planning Process when the relay service user operators are interested in securing relay services during the designated overflight but are open to negotiation. The submission of a proposal triggers a notification from MaROS to the relevant relay service provider operators. It is typically answered by a "plan" from the relay service provider operators, as in Section 8.1.5.
- Request: A request, sometimes called a *strategic request* or a *formal request*, is one published to MaROS during the Strategic Planning Process when the relay service user operators have definitively identified which relay sessions are desired for use. The submission of a request triggers a notification from MaROS to the relevant relay service provider operators. It is typically answered by an "implementation" from the relay service provider operators, as in Section 8.1.5.
- Tactical tentative: A tactical tentative request is similar to the tentative request, as above, but is only used during the Tactical Planning Process. It does not trigger any notifications from MaROS to the relay service providers and does not have a corresponding acknowledgement type, as in Section 8.1.5. In practice, this request type is rarely used.
- Tactical proposal: A tactical proposal is one published to MaROS during the Tactical Planning Process when the relay service user operators believe a change is needed to a previously scheduled relay session (as identified during the Strategic Planning Process). A tactical proposal represents an inquiry to the relay service provider operators to ask if such a change can be accommodated. The submission of a tactical proposal triggers a notification from MaROS to the relevant relay service provider operators. It is typically answered by a "tactical plan" from the relay service provider operators, as in Section 8.1.5.
- Tactical request: A tactical request is one published to MaROS during the Tactical Planning Process when the relay service user operators have definitively identified that a tactical change is needed. The submission of a tactical request triggers a notification from MaROS to the relevant relay service provider operators. It is typically answered by a "tactical implementation" from the relay service provider operators, as in Section 8.1.5.

8.1.4 Relay Service Request Categories

Relay service requests, as represented by the request types in Section 8.1.3, are always specified with a designated request category. These request categories are strictly defined by MaROS and are used to aid in the workflow of relay service negotiations. They may be of the following types:

- Normal: This designator specifies a routine relay session, as requested during either the Strategic or Tactical Planning Processes.

- Contingency: This designator specifies that the relay service user operators would like to reserve the designated time for relay services but do not wish the relay session to execute unless otherwise indicated. In practice, this is rarely used but would typically be indicated only during the Strategic Planning Process. A relay session of this nature may be activated via the Tactical Planning Process as a "normal" relay session.

- Alternate: This designator indicates that this relay service request may be considered in the context of other relay service requests and may be optionally selected for implementation by the relay service provider operators during the Strategic Planning Process. In practice, this request category is rarely used.

- Emergency: This designator indicates that this relay service request is greatly needed in the context of a spacecraft emergency. It is typically only used when requesting relay sessions that were not previously scheduled during the Strategic Planning Process. Adjustments to relay sessions that were already scheduled during the Strategic Planning Process would not need the "emergency" designator, even if a spacecraft emergency exists, but instead should use the "normal" designator, as above. In practice, this request category is rarely used and would only be used during the Tactical Planning Process.

- Withdrawn: This designator indicates that this relay service request is no longer needed by the relay service user. The relay service provider may continue with the implementation and eventual execution of the relay session, but the relay service user spacecraft would not be expected to respond. It may be used during either the Strategic or Tactical Planning Processes.

- Disable: This designator indicates that this relay session should not occur. In contrast to "withdrawn", where the relay service provider may or may not execute the relay session, this designator specifically asks that the relay session not be executed. In practice, this is rarely used but would typically be used only during the Tactical Planning Process.

8.1.5 Relay Service Acknowledgement Types

Relay service acknowledgements, as submitted to MaROS via an OAF, are always specified with a designated acknowledgement type. These acknowledgement types are strictly defined by MaROS and are used to aid in the workflow of relay service negotiations. They may be of the following types:

- Plan: A plan, sometimes called a *strategic plan*, is one published to MaROS during the Strategic Planning Process in response to a strategic proposal published by the relay service user operators. Used as part of the negotiation process, ideally the details of a published plan should match the details of the corresponding proposal precisely, with deviations (which would be identified by MaROS) suggesting the need for further negotiation. The submission of a plan triggers a notification from MaROS to the relevant relay service user operators, which would highlight any deviations from the proposal to which the plan responds.

- Implemented: An implemented acknowledgement, sometimes called a *strategic implementation*, is one published to MaROS during the Strategic Planning Process in response to a strategic request published by a relay service user. Ideally, the details of a published implementation should match the details of the corresponding request precisely, with deviations (which would be identified by MaROS) suggesting the need for further negotiation. The submission of an implementation triggers a notification from MaROS to the relevant relay service user operators, which would highlight any deviations from the request to which the implementation responds. Typically, the submission of an implementation that includes no deviations from the request to which it responds represents the very end of the Strategic Planning Process.

- Tactical plan: A tactical plan is one published to MaROS during the Tactical Planning Process in response to a tactical proposal published by the relay service user operators. Used as part of the negotiation process, ideally the details of a published tactical plan should match the details of the corresponding tactical proposal precisely, with deviations (which would be identified by MaROS) suggesting the need for further negotiation. The submission of a tactical plan triggers a notification from MaROS to the relevant relay service user operators, which would highlight any deviations from the tactical proposal to which the tactical plan responds. It is notable that tactical proposals may suggest changes to the relay plan that cannot be accommodated by the relay service providers.
- Tactical implementation: A tactical implementation is one published to MaROS during the Tactical Planning Process in response to a tactical request published by the relay service user operators. It effectively defines the end state of any negotiation for a change to the relay plan as previously established during the Strategic Planning Process. The submission of a tactical implementation triggers a notification from MaROS to the relevant relay service user operators, which would highlight any deviations from the tactical request to which the tactical implementation responds. It is notable that tactical requests may suggest changes to the relay plan that cannot be accommodated by the relay service providers, or may indicate changes that are necessitated by the operational conditions of the relay service provider spacecraft, such as a spacecraft emergency.

8.1.6 Relay Service Support Plans

Relay service acknowledgements, as represented by the acknowledgement types in Section 8.1.5, are always specified with a designated support plan. These support plans are strictly defined by MaROS and are used to aid in the workflow of relay service negotiations. They may be of the following types:

- Supported: This designator indicates that a routine relay session, as requested during either the Strategic or Tactical Planning Processes, will be supported as specified.
- Contingency: This designator specifies that the relay service provider will reserve the designated time for relay services but will not execute the relay session unless otherwise requested by the relay service user via the Tactical Planning Process. In practice, this support plan is rarely used but would typically be indicated only during the Strategic Planning Process. If such a relay session was later activated, it would typically be acknowledged during the Tactical Planning Process with the supported designator, as above.

- Unsupported: This designator indicates that the relay service request submitted by the relay service user operators cannot be accommodated in its present form. During the Strategic Planning Process, it might be used in response to a proposal that has been withdrawn by the relay service user operators. During the Tactical Planning Process, it might be used to indicate that a proposal is unacceptable in its present form, but that a modification may be accommodated.
- Denied: This designator crisply indicates that a relay service request as submitted by the relay service user operators cannot be accommodated in any way. It may be used during either the Strategic or Tactical Planning Processes. If used during the Tactical Planning Process, it may imply that even though the tactical change can't be implemented, the strategic plan (with any previously implemented tactical changes) would still occur as previously planned; here, careful communication between the operators of the relay service providers and the relay service users is needed to avoid ambiguity.
- Cancelled: This designator crisply indicates that a previously scheduled relay session will not occur, despite being planned during either of the Strategic or Tactical Planning Processes, as applicable. Operationally, this is typically used when there is a problem with the relay service provider spacecraft that necessitates the suspension of relay services.

8.1.7 Relay Negotiation Examples

To illustrate how the workflow of the relay service negotiations may occur, the following dialogs are included, from Table 8.3 through Table 8.9. MaROS does not mandate any particular use of the requests and acknowledgements nor prescribes any process. These examples are provided only to indicate possible workflows and to illustrate the flexibility of the system.

In practice, the rovers on Mars today both take great advantage of the ADR algorithm available from the majority of the Mars orbiters, circumventing the need to perform nearly all tactical changes. Almost exclusively, tactical changes are in response to a spacecraft emergency on an orbiter, as reflected in Table 8.9.

Table 8.3: Nominal Strategic Process

Actor	Intent	Type/Plan	File Format
User	We think we want this relay session.	tentative, normal	ORF
User	Can we have this relay session?	proposal, normal	ORF
Provider	Sure, you can have that relay session.	plan, supported	OAF
User	Okay, let's use that relay session.	request, normal	ORF
Provider	Okay, we'll use that relay session.	implemented, normal	OAF

Table 8.4: Strategic Process with a Withdrawal

Actor	Intent	Type/Plan	File Format
User	Can we have this relay session?	proposal, normal	ORF
Provider	Sure, you can have that relay session.	plan, supported	OAF
User	Never mind. We don't need it.	proposal, withdrawn	ORF
Provider	Okay.	plan, unsupported	OAF

Table 8.5: Nominal Tactical Change

Actor	Intent	Type/Plan	File Format
User	We need this change.	tactical proposal, normal	ORF
Provider	Sure, we can make that change.	tactical plan, supported	OAF
User	Great, please make this change.	tactical request, normal	ORF
Provider	Okay, we'll make that change.	tactical implemented, supported	OAF

Table 8.6: Unsupportable Tactical Change

Actor	Intent	Type/Plan	File Format
User	We need this change.	tactical proposal, normal	ORF
Provider	No, we can't do that at all.	tactical plan, denied	OAF

Table 8.7: Potentially Supportable Tactical Change

Actor	Intent	Type/Plan	File Format
User	We need this change.	tactical proposal, normal	ORF
Provider	No, we can't do that.	tactical plan, unsupported	OAF
User	Well, how about this different change?	tactical proposal, normal	ORF
Provider	Sure, we can do that.	tactical plan, supported	OAF
User	Great, please make this different change.	tactical request, normal	ORF
Provider	Okay, we'll make that different change.	tactical implemented, supported	OAF

Table 8.8: Relay Service User Emergency

Actor	Intent	Type/Plan	File Format
User	Something really bad happened and we need this new relay session.	tactical proposal, emergency	ORF
Provider	Okay, but we can only do it this way.	tactical plan, supported	OAF
User	That works, let's do it that way	tactical proposal, normal	ORF
Provider	Okay.	tactical plan, supported	OAF
User	Great, let's make it happen.	tactical request, normal	ORF
Provider	Okay!	tactical implementation, supported	OAF

Table 8.9: Relay Service Provider Emergency

Actor	Intent	Type/Plan	File Format
Provider	Something really bad happened. Sorry, but you can't have your relay session.	tactical implemented, cancelled	OAF

8.2 The Forward-Link Process

Table 8.10 indicates typical steps in the Forward-Link Process. All steps are relative to the LNUT for a given relay session, which represents the last opportunity for a forward-link file to be passed through MaROS with any assurance that the file will be received by the relay service user spacecraft as intended. All data exchanges are via MaROS unless otherwise noted. For every data publication, MaROS generates a notification that it sends to subscribed participants.

Table 8.10: Steps in the Forward-Link Process

Epoch	Event	Applicable File Formats
LNUT-hours or days	A relay service user operator identifies a need to transmit forward-link data and identifies a relay session which would be optimal for the transfer. The relay service user operator also considers the likelihood of a successful transfer given the time available to perform the transfer and any known restrictions identified by the applicable relay service provider operators, especially the proximity of the current time to the LNUT.	
LNUT-hours or days	The relay service user operator generates the forward-link data, independently performing all necessary internal approval processes, and stages it in MaROS as one or more distinct file(s).	
LNUT-hours or days	When ready, the relay service user operator initiates the transfer of the forward-link data to the relay service provider operators via MaROS. When doing so, the operator may also provide additional meta-data associated with the forward-link data, such as the relay session during which the data should be transferred, if desired, and any instructions for how the data should be handled.	FLT or FTF

Epoch	Event	Applicable File Formats
LNUT-hours or days	The relay service provider operators receive a notification from MaROS of the forward-link data transfer and subsequently extracts the forward-link data and any associated meta-data from MaROS.	FLT or FTF
LNUT-hours or days	The relay service provider operators package the data, as necessary, and transfers the data to the relay service provider spacecraft via a deep space tracking network, where it is de-packaged onboard the relay service provider spacecraft, as necessary, and otherwise prepared for transfer during a relay session.	
Time of relay session	The relay service provider spacecraft transmits the forward-link data to the relay service user spacecraft.	

8.3 The Return-Link Process

Table 8.11 indicates typical steps in the Return-Link Process, assuming the use of MaROS in the return-link data flow. In practice, the primary objective is to implement the fastest data return possible for all relay data.

Table 8.11: Steps in the Return-Link Process

Epoch	Event	Applicable File Formats
Time of relay session	The relay service provider spacecraft receives return-link data from the relay service user spacecraft, packages it as necessary, and begins transmitting it to a deep space tracking network when conditions merit doing so.	
First bit return-link time	The deep space tracking network begins receiving the return-link data from the relay service provider spacecraft.	
Last-bit return-link time	The deep space tracking network receives the last portion of the return-link data from the relay service provider spacecraft and delivers the data to the relay service provider MOC.	

Epoch	Event	Applicable File Formats
Last-bit return-link time	The relay service provider MOC performs any data transformations (such as de-packaging) to yield the data as-received during the relay session and delivers the return-link data to MaROS.	
Last-bit return-link time	When ready, the relay service provider operator initiates the transfer of the return-link data to the relay service user operators via MaROS.	FTF
Last-bit return-link time	The relay service user operators receive notifications from MaROS of the return-link data transfer from MaROS and subsequently extract the return-link data and any associated meta-data from MaROS.	FTF

8.4 The Relay Accountability Process

Table 8.12 indicates typical steps in the Relay Accountability Process. All data exchanges are via MaROS. For every data publication, MaROS generates a notification that it sends to subscribed participants. In principle, it is desired to publish this accountability data as soon after a relay session occurs as possible. MaROS supports the successive publication of data to facilitate scenarios where a partial set of accountability data comes available before the full data set does.

Table 8.12: Steps in the Relay Accountability Process

Epoch	Event	Applicable File Formats
Time of relay session	A relay session executes, with data exchanged between the relay service user spacecraft and the relay service provider spacecraft, in the return-link direction, the forward-link direction, or both simultaneously. Both spacecraft in the relay session collect and return the engineering data related to the relay session to the appropriate MOC as soon as practicable using any means available.	
First bit return-link time or later	The relay service user operators and the relay service provider operators process the engineering data in their MOCs and publish them to MaROS at the earliest possibility.	Scorecard, OPAF

Epoch	Event	Applicable File Formats
First bit return-link time or later	The relay service user operators and the relay service provider operators receive notifications from MaROS of the data publication.	

9 The ESA Relay Coordination Office (ERCO)

This section describes the functionality of the ESA Relay Coordination Office (ERCO), and outlines how it interacts with MaROS to facilitate cross-agency relay activities.

This section will be contributed later.

References

[1] CCSDS (July 2013). Proximity-1 Space Link Protocol – Physical Layer. CCSDS 211.1-B-4.

[2] CCSDS (October 2019). Proximity-1 Space Link Protocol – Coding and Synchronization Sublayer. CCSDS 211.2-B-3.

[3] CCSDS (July 2020). Proximity-1 Space Link Protocol – Data Link Layer. CCSDS 211.0-B-6.

[4] CCSDS (December 2021). Proximity-1 Space Link Protocol – Rationale, Architecture, and Scenarios. CCSDS 230.1-G-2.

[5] Taylor, F. J., e. a. (2014). DESCANSO Book Series Volume 13, Deep Space Communications, Chapter 6, pp. 229-237, Mars Reconnaissance Orbiter. Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, USA.

[6] A 48/56/64 kbit/s Data Circuit-Terminating Equipment Standardized for Use on Digital Point-to-Point Leased Circuits. ITU-T Recommendation V.38. Geneva: ITU, 1996.

[7] Performance Characteristics for Intermediate Data Rate Digital Carriers Using Convolutional Encoding/Viterbi Encoding. Revision 10. IESS 308. Washington, District of Columbia, USA. INTELSAT, 2000.

[8] Chamberlain, N. (March 2022). Implementing Low-Density Parity-Check Codes in the Mars Relay Network. IEEE Aerospace Conference. Big Sky, Montana, USA.

[9] Edwards, C. (2015). A first demonstration of Mars crosslink occultation measurements. Radio Science (50), 997-1007.

[10] Parrott, J. (2024). First Results of Mars Express—ExoMars Trace Gas Orbiter Mutual Radio Occultation. Radio Science, 59(7).

[11] CCSDS (July 2020). File Delivery Protocol (CFDP). CCSDS 727.0-B-5.

[12] CCSDS (October 2021). AOS Space Data Link Protocol. CCSDS 732.0-B-4.

[13] L3Harris. (n.d.). C/TT-505 Transceiver. Retrieved from https://www.l3harris.com/all-capabilities/spacecraft-communications-uhf-transceiver.

[14] L3Harris. (n.d.). Electra-Lite. Retrieved from https://www.l3harris.com/all-capabilities/mars-electra-lite-uhf-transceiver.

[15] Taylor, F.J. (October 2014). DESCANSO Book Series Volume 13, Deep Space Communications, Chapter 6, Mars Reconnaissance Orbiter. Retrieved from https://descanso.jpl.nasa.gov/monograph/series13/DeepCommo_Chapter6--141029.pdf.

[16] Friis, H. (May 1946). A Note on a Simple Transmission Formula. Proceedings of I.R.E. and Waves, 34(5), 254-256.

[17] Sauer, B. (May 2025). Mars Relay Operations Service (MaROS): Adapting to the Needs of the Mars Relay Network. SpaceOps 2025. Montreal, Quebec, Canada.

[18] Carruth, J. (May 2010). A Collaborative Scheduling Environment for NASA's Deep Space Network. SpaceOps 2010. Huntsville, Alabama, USA.

[19] Generic Kernels, The Navigation and Ancillary Information Facility, retrieved from https://naif.jpl.nasa.gov/naif/data_generic.html.

[20] CCSDS (April 2023). Orbit Data Messages. CCSDS 502.0-B-3.

[21] Berry, D. (May 2018). Automated Spacecraft Conjunction Assessment at Mars and the Moon – a Five Year Update. SpaceOps 2018. Marseille, France.

[22] SPICE Data (Spice Kernels), The Navigation and Ancillary Information Facility, retrieved from https://naif.jpl.nasa.gov/naif/data.html.

Acronyms

ADR – adaptive data rate
API – application programming interface
BER – bit error rate
BPSK – binary phase shift keying
C&DH – command and data handling
Caltech – The California Institute of Technology
CCSDS – Consultative Committee for Space Data Systems
CCITT – Consultative Committee for International Telephony and Telegraphy
CCR – communication change request
CLI – command-line interface
CRC – cyclic redundancy check
CSF – Contact Schedule File
CW – continuous wave
dB – decibel
DFE – direct-from-Earth
DSN – NASA's Deep Space Network
DTE – direct-to-Earth
DWE – direct-with-Earth
EDL – entry, descent, and landing
EEPROM – electrically erasable programmable read-only memory
EMI – electromagnetic interference
ERCO – ESA Relay Coordination Office
ESA – European Space Agency
ESOC – ESA's European Space Operation Centre
FDR – fixed data rate
FLB – forward-link binary (file)
FLT – Forward-Link Trigger
FSK – frequency-shift keying
FTF – File Transfer Trigger and Messaging File
GHz – gigahertz
GSOE – Generic Sequence of Events (file)
GUI – graphical user interface
HDO – half-duplex overlay
Hz – hertz
ID – identifier
IDD – Interface Definition Document
IESS – Intelsat Earth Station Standards
IMRCWG – International Mars Relay Coordination Working Group
IP – internet protocol
JPL – Caltech's Jet Propulsion Laboratory
JSON – JavaScript Object Notation
kbps – kilobits per second
kHz – kilohertz
km – kilometer
ksps – kilosymbols per second
LDAP – Lightweight Directory Access Protocol
LDPC – low-density parity check
LOAF – Location or Attitude File
LOPTG – Lander Orbit Propagation and Timing Geometry (file)
LPSNR – Link Power and Signal-to-Noise Ratio
MaROS – Mars Relay Operations Service
MaROS-Ops – the MaROS operations software environment
MaROS-TB – the MaROS testbed software environment
MAVEN – Mars Atmosphere and Volatile Evolution (orbiter)
Mbit – megabit ($1x10^6$ bits)
MER – Mars Exploration Rover
MHz – megahertz
MOC – mission operations center
MRN – Mars Relay Network
MRO – NASA's Mars Reconnaissance Orbiter
MPX – Metrics Predict Executive, an SPS service
NAIF – Navigation and Ancillary Information Facility
NASA – National Aeronautics and Space Administration (U.S.)
NRZ – non-return-to-zero
OAF – Overflight Acknowledgement File
OEM – Orbit Ephemeris Message
OPAF – Overflight Performance Assessment File
ORF – Orbiter Request File
OSF – Overflight Summary File
OSOE – Orbiter Sequence of Events (file)
PDU – protocol data unit
PL – path loss
PLL – phase lock loop
PSK – phase-shift keying
QoS – quality of service
QPSK – quadrature phase shift keying
R – slant range
ReST – representational state transfer
RF – radio frequency
RS – Reed-Solomon (error correction code)

RTP – Relay Telecom Predictor, a MaROS service
SANA – Space Assigned Numbers Authority
SCET – spacecraft ephemeris time
SCF – Spacecraft Configuration File
SCID – spacecraft identifier
SNR – signal-to-noise ratio
SPK – Spacecraft and Planet Kernel
SPS – Service Preparation Subsystem, a DSN-provided service
SRRC – Short-Range Relay Coordination (meeting)
SSNR – symbol signal-to-noise ratio
SSPA – solid state power amplifier
TGO – ESA's ExoMars Trace Gas Orbiter
TSP – technical support package
UAT – user acceptance test
UHF – ultra-high frequency
UTC – Universal Time Coordinated
WMRD – Where's My Relay Data, a MaROS service
XML – extensible markup language