17th International Conference on Space Operations, Dubai, United Arab Emirates, 6-10 March 2023.

Antenna Beam-Sharing: Progress Toward Multiple Uplinks Per Antenna

Douglas S. Abraham, Shakeh E. Brys, Jay L. Gao, Shantanu Malhotra, Leila Meshkat, David D. Morabito, James A. O'Dea, Emily R. Pascua, Marc Sanchez-Net, Dong K. Shin, and Zaid Towfic
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109

Within a decade, the number of spacecraft requiring support from NASA's Deep Space Network (DSN) is expected to increase by a factor of 2 to 3. Key drivers behind this expectation include NASA's increasing reliance on more affordable, targeted smallsat missions, a more robust science program in general, and the emergence of multi-element human lunar exploration missions. Given that a proportional increase in the number of antennas needed to communicate with all of these missions might prove prohibitively expensive, NASA has been exploring ways to gain increased capacity from its existing antennas. One way to do this is through antenna beam-sharing. While successfully employed for simultaneous downlink, beam-sharing for simultaneous uplink has not been an option, within the same frequency band, due to the severity of the intermodulation products. An alternative approach to achieving near-simultaneous uplink involves multiplexing the Forward Communications Link Transmission Units (FCLTUs) intended for all of the participating in-beam spacecraft onto a single uplink frequency, which they all receive. Each spacecraft then identifies which of the transfer frames within the received FCLTUs are intended for it on the basis of spacecraft ID. Known as Multiple Uplinks Per Antenna (MUPA), this technique has been under study for several years and is now at the point where key components are being prototyped and implementation pathfinding is occurring. This paper describes the progress in these efforts to date.

These efforts fall into three broad categories: the FCLTU multiplexer, ground system interfaces, and the spacecraft radio. The FCLTU multiplexer efforts include developing and testing the software needed to: receive FCLTUs from disparate Mission Operations Centers (MOCs), assess whether their destination spacecraft are in view or not, queue them accordingly, multiplex the in-view FCLTUs, send them on to the DSN under a virtual spacecraft ID for radiation, and feedback the appropriate radiated FCLTU statistics to each of the originating MOCs. The ground system interface efforts include identifying the modifications to DSN systems needed to successfully schedule MUPA sessions, distribute appropriate support data products to the uplink and tracking data systems, and ensure that the coherent turnaround of the shared uplink carrier frequency to each of the individual spacecraft downlink carrier frequencies is executed and communicated in ways that fully support simultaneous 2-way Doppler and ranging. The spacecraft radio efforts include designing, simulating, and/or testing radio features on the JPL Iris radio that will enable it to support variable turnaround ratios, selectable uplink frequencies, onboard Doppler compensation when trying to acquire and lock onto the shared uplink frequency, and filtering of the received transfer frames on the basis of spacecraft ID. These efforts have also involved examining MUPA's compatibility with different types of uplink encryption, coding, and modulation and how these factors would enter into the scheduling of a MUPA session.

Together, all of these efforts and the underlying system engineering for them, are "setting the stage" for the next key step on the path to implementing MUPA: flight demonstrations. Preparations with the SunRISE and EscaPADE projects for a multiplexer demonstration are already underway, and the team is working to identify other flight projects amenable to demonstrating the Iris radio's MUPA-related capabilities.

1. Introduction

NASA's Deep Space Network is currently supporting more spacecraft than it has at any other time in its 65-year history. And, it is doing so with fewer antennas than it has had in the past. As of the writing of this paper, the DSN is supporting 44 spacecraft with 14 antennas. A little over 15 years ago, it was supporting approximately 30 spacecraft with 16 antennas.[1] One of the DSN's keys to supporting more spacecraft with fewer antennas has been its development and application of the Multiple Spacecraft Per Antenna (MSPA) technique. As shown in the downlink portion of **Figure 1**, spacecraft that are within the half-power beamwidth of a single antenna can simultaneously downlink through it to separate receivers. Originally, only two spacecraft at a time could do so (2-MSPA). Now, four spacecraft at a time can do so (4-MSPA). And, by expanding the number "n" of available receivers to something greater than four in the future and also enhancing the network monitor & control and operator interfaces accordingly, "n" in-beam spacecraft could simultaneously downlink. While the occasions when multiple spacecraft are within the half-power beamwidth of a single DSN 34m diameter antenna are limited, they are numerous enough for MSPA to have a significant impact on network loading. At Mars, Venus, and more distant destinations, spacecraft are within the X-band (8.4 GHz) half-power beamwidth of a single antenna virtually all of the time. Spacecraft deploying as secondary payloads off of a common upper stage (such as the DSN-supported cubesats launching with Artemis-1) are in-beam together for many hours until their individual trajectory correction maneuvers take them out of beam. Constellations of spacecraft with inter-spacecraft separation distances designed to fall within the beamwidth of a single 34m antenna are also amenable to using MSPA. So, too, are lunar surface elements clustered together in a specific location, such as at the lunar south pole. Hence, there are enough opportunities to simultaneously service multiple spacecraft with a single antenna where MSPA can make a noticeable difference in network loading.

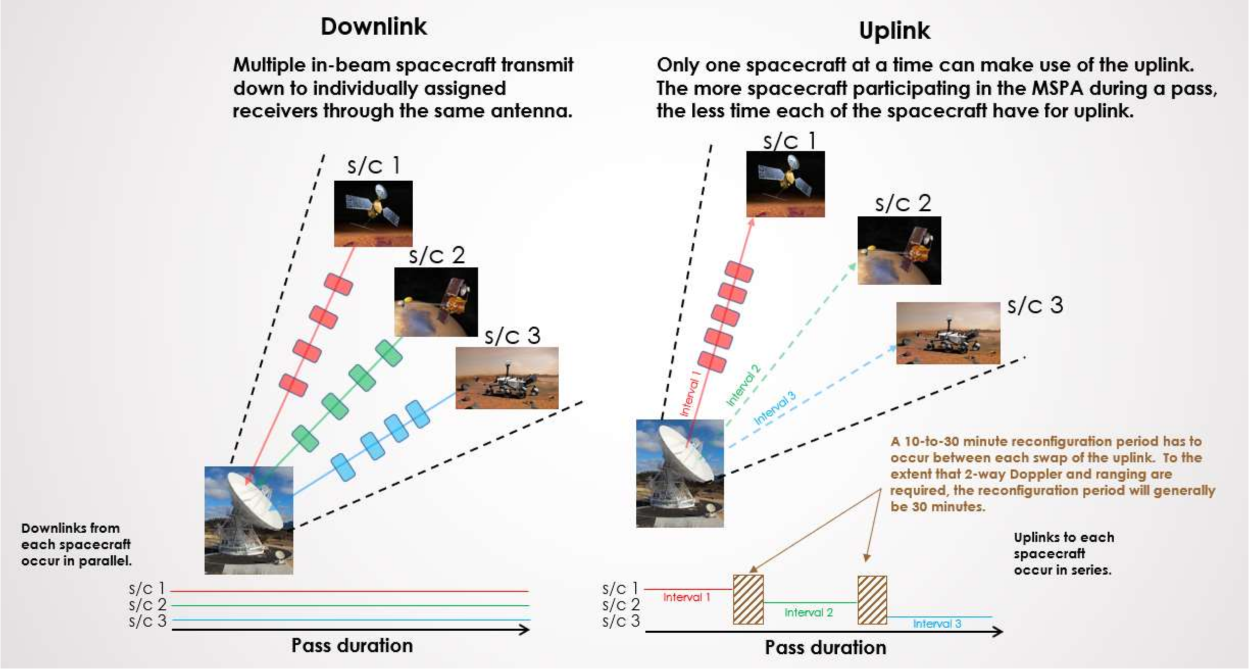

Figure 1. How the Multiple Spacecraft Per Antenna (MSPA) Technique Works

However, as the uplink portion of Figure 1 illustrates, MSPA has one key shortcoming that constrains its utility. Currently, during an MSPA pass, the uplink has to be shared between each spacecraft one at a time, with a 10- to 30-minute reconfiguration period between each handover. This significantly limits the amount of time each MSPA participant has available for commanding and, even more critically, for obtaining the two-way Doppler and ranging observables needed for navigation. If the DSN were to move to n-MSPA, this limitation would simply get worse because the time spent in reconfiguration would increase linearly with "n". As suggested in **Figure 2** by the number of DSN-supported spacecraft expected in the coming decade, finding a solution to this uplink problem is more urgent than ever.[2]

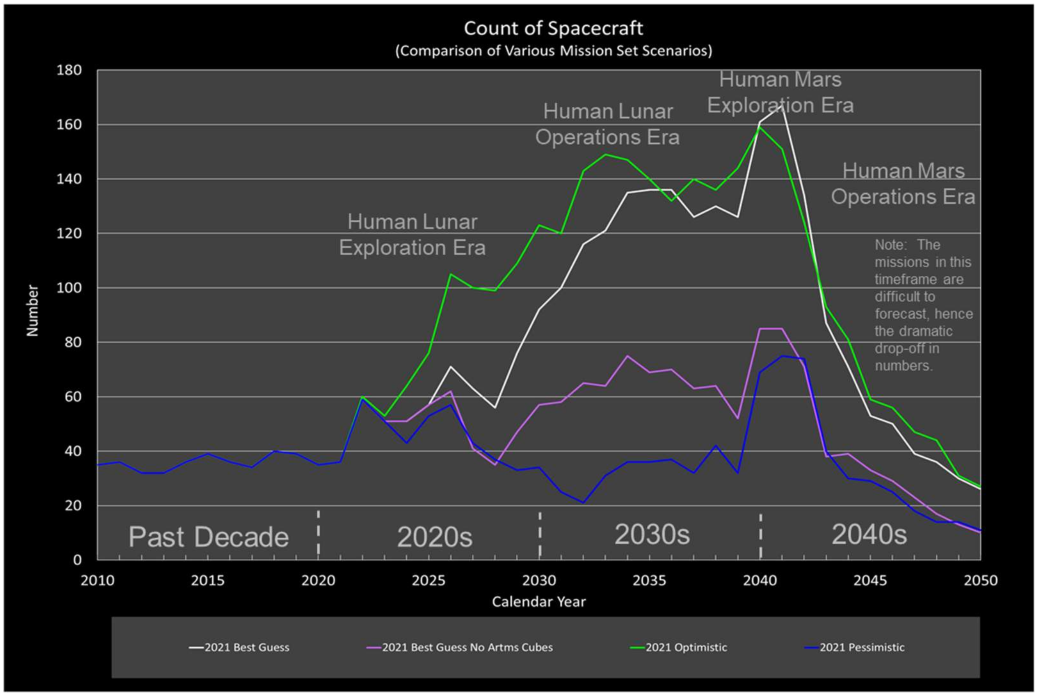

Figure 2. For Various Future Mission Set Scenarios, Projected Number of Spacecraft Requiring DSN Support as a Function of Time

To this end, NASA/JPL has been investigating ways to uplink to multiple in-beam spacecraft simultaneously.[3, 4] As illustrated in **Figure 3**, three techniques for achieving near-simultaneous uplink have been examined: (1) a "brute-force" approach involving the use of multiple exciters in conjunction with a single combiner and transmitter, (2) modulating subcarriers onto a single carrier and using each subcarrier as an assigned uplink channel, and (3) multiplexing the Forward Communications Link Transmission Units (FCLTUs) intended for each in-beam spacecraft onto a single uplink carrier with the radiated FCLTUs being received by all of the spacecraft and the recovered transfer frames differentiated on the basis of spacecraft ID. The multiple-exciter-single-transmitter approach tends to create intermodulation products that exceed the spectrum interference limits, requires an 80 kW transmitter to ensure that the simultaneous operation of the exciters occurs within a linear regime, and, given the operational complexity, would likely not be practical for more than a couple of simultaneous supports. In the subcarriers-modulated-on-a-carrier approach, a sine wave could be used to modulate subcarriers onto the carrier, with each subcarrier a fixed frequency distance from the other. However, this technique would likely require an 80 kW transmitter for power at interplanetary distances, has the potential for RF out-of-band "splattering," and makes Doppler shift accommodation more challenging, requiring changes to the way the DSN configures subcarrier tracking loops and, possibly, requiring advanced methods for Doppler compensation at the subcarrier level. For more than a couple of simultaneous links, this arrangement quickly becomes untenable. The third approach, multiplexed FCLTUs on a single uplink frequency, would likely be more straight-forward to implement on the ground, but requires variable turnaround ratios on the spacecraft radios for 2-way Doppler and ranging. It also requires a way to actively compensate for the Doppler differences each spacecraft's radio would experience when viewing the single-frequency uplink from an orbit different from that of the other participating spacecraft. An evaluation of the various technical, cost, and performance trades associated with these three approaches ultimately culminated in a decision to pursue the third one as the preferred MUPA approach. **Figure 4** illustrates how this approach would work in conjunction with MSPA. The remainder of this paper describes the ongoing prototyping and implementation pathfinding efforts associated with the approach. These efforts are organized in terms of three broad categories: the FCLTU Multiplexer, ground system interfaces, and spacecraft radio considerations and changes.

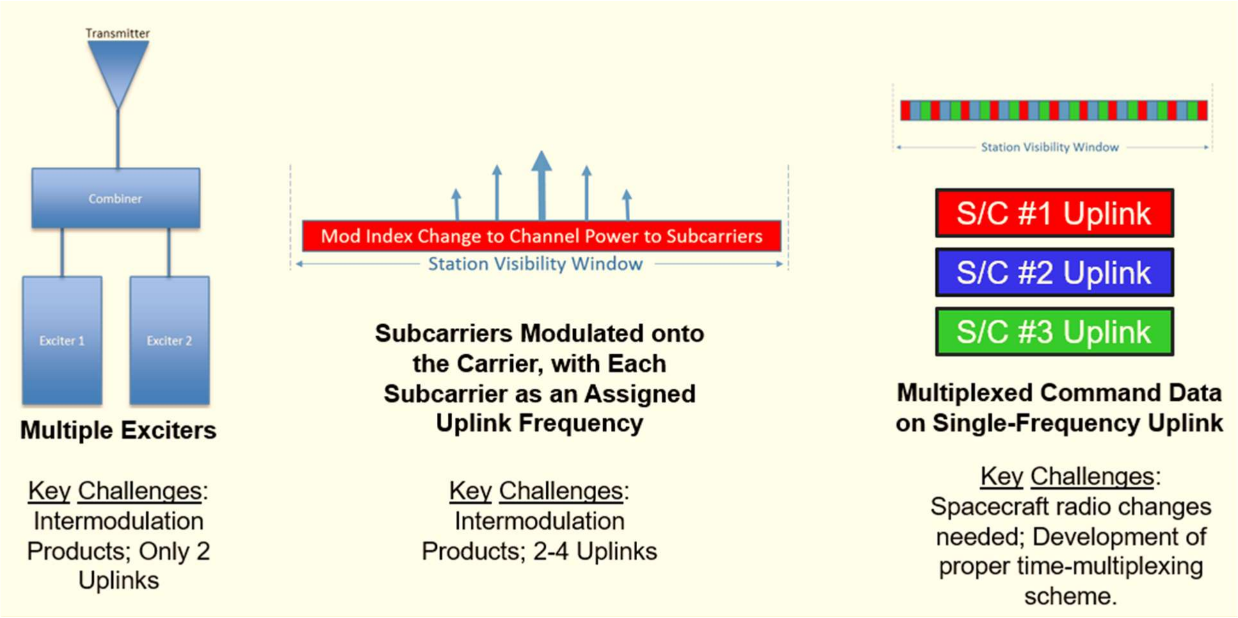

Multiple Exciters

Key Challenges: Intermodulation Products; Only 2 Uplinks

Subcarriers Modulated onto the Carrier, with Each Subcarrier as an Assigned Uplink Frequency

Key Challenges: Intermodulation Products; 2-4 Uplinks

Multiplexed Command Data on Single-Frequency Uplink

Key Challenges: Spacecraft radio changes needed; Development of proper time-multiplexing scheme.

Figure 3. Three Techniques for Near-Simultaneous, In-Beam Uplinks Explored During 2016

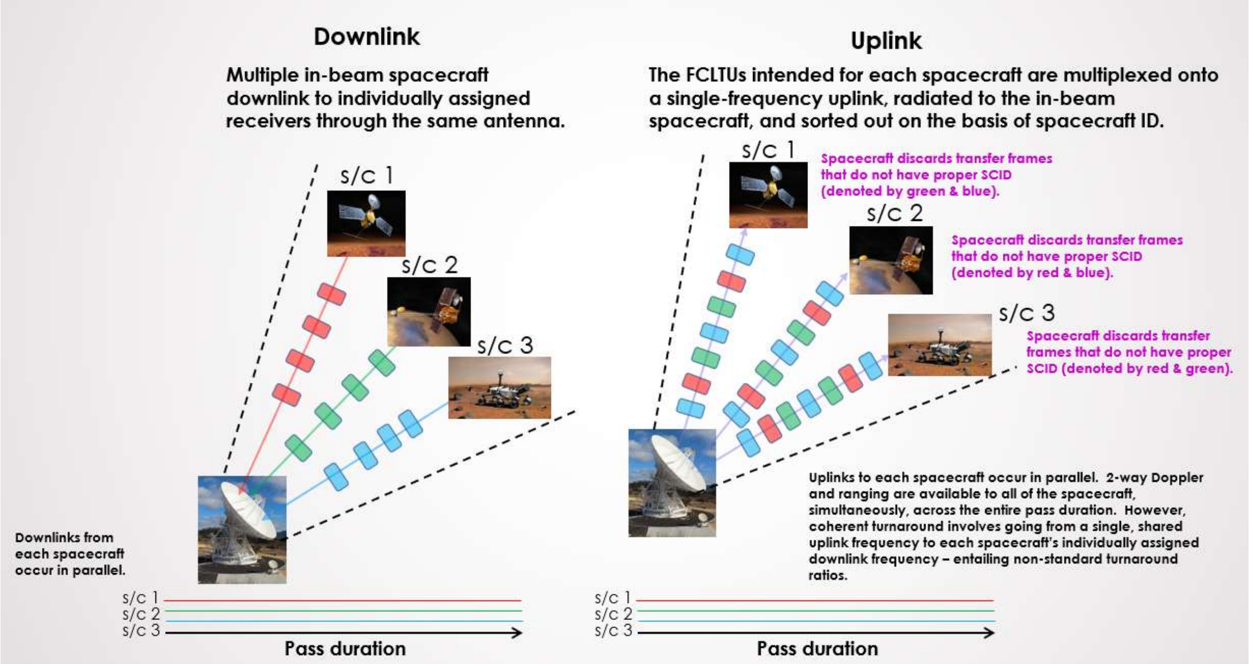

Downlink

Multiple in-beam spacecraft downlink to individually assigned receivers through the same antenna.

Downlinks from each spacecraft occur in parallel.

Uplink

The FCLTUs intended for each spacecraft are multiplexed onto a single-frequency uplink, radiated to the in-beam spacecraft, and sorted out on the basis of spacecraft ID.

Spacecraft discards transfer frames that do not have proper SCID (denoted by green & blue).

Spacecraft discards transfer frames that do not have proper SCID (denoted by red & blue).

Spacecraft discards transfer frames that do not have proper SCID (denoted by red & green).

Uplinks to each spacecraft occur in parallel. 2-way Doppler and ranging are available to all of the spacecraft, simultaneously, across the entire pass duration. However, coherent turnaround involves going from a single, shared uplink frequency to each spacecraft's individually assigned downlink frequency – entailing non-standard turnaround ratios.

Figure 4. Downlink MSPA and MUPA Being Used Together During the Same Pass

2. Design and Prototyping of the FCLTU Multiplexer

2.1 Establishing the Feasibility of the Multiplexing Approach

Before embarking on the design and prototyping of the FCLTU Multiplexer (FMUX), we first needed to establish two things: (1) that the multiplexed CLTUs intended for multiple, in-beam spacecraft could all be radiated within the timeframe of a single pass and (2) that different types of forward-error correction coding would be compatible with the multiplexing. Elaborating on the first feasibility concern, if participating mission users typically radiate CLTUs to their spacecraft over the entire duration of their passes, then fitting all of those CLTUs onto a single-frequency uplink during a shared MSPA pass of similar duration would prove problematic. The effective uplink rate for any given user would likely be very slow, and no one might actually finish uplinking all that they needed to send.

To investigate this issue, we analyzed a year's worth of CLTU radiation statistics for all of the DSN's supported missions to ascertain whether CLTU radiation times are generally a small enough portion of the total pass duration to allow CLTU multiplexing to occur during an MSPA pass. **Figure 5** shows that the average CLTU radiation time across a year's worth of passes is only 1.75% of the total time that the transmitter is active during the pass. And, in **Figure 6,** we see that 95% of all these passes entail CLTU radiation times less than 5 seconds in duration. Hence, multiplexing the CLTUs intended for several different in-beam spacecraft appears to be quite feasible. Even with 16 average spacecraft, the total multiplexed CLTU radiation time over an 8-hour pass would be on the order of only a couple of minutes. In fact, with the right multiplexing algorithm, all of the participating spacecraft might be commanded virtually simultaneously.*

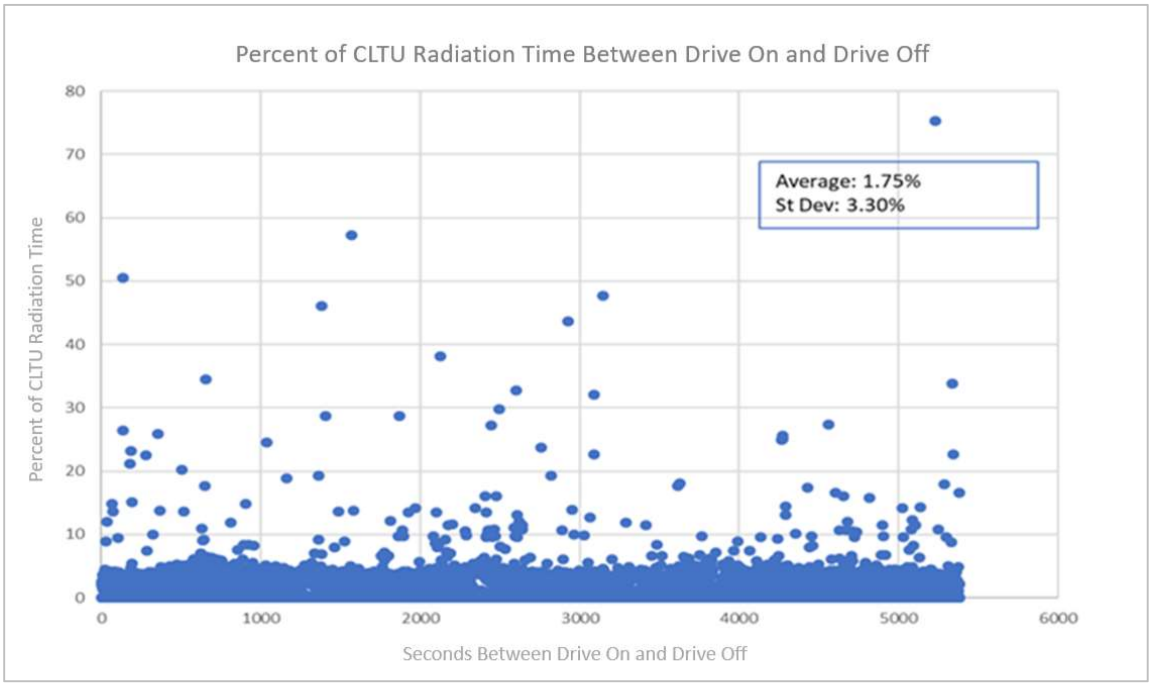

Figure 5. The Percent of CLTU Radiation Time While the Transmitter Is On, Per Pass, Across a Year's Worth of Uplink Passes

---

* In continuous communication cases (e.g., audio, video, large software uploads, etc.), CLTUs might be radiated over the entire pass duration. In such cases, if all of the mission participants have similar communication needs, a "round-robin" approach to multiplexing might prove desirable, since the perceived "simultaneity" of the communications would then be a function of the achievable data rate. In a mixed customer case (e.g., large software uploads vs. commanding), a more sophisticated priority-based or weighted multiplexing algorithm would likely be needed. In all cases, since all of the participants have their CLTUs multiplexed on a single-frequency uplink with which all of the participating spacecraft are in carrier-lock, 2-way Doppler and ranging would be available to all of them over the entire pass duration.

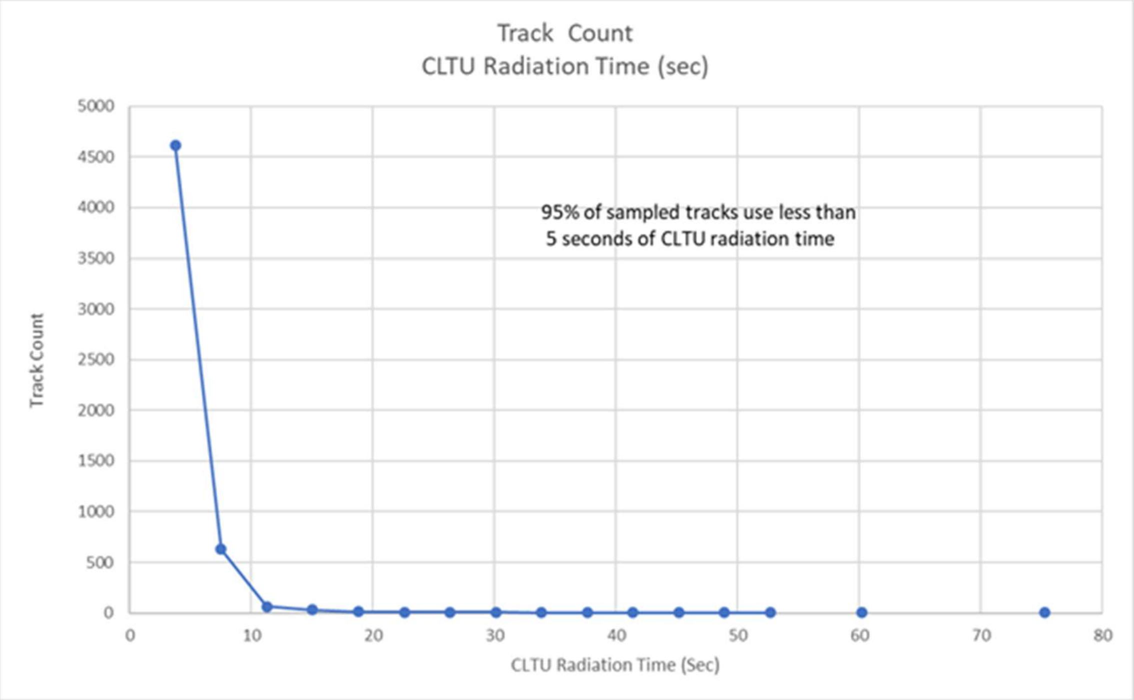

Figure 6. Track Count as a Function of CLTU Radiation Time

The second feasibility issue pertains to the compatibility of forward-error correction coding with CLTU multiplexing. Today, when a spacecraft radio is receiving an uplink, it first searches for the CLTU's start by correlating the received symbols with a known sequence referred to as the "start sequence" (**Figure 7**). On the basis of this start sequence, it then decodes the codewords to reconstruct the frames contained within the CLTU. It then examines the frame headers to see whether or not they contain the proper spacecraft ID. For those frames that do, the radio then proceeds to further process them. Analysis to date suggests that, if all the smallsats participating in a MUPA session utilize the same coding and associated start sequence (e.g., 16 bits for BCH[†] code or 64 bits for the two LDPC[‡] short codes), as well as the same type of transfer frame, there should be no issue with receiving the multiplexed CLTUs and sorting them out by spacecraft ID. Even for the case where different MOCs are sending CLTUs to the multiplexer with codes, start sequences, and transfer frames that differ from one another, analysis suggests that the MUPA session could still work (though, careful design and testing of the shared uplink would be needed to verify the correct system operation). Each spacecraft would search for the start sequence corresponding to its own code, and then filter on the spacecraft ID from the reconstructed frames. One proviso, however, would be that each MOC's ground system make sure that each frame is always entirely contained within a single CLTU. Otherwise, the portion of the frame lacking the spacecraft ID in the header would not filter correctly, causing it to be discarded and making the other part useless. Although the current CCSDS[§] standard does not explicitly prohibit splitting a frame across multiple CLTUs, for MUPA this prohibition is required to ensure proper functioning. Hence, LDPC long codes that involve a split-frame implementation may not be compatible with MUPA. Implicit in all of this, of course, is the notion that the missions participating in a MUPA session are all using the same modulation scheme and the lowest common-denominator data rate. While this may not be an absolute necessity if there are no split-frame situations, not doing so entails a level of complexity that the team did not think wise to embed into MUPA's workings before testing the more straightforward cases.

[†] BCH code = Bose–Chaudhuri–Hocquenghem code, a widely used, multi-bit random-error-correcting code.

[‡] LDPC code = Low Density Parity Check code, a linear error-correcting block code suitable for transmitting data over a noisy channel.

[§] CCSDS = Consultative Committee for Space Data Standards, a multi-national forum for the development of communications and data systems standards for spaceflight.

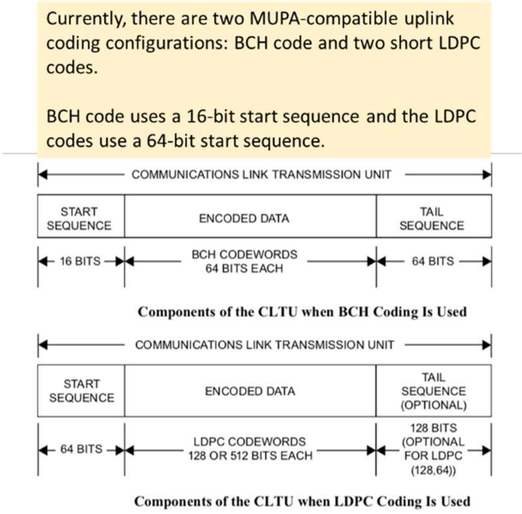
Currently, there are two MUPA-compatible uplink coding configurations: BCH code and two short LDPC codes.

BCH code uses a 16-bit start sequence and the LDPC codes use a 64-bit start sequence.

Figure 7. CLTU Structure for BCH and LDPC Short Code Coding Schemes. Frames Are Entirely Contained Within Such CLTUs.

2.2 Designing FMUX

Prior to developing a prototype system, a notional concept of operations had to be sketched out, system operational requirements defined, and functional requirements enumerated. **Figure 4**, shown earlier, summarizes the notional operations concept. While the MUPA requirements and those specifically allocated to FMUX are too numerous to list, there are four key tenets that have guided our design. First, MUPA sessions are typically complements to MSPA sessions. Second, as part of scheduling these MUPA-MSPA sessions, session participants are assumed to agree to a common waveform and lowest common-denominator data rate – an agreement that would likely be worked well in advance via the DSN's scheduling system. Third, missions desiring to participate in MUPA sessions should not have to make any significant changes to their ground systems in order to do so. And, fourth, the FMUX interfaces with the DSN should be designed in a way that minimizes any impact to existing DSN systems and services. With these four tenets in mind, the FMUX system was developed.

2.3 Description of the System

As illustrated in **Figure 8**, FMUX receives CLTUs from disparate Mission Operations Centers (MOCs), multiplexes them and outputs a single stream of interleaved CLTUs. All of the participating spacecraft that are in-beam receive this multiplexed transmission, with each spacecraft needing to differentiate which commands and/or data are intended for it on the basis of the spacecraft ID within each transfer frame header.

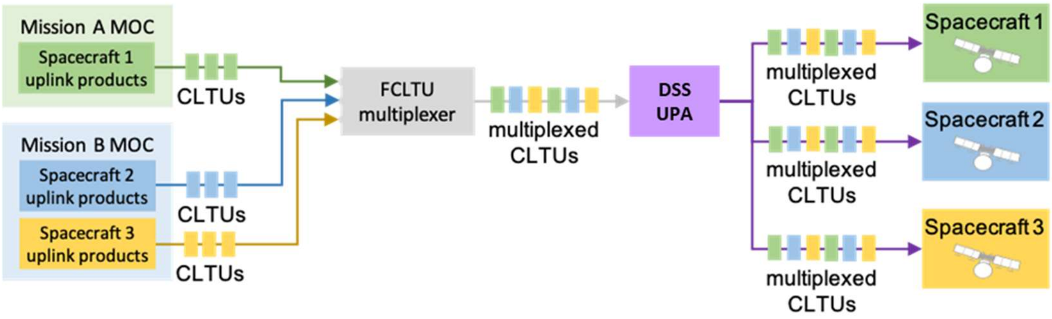

Figure 8. FCLTU Multiplexer Functional Flow (forward).

In delivering this stream of interleaved CLTUs to the Deep Space Station (DSS) Uplink Processor Assembly (UPA), FMUX utilizes what is called the Forward CLTU (FCLTU) service -- a Space Link Extension (SLE) transfer service[**] that enables a mission to send CLTUs to a spacecraft, as specified in the SLE Forward CLTU CCSDS standards (Bluebook version 4).[5] The interaction between each MOC and FMUX, and between FMUX and a DSN antenna is shown in **Figure 9.** Each MOC implements an SLE user and establishes an SLE-FCLTU session with an SLE provider running on FMUX. This session is then used by FMUX to deliver the stream of multiplexed CLTUs to the DSS UPA in much the same way that it would happen if no MUPA session was scheduled. To do so, FMUX instantiates a DSN-internal SLE user that connects to the SLE server running in the UPA associated with the radiating antenna. This latter connection is then used by FMUX to deliver the stream of multiplexed CLTUs to the antenna and obtain reports on radiation status. Finally, FMUX stores the state information used to associate SLE sessions from MOCs with their CLTUs, as well as the SLE PDUs[††] used in the internal connection to the UPA. This allows FMUX to forward reports arriving from the UPA to the appropriate MOC – thus, making it seem from the MOC's perspective that its CLTUs are traveling directly to the UPA without FMUX in the loop.

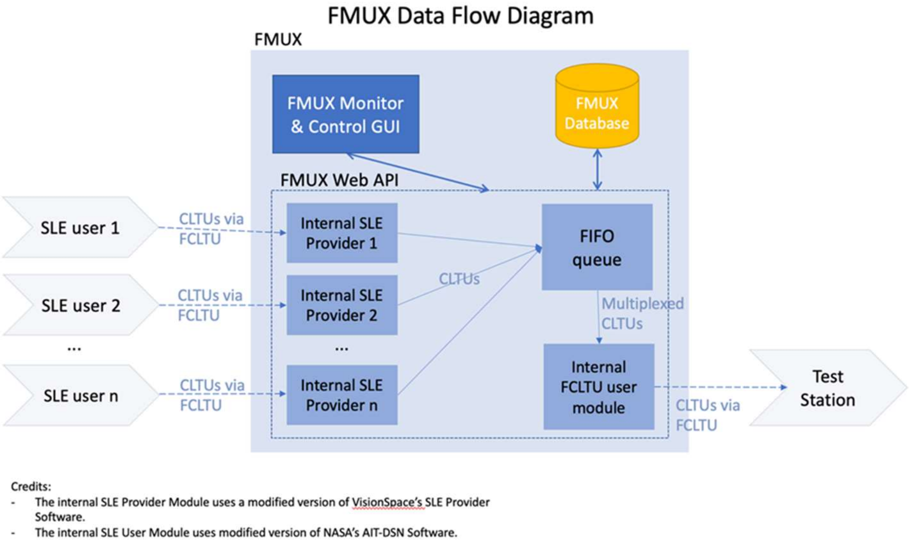

Credits:
- The internal SLE Provider Module uses a modified version of VisionSpace's SLE Provider Software.
- The internal SLE User Module uses modified version of NASA's AIT-DSN Software.

Figure 9. FMUX Design.

The operators in each Mission Operations Center (MOC) need to receive information regarding the status of the CLTUs sent to FMUX and then onto the DSS as part of the multiplexed CLTU stream. This information includes their placement in the uplink queue and subsequent radiation status. (Note that the radiation status of a CLTU describes whether it was sent from the DSN to the spacecraft and does not include any data regarding whether or not it was received by the spacecraft.) FMUX automates this flow of return data using basic SLE FCLTU operations as shown in **Figure 10**. To do this, FMUX must convey invocations and status returns of the Forward CLTU service operations between the service users and providers. Key service operations defined in the CCSDS standard have been implemented as part of the FMUX system to ensure adequate functionality between incoming and outgoing connections to and from FMUX. An invocation requests a particular service operation to be carried out and can be either confirmed or unconfirmed. A confirmed operation requires the performer to return a report of the outcome. To support such a capability, the FMUX utilizes a database, coined the "accountability database," to trace the lineage of data and allow for compiling reports for confirmed actions to the originating SLE user. Each and every CLTU that flows through the system is captured with its appropriate identifiers and timestamped at each stage to confirm with

---

[**] The Space Link Extension (SLE) transfer service is a CCSDS standard for transporting spacecraft forward and return data between tracking stations, mission operation control centers, and data-user facilities irrespective of their organizational affiliation or the particulars of their implementation. To do this, SLE makes use of protocols that utilize TCP/IP and run over existing communications infrastructure. In so doing, SLE enables unaffiliated, CCSDS-compliant ground stations to back each other up in the event of a failure or to augment one another when excess antenna demand necessitates it.

[††] PDU = Protocol Data Unit

great confidence that no data are lost and/or altered. The system does not interpret or modify the contents of a CLTU. Each CLTU is transmitted bit for bit.

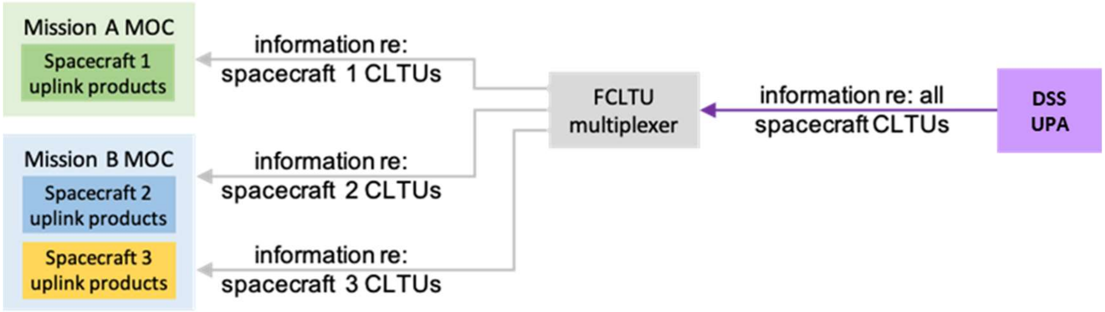

Figure 10. FCLTU Multiplexer Functional Flow (return).

Since the SLE Forward CLTU service requires a strict sequence of operations to support data integrity, sequence numbers and invocation identifiers must be assigned to each transferred CLTU. FMUX uses this metadata store in the accountability database to accurately relay information about each DSN-deposited/radiated CLTU to the CLTU's original SLE user. Async notifications and transfer data confirmation are sent to the correct originating participant by generating FMUX sessions and collecting vital CLTU metadata during the data transfer phase. Once a confirmation for data transfer is received by FMUX from the antenna transmit subsystem, FMUX translates the SLE-PDU to match the appropriate sequence number and invocation identifier for a particular originating user's CLTU. With this design, we hope to minimize modifications to existing systems such as the Advanced Multi-Mission Operations System (AMMOS) SLE Command Client (SCC), AMMOS Instrument Toolkit (AIT) Deep Space Network (DSN) Interface, and those of the DSN. Major capabilities should fit in seamlessly with existing tools and services that follow the FCLTU service specification. Data provenance also gives users better confidence that FMUX operates properly and allows for ease of debugging when an issue does arise. Users can easily spot CLTUs that were dropped based on what the accountability database reports.

2.4 The Web Application Programming Interface (API)

One of the main features that FMUX provides is a web Application Programming Interface (API). This allows users the ability to utilize FMUX's functionality without being tied to the default graphical user interface (GUI) that FMUX provides. The FMUX web API exposes the following functionality through server endpoints:

- The ability to perform implemented FCLTU service operations (i.e., Bind, Start, Stop, Unbind, Peer Abort, Status Report) from FMUX to the DSSUPA.
- The ability to invoke an SLE peer-abort from FMUX to any connected SLE user.
- The ability to query the following monitor data from FMUX:
    - FMUX Logs
    - SLE-User CLTU Metadata
    - Mission Connection Details
    - Station Connection Details
    - Insight into the multiplexer/queue
- The ability to add and modify FMUX mission and station configurations.

All responses from the FMUX API are returned in JavaScript Object Notation (JSON).

2.5 The Graphical User Interface

Operators need to be able to configure the Forward CLTU multiplexer, monitor its behavior, and have sufficient control to identify and resolve issues within the tool or with the external interfaces through a simple graphical user interface. FMUX provides this GUI for monitor and control of the system as shown in **Figure 11**. It allows for a more intuitive way to configure the application and services without having to edit multiple configuration files somewhere on a server. It also provides detailed reports regarding CLTU production and uplink status. And, the GUI provides ways to bind/unbind with the DSS UPA, start/stop CLTU radiation, and monitor each MOC's FCLTU states and how they have transferred and queued into the system. All the operations and interactions between external services are also logged by the system and available to browse via this FMUX GUI, providing a way for operators to quickly discover warnings or errors and try to resolve issues on the fly with very little delay.

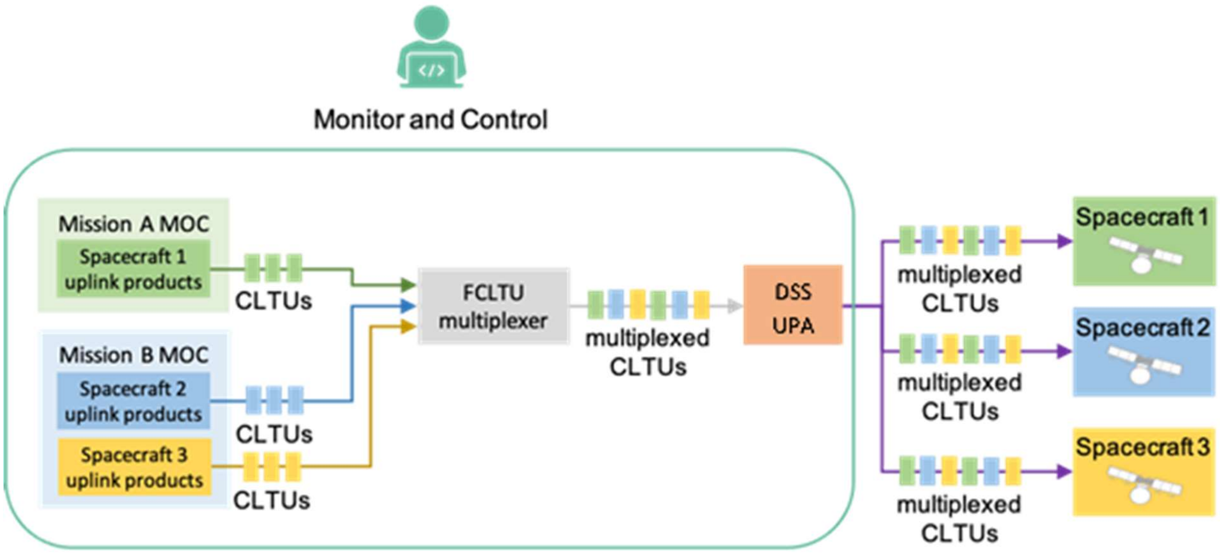

Figure 11. The Monitor and Control Span of the FMUX GUI.

2.6 Development and Testing Approach

An end-to-end agile software development approach is currently being used to implement FMUX. The use cases and requirements are developed incrementally, and the associated software units are also developed in an agile manner. The FMUX software installation process has been Dockerized to provide users with a quick and easy way to install the application for use. Docker is used to: (1) package an application with all the essential dependencies (such as libraries) needed to run it and (2) ship it as one package, a "container." The Docker container allows another user to quickly recreate the environment needed to run FMUX. This results in fewer deployment issues and saves time. Hence, we have Dockerized the testing environment to run all integration tests in one simple step. We have also implemented a large set of critical integration tests and used a continuous integration and validation tool for automated validation. Unit tests are automatically conducted as new software units are added to the code. Continuous integration helps reduce the chances of new bugs being introduced during each development phase.

For full integration testing, a virtual machine on JPL's Flight Operations network has been used to run FMUX. Tests using DTF-21[‡‡] have been conducted in this environment. These tests have helped identify and resolve issues as well as demonstrate FMUX's feasibility to the task's stakeholders. In particular, two types of tests have been conducted:

- Functional tests. Their main purpose is to test the basic functionality of FMUX. Success is declared when the functionality can be successfully executed.
- Performance tests. Their main purpose is to test the capabilities of FMUX and help identify performance "bottlenecks" or breaking points.

[‡‡] DTF-21 is the DSN Test Facility in Monrovia, California that provides, for testing purposes, the same uplink and downlink equipment as that used at the three DSN sites in Goldstone, California, Madrid, Spain, and Canberra, Australia – with the exception of the antennas themselves and their Low Noise Amplifier (LNA) front-ends.

Examples of functional tests conducted include:

- Bind/start/upload/stop/unbind with Station
- Bind/start/upload/stop/unbind from one MOC
- Bind/start/upload/stop/unbind with 2 MOCs
- Confirmation of CLTU data transfer return information being sent to the correct SLE user with accurate identifiers
- Confirmation of async notifications being sent to the correct SLE user with accurate identifiers.

Examples of performance tests conducted include:

- 2 MOCs simultaneous with interleaved CLTUs
- 4 MOCs simultaneous with interleaved CLTUs using typical CLTU information
- 1000 CLTUs from 2 connections - successfully interleaved data
- Testing of over 3000 CLTUs in one session.

End-to-end demonstrations involving actual mission telecommand data and flight/testbed hardware are currently being pursued with the SunRISE and EscaPADE missions. To help prove FMUX's mission operational readiness, ground demonstrations are being planned, scheduled, and performed, where FMUX takes the FCLTUs intended for multiple spacecraft, multiplexes them, and sends them on to DTF-21 for "radiation," while providing each original mission user with the appropriate information on each CLTU's progress to the point of radiation. The participating missions then use their flight/testbed hardware to "receive" the "radiated" CLTUs and verify correct receipt of the associated data. To the extent that these tests go well and any "bugs" that arise are resolved, these demonstrations will then progress to actual in-flight MUPA demos.

3. Ground System Interfaces

At present, most DSN systems implicitly assume that a single uplink pairs with a single downlink. In the case of MUPA, however, a single, interleaved uplink pairs with multiple downlinks as part of an MSPA session. This difference creates at least three significant challenges: (1) getting the DSN's Service Management System (SMS) to create and load the appropriate Support Data Packages (SDPs) needed to schedule and configure the MUPA uplink, (2) having SMS appropriately populate an API[§§] that tells FMUX when spacecraft participating in a MUPA session are behind a planet or otherwise out of view so that their CLTUs are not multiplexed and radiated during such time periods, and (3) ensuring that the Tracking Data Delivery System (TDDS) is still able to deliver 2-way Doppler and ranging data to the navigation teams when there are non-standard turnaround ratios involved. In the subsections that follow, we address each of these three ground system interface challenges in turn.

3.1 Making SMS Work with MUPA

The multiplexed output from FMUX looks to the DSS UPA like the normal input from a single user. This virtual single user also needs to have a virtual spacecraft ID associated with it in the SMS's Service Preparation Subsystem (SPS) so that the proper scheduling and execution of the MUPA session can occur. Other recommended changes that need to occur in relation to this subsystem are summarized in **Figure 12.** The first change involves providing a MUPA flag in the Service Scheduling Software (SSS) to allow for simultaneous uplink scheduling, generally as part of an MSPA. An appropriate schedule would then be generated and sent to the Service Data Management (SDM) subsystem. In order for SDM to properly utilize this schedule, changes need to occur that provide a multiple-concurrent-uplinks alternative to the sequential commanding and tracking approach for an MSPA currently assumed in the software. With this change instituted, the SDM would then generate an SDP for the virtual spacecraft. The SDP would be sent to the Support Data Loader (SDL). From there, the uplink would execute within the various DSN systems as normal -- the only exception being that "throw events" (i.e., "on-the-fly" changes to the uplink) by a participating mission would not be allowed since it would also impact the other missions participating in the MUPA session. If a given spacecraft in a MUPA session should need emergency services, it would be required to exit the session and establish a separate communications link.

[§§] API = Applications Programming Interface

To date, a prototype SDP modifier program has been developed to allow for SPS-UPL interface testing. Some of this testing has occurred at DTF-21, concurrent with the FMUX testing. Additional scheduling considerations that still have to work their way into SPS and the SDP generation include a method for ensuring that the MUPA participants in an MSPA-MUPA session all utilize the same waveform and lowest common denominator data rate. And, an interface between the SPS and FMUX needs to be developed so that both the FMUX forward flow to the UPA and the return flow to the MOCs are consistent with what has been scheduled via SMS.

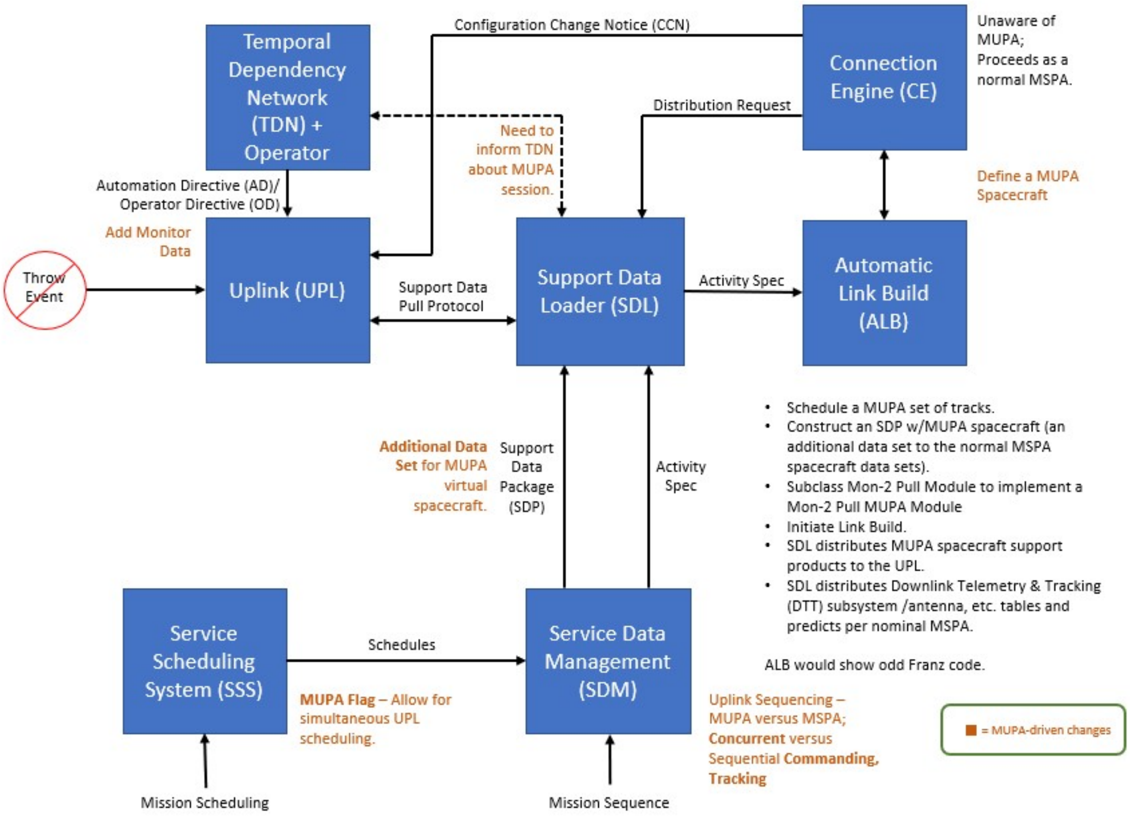

- Schedule a MUPA set of tracks.
- Construct an SDP w/MUPA spacecraft (an additional data set to the normal MSPA spacecraft data sets).
- Subclass Mon-2 Pull Module to implement a Mon-2 Pull MUPA Module
- Initiate Link Build.
- SDL distributes MUPA spacecraft support products to the UPL.
- SDL distributes Downlink Telemetry & Tracking (DTT) subsystem /antenna, etc. tables and predicts per nominal MSPA.

ALB would show odd Franz code.

Figure 12. SMS-Related Subsystem Changes Needed to Support MUPA

3.2 Developing an SPS-FMUX API for Participating Spacecraft Visibilities

Spacecraft orbiting planets and other celestial bodies are not always visible to the Earth. Sometimes, their orbits take them behind the bodies that they are orbiting. When this occurs, it would not make sense to include CLTUs bound for such spacecraft in the multiplexed CLTU stream. So, FMUX needs a way to know when such CLTUs need to be held back in a queue until visibility returns to the DSN station(s) participating in the MUPA session. To this end, a prototype API has been developed that would draw upon the participant trajectory information available from SPS, ascertain the visibility outages relative to the participating DSN station(s), and communicate this information to FMUX to apply to its CLTU queueing management. Our next step involves incorporating the ability to make use of this prototype API in FMUX – then, testing it with an appropriate MUPA session scenario in DTF-21.

3.3 Ensuring that TDDS Can Deliver 2-way Doppler and Ranging

Because the FCLTUs intended for each of the spacecraft participating in a MUPA session are multiplexed on a single uplink frequency associated with a virtual spacecraft, that single uplink frequency has to be turned around on each spacecraft's own downlink frequency in order to obtain 2-way Dopper and ranging for each of the participating spacecraft. Aside from each spacecraft's radio needing to accommodate non-standard turnaround ratios, the DSN's TDDS needs to have a way to communicate to the navigation team what the non-standard turnaround ratio is.

Currently, the downlink Doppler and ranging data are sent from the Signal Processing Center (SPC) associated with the participating ground station to the TDDS at JPL. Using internal, manually updated tables of spacecraft turnaround ratios, the TDDS supplements the radiometric data it sends to the navigation team with this turnaround information. For MUPA, we recommend that the turnaround ratios be provided in an automated manner that dispenses with the internal table. In this proposed approach, the SPS would draw upon, as appropriate, the information for each spacecraft being delivered to the Downlink Tracking & Telemetry (DTT) subsystem. The SPS would then send this information to the SPC of the participating ground station which would then send it, along with the received radiometric data, to the TDDS. The TDDS would then append the normal file conversions that it sends to the navigation team with the appropriate spacecraft turnaround ratio information. While it sounds a bit complex, the proposed approach would minimize the overall magnitude of any required changes to DSN systems. The current and proposed new processes are summarized in **Figure 13**.

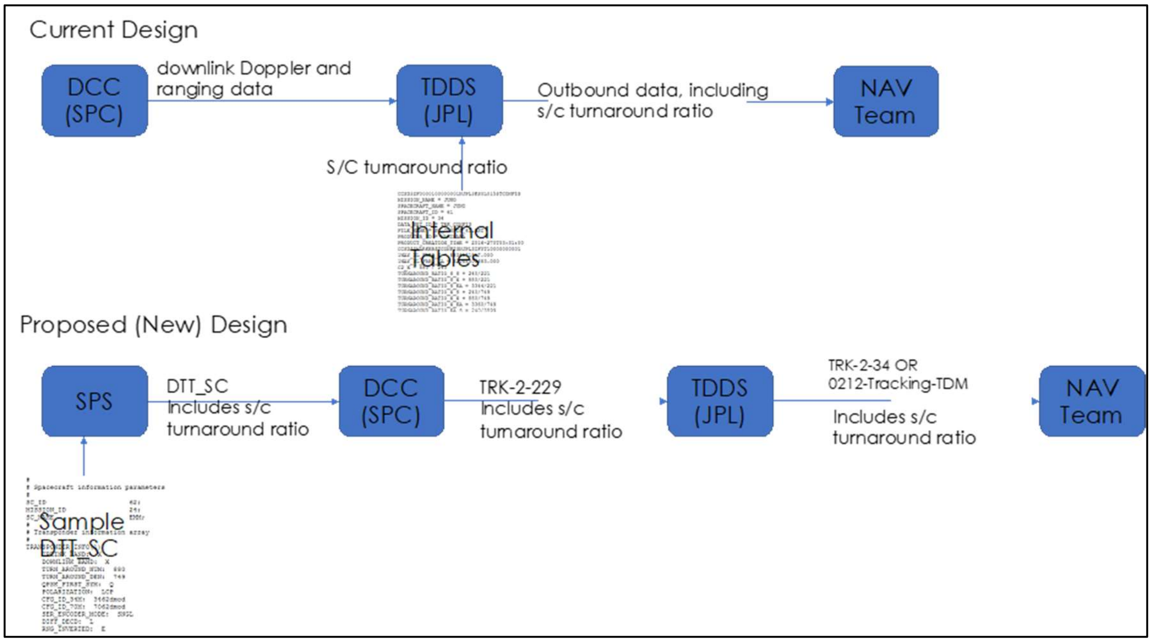

Figure 13. Proposed Change in the Way TDDS Appends Radiometric Data with Turnaround Ratio Information

4. Spacecraft Radio Considerations

MUPA's multiplexed-FCLTUs-on-a-single-uplink-frequency approach poses five key requirements on the participating spacecraft: (1) an onboard capability to compensate for the relative Doppler dynamics in the uplink signal, (2) a spacecraft radio capable of accommodating non-standard turnaround ratios, (3) acquisition of an additional uplink frequency assignment for shared use during a MUPA session, (4) the ability of the spacecraft radio to switch between the mission's standard uplink frequency and the shared, MUPA-session uplink frequency, and (5) a spacecraft radio or C&DH capable of filtering out received transfer frames on the basis of spacecraft ID.

4.1 Onboard Doppler Compensation

The first requirement, onboard Doppler compensation, probably poses the largest technical hurdle. When multiple, in-beam spacecraft participate in a MUPA session, they may each see the same, single uplink frequency with a different Doppler shift and associated rate if their trajectories or orbits around a planet or other body are significantly different. The Mars example shown in **Figure 14** provides a good case in point.[6] As **Figure 15** shows, the relative Doppler difference between Mars Odyssey and other in-beam spacecraft is, at times, on the order of + or – 150 kHz. Even with the ground station sweeping the carrier, the majority of the spacecraft would be unable to acquire and/or maintain lock.[***] In 2016-2018, three onboard Doppler compensation approaches were investigated to remedy this issue: (1) incorporating a step-and-sweep capability in the radio, (2) incorporating an FFT[†††] search capability in the radio, and (3) incorporating a trajectory-aided tuning capability in either the radio or the spacecraft Command & Data Handling (C&DH) subsystem. Extensive initial simulations and tests focused on the step-and-sweep and FFT performance for different frequency offsets, Doppler rates, symbol rates, signal strengths, loop bandwidths, FFT search ranges, etc.[7] Performance was measured in terms of acquisition times, ability to maintain lock, and associated static phase error/phase jitter. The results of these efforts indicated both approaches to onboard Doppler compensation to be technically feasible and suggested that the next logical step would be to try implementing one or more of these approaches in an actual software defined radio.

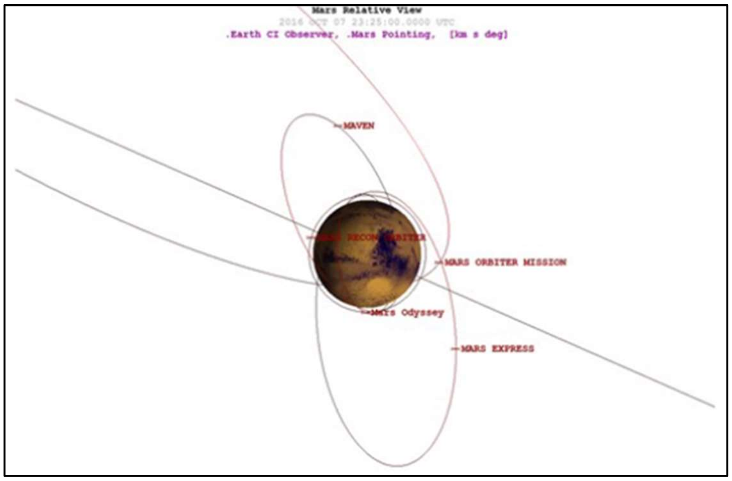

Figure 14. A Mars Example of the Diversity of Orbits Associated with Spacecraft that Are All Within the Beam of a Single 34m Antenna [Credit: Morabito, et al., 2018]

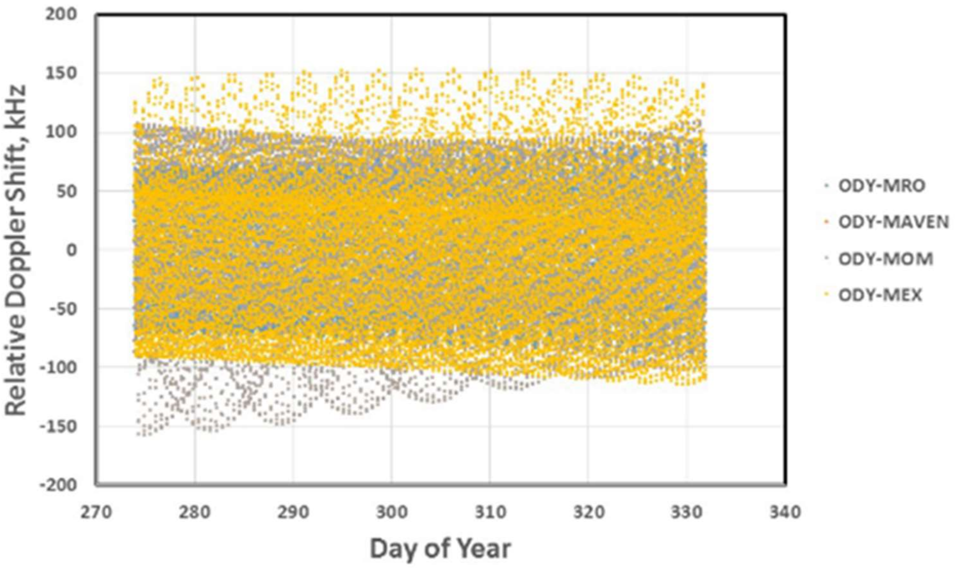

Figure 15. Relative Doppler Shift Between Odyssey and Other Mars Spacecraft as a Function of Time [Credit: Morabito, et al. 2018]

[***] A spacecraft can temporarily acquire a signal and then lose lock depending upon the dynamics or other changes. Considerations such as loop type can be important.

[†††] A Fast Fourier Transform (FFT) allows the analysis of a time-varying signal in terms of its spectral components, or the reconstruction of a signal from such components.

In 2018-2020, we began working with NASA's Iris Radio task at JPL to explore implementation of a sweep capability within the Iris radio. As illustrated in **Figure 16**, a frequency sweep generator was programmed into Iris's FPGA[‡‡‡] with user-definable inputs for sweep rate, sweep range, and dwell time (for stepping if desired). The sweep rate was user-definable between 24 Hz/sec and 200 Hz/sec. The allowable sweep range was 0 Hz to 524 kHz. And, the dwell time was set to zero seconds since the uplink sweep traditionally performed on the ground is continuous. In subsequent DTF-21 testing, this new Iris sweep capability was able to detect and lock onto the uplink carrier under a variety of assumed Doppler offset and received power conditions. In the most extreme of these assumed conditions, the uplink carrier had a Doppler shift of -150 kHz and a received power of -130 dBm. Using a sweep rate of 200 Hz/sec, Iris took 250 seconds to achieve lock, after which 10 no-op commands were successfully transmitted to it.

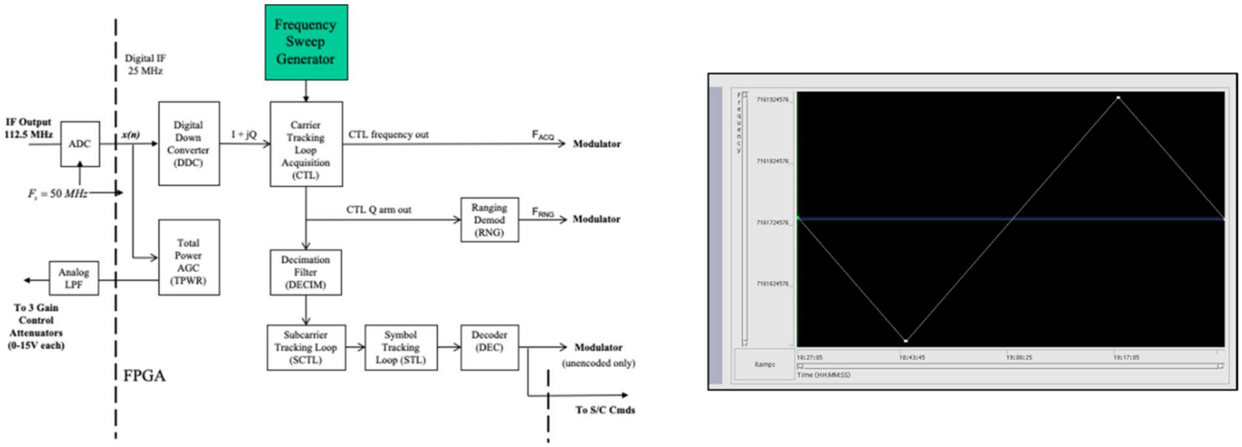

Figure 16. Addition of Frequency Sweep Generator into Iris FPGA and Example Sweep Pattern with Dwell Time Set to Zero

While sufficient for many potential MUPA use cases, additional simulations in the 2018-2021 timeframe suggested that the basic sweep capability discussed above might not be sufficient in situations involving high Doppler rates and/or very large Doppler offsets.[8] As shown in **Figure 17**, just relative to Mars-center, a spacecraft can have significant Doppler rates associated with the change in the Doppler offset as a function of time. At times, these Doppler rates can also have significant Doppler accelerations associated with them. Such Doppler dynamics hinder the ability of the "normal" Phased-Lock Loop (PLL) to achieve and/or maintain lock even though the onboard sweep has located the signal at the offset frequency. Here, we are defining the "normal" PLL as one which utilizes a Type 2 (2nd order polynomial) loop filter, where the purpose of the loop filter is to help drive the phase-error output from the PLL mixer to as close to zero as possible. To the extent that the literature suggested that a Type 3 (3rd order polynomial) loop filter might better manage high Doppler dynamics, our efforts then focused on conducting simulations and Iris radio tests that combined the onboard sweep capability with a Type 3 PLL.[9] These simulations and tests demonstrated that the Type 3 PLL definitively outperforms the Type 2 PLL in terms of acquisition time and static phase error when Doppler rates are high, and particularly when Doppler acceleration is present. In the absence of those special Doppler conditions, the Type 2 PLL performs at least as well. Given these findings, we recommended that the Iris radio be equipped with a Type 3 PLL option for use in conjunction with the sweep function to accommodate situations with high Doppler rates and/or Doppler accelerations.

[‡‡‡] A Field Programmable Gate Array (FPGA) is an integrated circuit containing an array of programmable logic blocks that can be reconfigured to perform a variety of different operations.

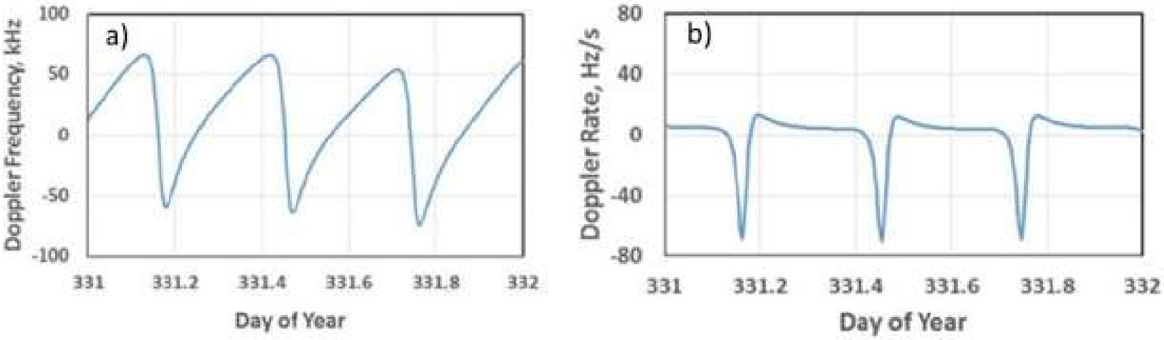

Figure 17. (a) Doppler frequency versus time for example spacecraft orbit. (b) Doppler rate versus time for the case of a Mars orbiter (relative to Mars-center). (Credit: Morabito, et al. 2021)

Over the 2020-2021 time period, we also wanted to see if we could achieve faster lockup times for large Doppler offsets than what our basic sweep capability would allow. To do this, we designed, simulated, and tested an FFT implementation in Iris to replace the sweep capability. Key features included: an FFT acquisition with averaging capability, an ability to switch from acquisition mode to tracking mode in the carrier tracking loop, and an optimization of the carrier tracking loop. In designing the FFT acquisition with averaging, a self-imposed limit of + or – 2 MHz around the best-lock frequency was used. To reliably capture the carrier down to approximately the carrier tracking threshold, the ability to average up to 4,096 FFTs was built into the firmware programming. With respect to the ability to switch from acquisition to tracking, the carrier tracking loop was designed with the capability to capture far away carriers. This ensures that there is no need to sweep in Iris (or, for a single spacecraft uplink, sweep on the ground via the DSN), since the FFT can estimate the carrier frequency up to an error of +/- 750Hz, and the acquisition bandwidth can then capture the carrier -- at which point, the carrier tracking loop can switch to a much narrower tracking loop bandwidth to improve the signal-to-noise ratio and associated performance. Optimization of the carrier tracking loop involved reworking its FPGA implementation to remove unnecessary delays. **Figure 18** provides a summary illustration of the FFT implementation within Iris. In subsequent DTF-21 testing, combined with a Type 3 PLL, this new prototype version of the Iris radio was able to acquire a static carrier offset of 1.5 MHz from the best-lock frequency in less than a second. The lock occurred at a signal power of -70 dBm and was maintained down to approximately -143 dBm (with the FFT lock remaining at significantly lower power). During the test, the carrier tracking loop bandwidth was configured with a 200 Hz acquisition mode and a 20 Hz carrier tracking mode. In a separate test pertinent to the Doppler rate, the improved Type 3 carrier tracking loop was able to achieve lock at 2000 Hz/s at -80 dBm. And, at 200 Hz/s, it achieved lock with a 10 dB lower signal power threshold than was the case without the improved carrier tracking loop. A final test of the FFT and improved carrier tracking loop combination was conducted in which the the signal was offset by 1.5 MHz with a Doppler rate of 1000 Hz/s. Iris acquired the carrier in less than a second at a signal power of -120 dBm. The number of FFTs averaged for the test was set to 128.

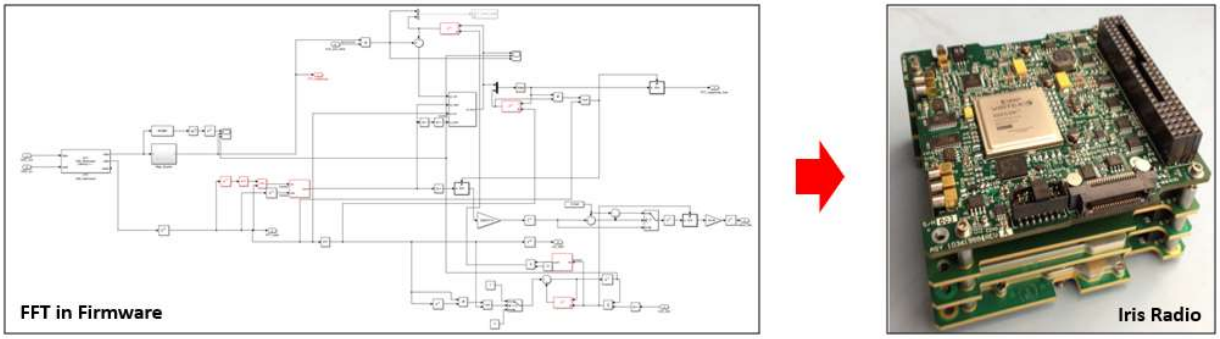

Figure 18. Example Illustration of FFT Incorporation into the Iris Radio

At present, through the efforts of the JPL-led Iris radio task team, the above improvements to the Iris radio have now been incorporated by the Space Dynamics Laboratory (SDL, located in North Logan, Utah) into a commercially available version of the radio. Hence, a radio with an onboard Doppler compensation capability for a MUPA uplink is now commercially available. Nonetheless, to make a potential DSN MUPA service as broadly accessible to the future mission user base as possible, we need to have an onboard uplink Doppler compensation capability available in more software-definable radios than just Iris. To this end, we have begun looking at "smart" sweep algorithms with fast acquisition times that might be readily applied to other radios. We are also working to test one or more prototype trajectory-aided tuning techniques. In the "smart" sweep algorithm realm, a separate JPL research effort has developed an "adaptive sweeping carrier acquisition and tracking" algorithm that shows substantial promise for rapid uplink acquisitions under high Doppler dynamic conditions.[10,11] Preliminary simulation comparisons between this algorithm and the FFT-Type 3 PLL combination suggest comparable acquisition-time performance over a wide range of test conditions, with the algorithm running just slightly slower than the FFT- Type 3 PLL (1-2 sec. vs. 0.5 sec) when the signal-to-noise ratio is very low. In FPGA resource utilization comparisons, however, the adaptive-sweep algorithm made more efficient use of the FPGA. So, a logical next step might be to prototype an implementation of the adaptive-sweep algorithm in the Iris radio for additional testing, perhaps replacing the basic sweep capability originally implemented.

In the trajectory-aided tuning realm, a prototype program has been written that converts spacecraft trajectory information into uplink frequency predicts, then expresses those frequency predicts in the form of a memory-efficient Everette table of coefficients. The spacecraft C&DH or software-defined radio can periodically perform interpolations of these coefficients as a function of time to derive the uplink frequency predicts. These predicts can then be applied through the radio's numerically controlled oscillator to tune the receiver to the anticipated Doppler-shifted uplink frequency. On some or all of the regularly scheduled antenna passes, the table of coefficients could be updated with the latest trajectory information to ensure that the receiver remains properly tuned for the anticipated Doppler-shifted uplink. While one would probably want to also have some sort of onboard sweep or FFT capability in the event of deviations from the planned/updated trajectory information, trajectory-aided tuning might have a standalone application in the case of legacy spacecraft already in well-defined orbits. Some of these spacecraft may have the capability to tune their receivers given appropriate information from the C&DH. Current efforts are focused on ascertaining which, if any, legacy spacecraft could make use of such a technique (and, therefore, potentially be able to make use of MUPA) and whether, for newer spacecraft with software defined radios, the memory and computational resource vs. performance trade would make trajectory-aided tuning a sensible augmentation for a sweep or FFT signal search capability. For the latter endeavor, efforts are currently underway to prototype a trajectory-aided tuning capability within Iris to facilitate conduct of this computational resource vs. performance trade. As indicated previously, however, the ultimate goal associated with these efforts is to enable onboard uplink Doppler compensation and, hence, potential MUPA capability in a broader swath of spacecraft radios than just Iris.

We are also in the process of exploring the performance of these Doppler compensation techniques during extreme MUPA events such as during solar superior conjunctions. In addition to the high Doppler offsets and dynamics induced by the different trajectories, the intervening charged particles of the solar corona will induce amplitude and phase scintillation effects onto the uplink signal.

4.2 Accommodating Non-standard Turnaround Ratios

Because MUPA involves multiplexing the CLTUs intended for different participating spacecraft on the same uplink frequency, the turnaround on each spacecraft's assigned downlink frequency for 2-way Doppler and ranging entails a non-standard ratio (i.e., a ratio different from the table of spacecraft transponder turnaround ratios recommended in the *DSN Telecommunications Link Design Handbook*) . And, to the extent that the uplink frequency assigned to the MUPA-MSPA session may vary, participating spacecraft radios will generally need to be equipped with variable turnaround ratios. To this end, the Iris radio was equipped with variable turnaround ratios at the same time the basic sweep and FFT capabilities were implemented. Ensuring that other software-defined spacecraft radios come equipped with variable turnaround ratios in the future may prove to be a MUPA-related CCSDS standards challenge going forward.

4.3 Acquisition of an Additional Uplink Frequency Assignment for Shared Use During a MUPA Session

NASA and other government agency missions must obtain authorized frequency assignments through the National Telecommunications and Information Agency (NTIA). U.S. commercial and university missions (not sponsored by a U.S. government agency) must obtain authorized frequency assignments through the Federal Communications Commission (FCC). For the NASA-sponsored missions, the appropriate NASA Center Spectrum Managers typically work with them to assist in obtaining their uplink and downlink frequency assignments. In the case of MUPA, however, there are two uplink frequencies – one that is uniquely assigned to a mission and one that is shared between the MUPA-session participants. The process and particulars associated with obtaining the frequency license for this shared uplink frequency is something that still has to be worked out. In May 2022, the NASA National Spectrum Program Manager was briefed on MUPA and the uplink assignment question discussed amongst the NASA Center Spectrum Managers. While they saw no insurmountable issues, they did note that each MUPA frequency will need to be coordinated with other space agencies making use of the associated band. Hence, it may not be possible to have just a single assigned MUPA uplink frequency per band for use across all potential MUPA sessions in the future; it may be necessary to have some alternative frequencies in the event there is a conflict. The designation of these frequencies and the particulars for how prospective mission participants in MUPA sessions get licensed for such shared uplink frequencies remain items to work going forward.

4.4 Ability of the Spacecraft Radio to Switch Between Individual and MUPA Uplink Frequencies

To the extent that a spacecraft must operate with an individual uplink frequency when the sole participant in a link and with a shared uplink frequency when participating in a MUPA-MSPA session, the spacecraft radio needs to be able to: (1) support at least two different uplink frequencies and (2) switch between these two frequencies when switching from a pass involving an individual communications link to a pass involving a MUPA session (or vice versa). At this point in time, most spacecraft radios lack these attributes for their direct-from-Earth link capability. Ensuring that future software-defined spacecraft radios come equipped with these capabilities is probably another MUPA-related CCSDS standards challenge going forward. [§§§]

For multi-spacecraft missions that are the sole participants in a MUPA session and wish to use the same uplink frequency for all of their spacecraft, the above issue is irrelevant (as is the MUPA frequency assignment issue). Examples of such missions include SunRISE and EscaPADE – each of which can benefit from the multiplexing of the FCLTU streams intended for each of its spacecraft. Because each of these missions already plans to use the same uplink frequency for each of its spacecraft, they make excellent candidates for the demonstrations of the FCLTU Multiplexer now in work.

4.5 *Ability to Filter Out Received Transfer Frames on the Basis of Spacecraft ID*

Since all of the spacecraft participating in a MUPA-MSPA session receive the same FCLTU-multiplexed uplink signal, they will all demodulate and decode it to obtain the same set of transfer frames. They will then have to filter out the frames that pertain to other spacecraft and only apply the contents of those frames intended for their spacecraft. This filtering will need to be done on the basis of the spacecraft ID in the transfer frame headers and may either occur in the spacecraft radio or in the C&DH. In the case of the Iris radio, this capability to filter frames on the basis of spacecraft ID already exists. In other radios, that is not necessarily the case. But, somewhere, either via the radio or the spacecraft C&DH, the capability usually exists in order to preclude the receipt of errant spacecraft commands from either accidental or intentional interference. So, from a MUPA perspective, this filtering capability is not likely to be a barrier to participation for most missions. However, from a mission operations perspective, the gound system and associated operators may need to build the expectation of a high frame rejection rate during MUPA sessions into their processes and procedures.

4.6 Security

On October 29, 2019, NASA released its “Space System Protection Standard,” NASA-STD-1006.[12] Among other things, this standard mandates that NASA programs and projects “shall protect the command stack with encryption

---

[§§§] There may also be a technical challenge for the radio hardware if the two frequencies are far apart. In such a case, certain RF elements like filters might not have sufficient bandwidth to work well at both frequencies.

that meets or exceeds the Federal Information Processing Standard (FIPS) 140, Security Requirements for Cryptographic Modules." The objective of this standard is to ensure missions "maintain command authority to prevent unauthorized access" so as to prevent such access from resulting in "mission loss and/or damage to other space systems." The manner in which missions implement this command encryption, however, may vary. CCSDS has issued the Space Data Link Security (SDLS) protocol as a recommended means for "providing authentication and/or confidentiality to the contents of transfer frames."[13] While the content of these frames is encrypted, the header and associated spacecraft ID information is not. Hence, in a MUPA session in which the participants are all using SDLS, they can still filter these frames on the basis of spacecraft ID, then go on to decrypt only those that are appropriate. However, if a mission elects to use bulk encryption (i.e., encryption at the physical link layer) the CLTU, the associated acquisition and idle sequences, the transfer frame headers, and the transfer frame content are all encrypted. At present, it is not clear how MUPA could be made to work with bulk encryption – particularly if the participating spacecraft all belong to different projects. This topic is currently the subject of on-going study.

5. Conclusion

Demand for the DSN's antennas has grown significantly over the past 15 years and is expected to grow even more rapidly over the next ten. Hence, it behooves the DSN to increase the utility of its existing antennas. Historically, one way this has been accomplished is through MSPA. But, MSPA currently allows multiple, in-beam spacecraft to make simultaneous use of the DSN antenna only on the downlink. The uplink still occurs one spacecraft at a time. With MUPA, we aim to change this by enabling virtually simultaneous commanding to all of the spacecraft participating in the simultaneous MSPA downlink. In so doing, all of the MUPA-MSPA session participants will also be able to simultaneously obtain 2-way Doppler and ranging data.

Three areas of endeavor are key to realizing these objectives: developing FMUX, modifying the ground system interfaces needed to seamlessly schedule and execute the MUPA sessions, and augmenting spacecraft radios and/or the associated spacecraft C&DH subsystems with the abilities needed to acquire and maintain lock with the shared single-frequency MUPA uplink (as well as filter out the received frames on the basis of spacecraft ID). With respect to FMUX, we have succeeded in developing, testing, and refining the software to the point where we are now beginning the process of ground-testing it with two flight projects – SunRISE and EscaPADE. After working out any issues arising from these tests, the next step will be to conduct MUPA flight demonstrations with it. These will occur once the SunRISE and EscaPADE spacecraft have successfully launched and reached the stages in their respective missions when such demonstrations can be safely conducted.

With respect to modifying the ground system interfaces, we have identified some of the key changes in the DSN's Service Management System and Tracking Data Delivery Service needed to successfully implement MUPA. A prototype Support Data Package modifier has been developed to facilitate testing of the various SMS-FMUX-UPL interfaces needed to support upcoming ground and flight demonstrations. In addition, a prototype API has been developed to allow SMS to inform FMUX as to when each spacecraft participating in a MUPA session is occulted by the planet or body it is orbiting so that its FCLTUs do not get multiplexed during such times. For TDDS, a conceptual solution to a turnaround-ratio table impediment to the delivery of 2-way Doppler and ranging data to the navigation teams has been identified and is now being worked. All of these efforts are aimed at helping to define a potential "path" for MUPA's operational implementation within the DSN.

With respect to augmenting the spacecraft radios and/or spacecraft C&DH with the abilities needed to acquire and maintain lock with the shared single-frequency MUPA uplink, significant effort has gone into researching, simulating, designing, testing, and implementing an onboard Doppler compensation capability within the Iris radio. With the help of JPL's Iris radio task team, this effort has culminated in a commercially available version of Iris with an FFT-Type-3-carrier-tracking-loop-combination that, in ground tests, adequately compensates for anticipated Doppler dynamics in the single-frequency MUPA uplink. This version of Iris also includes the variable-turnaround-ratio and received-frame-filtering-by-spacececraft-ID capabilities so essential to making MUPA work. We are now actively seeking opportunities to flight-test the MUPA-related performance of this latest version of the Iris radio.

Other ongoing spacecraft radio efforts include investigating how "smart" sweep algorithms and trajectory-aided tuning might also be used to confer MUPA utility to a larger subset of spacecraft radios and/or C&DH subsystems – such that MUPA-MSPA sessions in which significant Doppler dynamics are present can involve more spacecraft than just those using Iris radios. Also, key to making Iris and other radios fully MUPA-applicable is the need for them to incorporate an ability to switch between their individually assigned uplink frequency and a shared MUPA uplink frequency. To this end, development of a MUPA-related CCSDS standard for spacecraft radios is needed – one which encompasses not only uplink switching but, also, the need for variable turnaround ratios and inclusion of the ability to filter received frames on the basis of the spacecraft ID in the frame header.

Some attention will also be needed to comply with security requirements in a MUPA concept of operations where multiple spacecraft are receiving one shared uplink stream that includes commands and uplink tracking data intended for other spacecraft. There are various technical approaches that can be adopted for this and in the future, trades will be conducted to evaluate the options and make the best design decisions going forward.

Consistent with the MUPA-related uplink provisions just discussed, we also recommend that NASA and the other space agencies participating in the Space Frequency Coordination Group (SFCG), develop a recommendation for what frequencies to use for shared uplink during MUPA sessions.

In the interim, FMUX, ground system interface work, and MUPA-related spacecraft radio work are far enough along to support demonstrations with flight projects – both on the ground and in space. The lessons that we learn from these demonstrations can then inform the operational implementation of MUPA – an implementation that needs to occur by the middle of this decade if it going to help meet some of the significant increase in DSN antenna demand forecast to occur over the next 10 years.[14]

Acknowledgments

The authors wish to thank John Baker and Mazen Shihabi for their sponsorship of the Iris radio team work needed to incorporate the MUPA-related improvements into the Iris radio. The authors also thank Brian Giovannoni and Jay Wyatt for their sponsorship of much of the FCLTU Multiplexer development. In addition, the authors thank Stephen M. Lichten and Sami W. Asmar for their sponsorship of much of the other work discussed in this paper. The research described in the paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. A special thanks goes to NASA's Science Mission Directorate and NASA's Space Communications and Navigation (SCaN) Office for their support of this research.

---

[1] Abraham, Douglas S., "Future Mission Trends and Their Implications for the Deep Space Network," Space 2006 Conference, Session: SPS-1: Improvements, Innovations and Trends in Space Operations, San Jose, California, 19-21 September 2006, URL: https://doi.org/10.2514/6.2006-7247 [cited 13 November 2022].

[2] Abraham, Douglas S., Bruce E. MacNeal, David P. Heckman, Kristy Tran, Yijiang Chen, Andrew Kwok, Leslie Heid, and Hua Xie, "NASA Deep Space Communications: Future Mission Trends and Their Implications," SpaceOps 2023 Conference, ID# 137, Dubai, United Arab Emirates, 6-10 March 2023, [cited pre-publication, 14 November 2022].

[3] Abraham, Douglas S., Bruce E. MacNeal, and David P. Heckman, "Enabling Affordable Communications for the Burgeoning Deep Space Cubesat Fleet," SpaceOps 2016 Conference, Session: SSCSO - SmallSat Communications - Standards & Technologies, AIAA 2016-2625, Daejeon, Korea, May 2016, URL: https://doi.org/10.2514/6.2016-2625 [cited 14 November 2022].

[4] Morabito, David D. and Douglas S. Abraham, "Multiple Uplinks Per Antenna (MUPA) Signal Acquisition Schemes," SpaceOps 2018 Conference, Session: CAN - Communications Architectures and Spectrum Allocations, AIAA 2018-2610, Marseille, France, 28 May – 1 June 2018, URL: https://doi.org/10.2514/6.2018-2610 [cited 14 November 2022].

---

[5] Consultative Committee for Space Data Standards (CCSDS), *Recommnedation for Space Data System Standards: Space Link Extension – Forward CLTU Service Specification*, Recommended Standard 912.1-B-4, Blue Book, August 2016, URL: https://public.ccsds.org/Publications/BlueBooks.aspx [cited 27 November 2022].

[6] Morabito, David D. and Douglas S. Abraham, "Multiple Uplinks Per Antenna (MUPA) Signal Acquisition Schemes," [cited previously].

[7] *Ibid.*

[8] Morabito, David D., Zaid Towfic, and Douglas S. Abraham, "Comparison of Type 2 versus Type 3 Carrier Tracking Loops under High Dynamic Signal Conditions," SpaceOps 2021 Conference, Session 8, Cape Town, South Africa, 3-5 May 2021, URL: http://hdl.handle.net/2014/54759 [cited 30 November 2022].

[9] *Ibid.*

[10] Divsalar, Dariush, Marc Sanchez-Net, and Kar-Ming Cheung, "Adaptive Sweeping Carrier Acquisition and Tracking for Dynamic Links with High Uplink Doppler," 2020 IEEE Aerospace Conference, Big Sky, Montanna, 07-14 March 2020, DOI: 10.1109/AERO47225.2020.9172578, URL: https://ieeexplore.ieee.org/document/9172578 [cited 2 December 2022].

[11] Ortega, Tomas, Marc Sanchez-Net, Dariush Divsalar, and Kar-Ming Cheung, "Acquisition and tracking of high dynamics Doppler profiles for space applications," 2021 IEEE Aerospace Conference (50100), Big Sky, Montanna, 06-13 March 2021, DOI: 10.1109/AERO50100.2021.9438418, URL: https://ieeexplore.ieee.org/document/9438418 [cited 2 December 2022].

[12] Office of Chief Engineer, National Aeronautics and Space Administration, "Space System Protection Standard," NASA-STD-1006, 29 October 2019, URL: https://discovery.larc.nasa.gov/PDF_FILES/2019AO/nasa-std-1006.pdf [cited 23 December 2022].

[13] Consultative Committee for Space Data Standards (CCSDS), *Report Concerning Space Data System Standards: The Application of Security to CCSDS Protocols,"* Informational Report CCSDS 350.0-G-3, Green Book, March 2019, URL: https://public.ccsds.org/Pubs/350x0g3.pdf [cited 23 December 2022].

[14] Abraham, Douglas S., Bruce E. MacNeal, David P. Heckman, Kristy Tran, Yijiang Chen, Andrew Kwok, Leslie Heid, and Hua Xie, "NASA Deep Space Communications: Future Mission Trends and Their Implications," [cited previously].