17th International Conference on Space Operations, Dubai, United Arab Emirates, 6-10 March 2023.
Copyright 2023. All rights reserved. Published by the Mohammed Bin Rashid Centre (MBRSC) on behalf of SpaceOps, with permission and released to the MBRSC to publish in all forms.

SpaceOps-2023, ID # 583

A Demand Access System for Deep Space Operations

Marc Sanchez Net, Nathaniel Richards, Jay E Wyatt, T. Joseph W. Lazio, Rebecca Castano, Mark Johnston, Tim Pham, Nihal Dhamani, Gigi Lucena, Costin Radulescu
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109

Benjamin K. Malphrus, Jeffrey A. Kruth, Chloe Hart, Emily Mattle, Emily Walter
Space Science Center, Morehead State University, Morehead, KY 40351

Using small spacecraft for deep space exploration is a topic of growing interest. This is exemplified by the success of the Mars Cubesat One (MarCO) spacecraft, as well as by other missions such as the Lunar Trailblazers, the Escape and Plasma Acceleration and Dynamics Explorer (EscaPADE), and the thirteen Artemis 1 cubesats. While exploring deep space with smallsats offers clear scientific advantages, the prospect of a possibly (much) larger mission suite raises significant challenges for our current approach to ground support and mission operations.

This paper summarizes the efforts conducted at the Jet Propulsion Laboratory (JPL) to define and prototype a new way of conducting mission operations, which we term "demand access". In demand access, ground support is provided to missions in near real-time upon request from the spacecraft, including both the Deep Space Network (DSN) and the ground data system (GDS). This contrasts with the current way of performing operations, which entails long sequencing activities and contact opportunities with the ground that are scheduled weeks if not months in advance.

The demand access system is built around three core and fully integrated capabilities: (1) A queuing antenna that, following a schedule, periodically "polls" spacecraft operating in demand access mode whenever they point to Earth; (2) a suite of flexible scheduling tools that reserve time on DSN antennas ahead of each day of operations and then allocate time to spacecraft in near real-time based on requests received via the queuing antenna; and (3), optionally, a cloud-based GDS that is spins up resources for real-time telemetry and science processing in response to newly scheduled demand access tracks.

To demonstrate the feasibility of implementing and operating the demand access system, we show results from several prototyping efforts and tests conducted over the last two years in conjunction with the Morehead Lunar IceCube (MLIC) team at Morehead State University (MSU). These include, for instance, the use of MSU's 21-meter antenna (DSS-17) as a queuing antenna to support MLIC in demand access mode during their long cruise to the Moon; as well as tests conducted to demonstrate real-time delivery and processing of MLIC telemetry in the cloud using Delay Tolerant Networking (DTN) and the Advanced Multi-Mission Operations System (AMMOS).

1. Introduction

NASA's Deep Space Network (DSN) is currently supporting ~ 40 spacecraft with 14 antennas located at three complexes, Goldstone, CA, Madrid, Spain, and Canberra, Australia (**Figure 1**) [1]. Despite the fact that this is already the largest fleet of spacecraft ever supported by the DSN, space exploration trends suggest that the need for deep space communication and tracking services will only increase in the future. Three factors drive this need:

1. A growing interest in exploring the inner solar system using constellations and/or fleets of smallsats, as exemplified by mission concepts such as Janus [2], the Escape and Plasma Acceleration and Dynamics Explorer (EscaPADE) [3], the SunRISE CubeSats [4], or the already successful Mars Cubesat One (MarCO) spacecraft [5];

2. Renewed interest in exploring the Moon using both orbital and landed assets, as demonstrated by the Artemis-I mission, which included 10 CubeSats as secondary payloads, as well as the Commerical Lunar Payload Services (CLPS) program, which provides commercial transportation services to the lunar surface.
3. A new suite of deep space missions to explore the outer solar system, which includes a mix of flagship missions to outer planets (e.g., Uranus Orbiter & Probe, Europa Clipper, Enceladus Orbilander) together with cost-capped missions that are part of NASA's New Frontiers program (e.g., Dragonfly), among others.

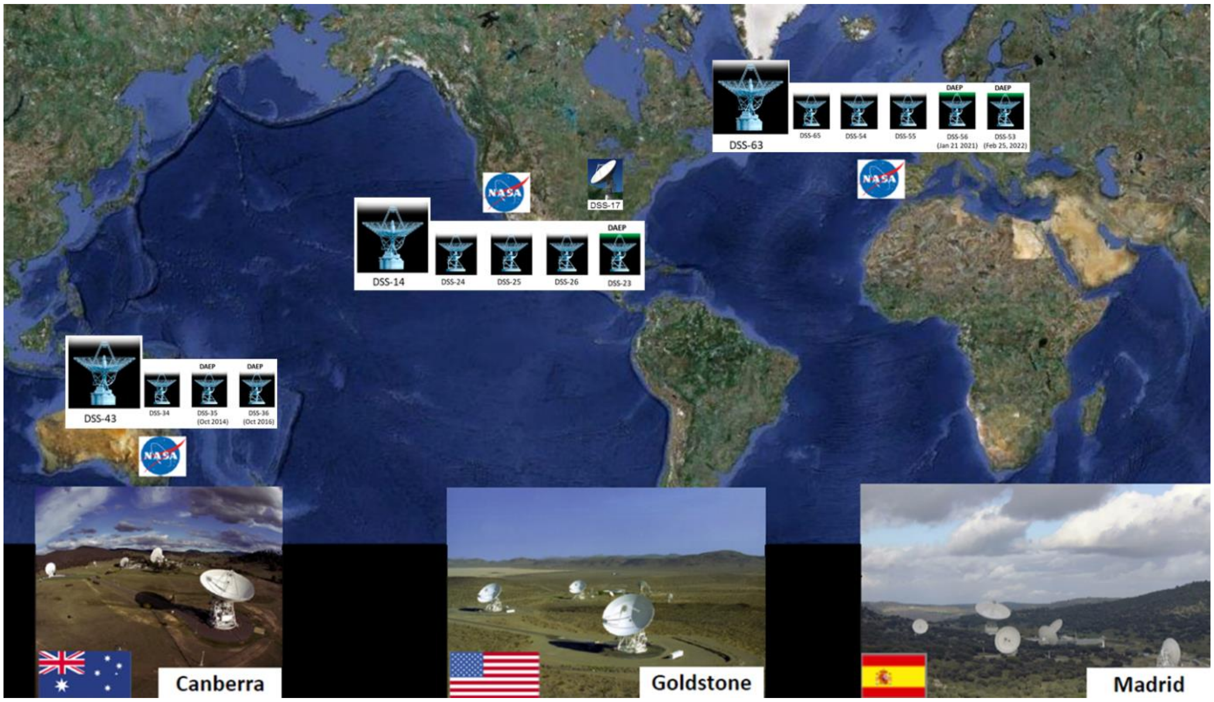

Figure 1. NASA's Deep Space Network (Image Credit: Stephen Lichten, JPL) [1]

To cope with the increase in demand of tracking time, the DSN is currently undertaking several technology developments to maximize the number of spacecraft that can be concurrently supported. In general, these efforts can be categorized in three areas: (a) Build new and/or upagrade existing antennas under the umbrella of the DSN Aperture Enhacement Project; (b) increase the number of spacecraft that can be supported simultaneously using techniques such as Multiple Spacecraft per Antenna (MSPA), Multiple Uplink per Antenna (MUPA), and Opportunistic Multiple Spacecraft per Antenna (OMSPA); (c) and develop new operational paradigms that more effectively utilize the available DSN assets. This paper focuses on this latter approach. In particular, we describe the efforts undertaken to date to define and prototype a new DSN demand access system (DAS) that is suitable for deep space operations.

The intent of a demand access system is to provide a way for spacecraft to request ground support whenever they need it, even if all of the specifics of that support were not fully planned and scheduled in advance. In essence, a DAS is akin to a paging system: Users (spacecraft) first notify the DSN of the need for tracking time and, in response to that event, DSN and ground data system (GDS) resources are allocated. This mode of operations is enabled by improvements in ground segment automation, as well as a lower risk posture for small missions.

The rest of this paper is organized as follows: In Section 2 we provide an overview the DAS architecture and describe in detail its concept of operations. We also summarize related operational concepts that are enabled by the technical capabilities required to build the DAS. Next, Section 3 describes the prototyping effort conducted to date to validate the different capabilities required for the DAS. Finally, Section 4 summarizes our conclusions and future work.

2. The DSN Demand Access Service

This architecture of the proposed DSN demand access system has been selected in order to satisfy four high-level requirements:

1. The system must be able to provide end-to-end ground support in response to possibly unanticipated and/or unplanned spacecraft requests. Here, end-to-end encompasses the DSN, the ground data system and, ultimately, data delivery to the end user, be it a mission operator or a scientist.
2. The system must offload tracking time from traditional DSN 34m and 70m antennas to a less capable but more affordable queuing antenna, especially for certain mission classes (e.g., smallsats) and/or during certain mission phases (e.g., long cruises to the outer planets).
3. The system must be able to service missions in deep space, for which long propagation delays render round-trip acknowledgements impractical. As a result, requests for ground support are placed by the spacecraft in a unidirectional manner, without explicit acknowledgement from the ground that they will be satisfied (or, for that matter, that they have even been received).
4. The system must enchance current operations, not disrupt them. Therefore, missions using the DAS must coexist with missions operating under the traditional paradigm of pre-scheduled contacts and normally sequenced activities.

A high-level comparison of how missions conduct day-to-day operations today vs. in demand access mode is provided in **Figure 2**. In normal scheduled operations, missions typically request DSN contact time with either 34m or 70m antennas at approximately regular intervals.[*] The cadence of contacts varies by mission and mission phase, but each contact opportunity is used for mission commanding and telemetry, as well as to obtain tracking observables. In contrast, in demand access mode missions obtain *periodic but sporadic* scheduled contacts with the DSN for the purpose of telemetry, tracking and commanding (TT&C). These scheduled contacts are interspersed with frequent contacts with the queuing antenna,which is used to *poll* the spacecraft and assess whether it needs additional ground support. Only if that is the case, extra TT&C tracks (which we call *demand access tracks* and are different from scheduled tracks represented in blue) with larger 34m and 70m antennas are allocated in near-real time from pre-reserved resources, a step not shown in **Figure 2** for the sake of clarity.

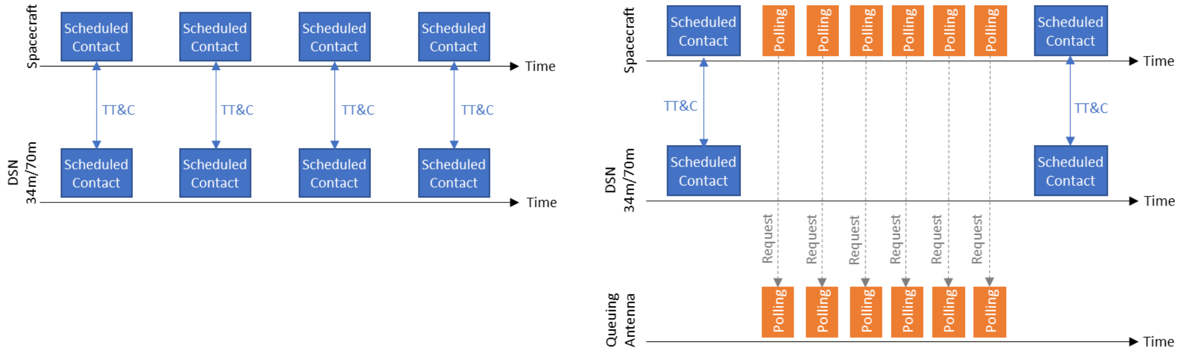

Figure 2.a. Operations with Scheduled Contacts

Figure 2.b. Operations in Demand Access

Figure 2. Operations in DAS vs. Scheduled Contacts

Central to this DAS is the notion of a "request", which we define as the act of transmitting a signal from the spacecraft to the queuing antenna for the purpose of requesting/securing tracking time on a DSN 34m or 70m antenna and, possibly, obtaining simultaneous computational resources on a ground data system. To transmit the request efficiently over a deep space link, it is encoded using a DSN beacon tone, a signal format that is inherited from the already operational DSN Beacon Tone Service, and that can already be detected using today's DSN receivers.[†] Also, the actual meaning of the request is mission dependent and can elicit different ground responses. For example, requests can be used to encode the urgency with which ground support is required according to the following levels:

---

[*] Actual mission operations are certainly more complicated than the basic view from **Figure 2**, especially for critical events. However, the example provided here is simplified for the sake of clarity and because elevated DSN support for critical events tends to happen a handful of times during a mission's lifetime.

[†] Because of that, the terms "request", "tone", "beacon", or "beacon tone", will be assumed synonym for the rest of the discussion.

Copyright 2023. All rights reserved. Published by the Mohammed Bin Rashid Centre (MBRSC) on behalf of SpaceOps, with permission and released to the MBRSC to publish in all forms.

- **Nominal request**: The spacecraft does not request ground support at this time. It is in nominal state and therefore there is no need for a demand access TT&C track.
- **Interesting request**: The spacecraft requests ground support whenever convenient/possible for the purpose of conveying an interesting on board event.
- **Important request**: The spacecraft requests ground support within a certain time window or its state could deteriorate or critical data could be lost.
- **Urgent request**: The spacecraft requests ground support immediately. This is equivalent to detecting a spacecraft anomaly and thus triggers emergency recovery procedures from the MOC and, if necessary, elevated levels of support by the DSN.
- **No request**: The tone service is not operating or another critical anomaly has occurred on board the spacecraft. The spacecraft concept of operations dictactes how the DSN and MOC respond to not receiving a request.

Alternatively, requests can be used to provide information on a limited set of vehicle states, which are then mapped by operational policy to a set of ground responses that can be executed in near-real time in a possibly automated manner. Note that, regardless of whether requests encode urgency or state information, they are decidedly different from traditional spacecraft-DSN interactions which aim at exchanging telemetry and commanding information and obtaining tracking observables.

Requests are received via an specially designated antenna, henceforth called queuing antenna, which *polls* spacecraft by pointing towards them for a certain period of time and attempting to receive a request. This queuing antenna may be implemented in several ways by:

- Reserving time on DSN 34m antennas for the purpose of performing queuing antenna operations. This is not the preferred alternative because it does not reduce operational burden on the DSN 34m antennas, but may prove useful as a way to deploy the DAS while dedicated queuing antennas are not available to provide support over the entire sky.
- Building dedicated queuing antennas at each of the DSN complexes. In that sense, previous analysis have shown that an 18 to 20m antenna with a cryo-cooled received is enough to provide service to the inner solar system and even parts of the outer solar system [6].
- Partnering with external institutions that have 18m or larger antennas available. This approach will be described in further detail in Section 3, when we describe the prototying efforts to date with Morehead State University (MSU).
- Reserving time on radio astronomy observatories. For example, the current ngVLA design calls for 244 18m receive-only antennas [7]. Because our proposed DAS does not have an acknowledgement, it does not require a transmit mechanism, thus making these receive-only antennas a good alternative.

The different elements that are required to implement the DSN demand access system are shown in **Figure 3**. Observe that the spacecraft communicates with the ground using two types of interactions, “requests” and normal TT&C.[‡] Incoming requests are first received by the queuing antenna and then disseminated via the DSN network operations center (NOC) to the DSN Service Scheduling Software (SSS) and the Mission Operations Center (MOC). SSS is then responsible for finding time on a larger 34m or 70m antenna, blocking it for exclusive use by the mission, notifying the MOC of the newly allocated track and, if necessary, communicating the schedule changes to the GDS so that computing resources are available to process TT&C data as it is received.

[‡] Although requests and TT&C are shown together in the **Figure 1**, they actually occur at different moments in time, as explained in Section 2.3.

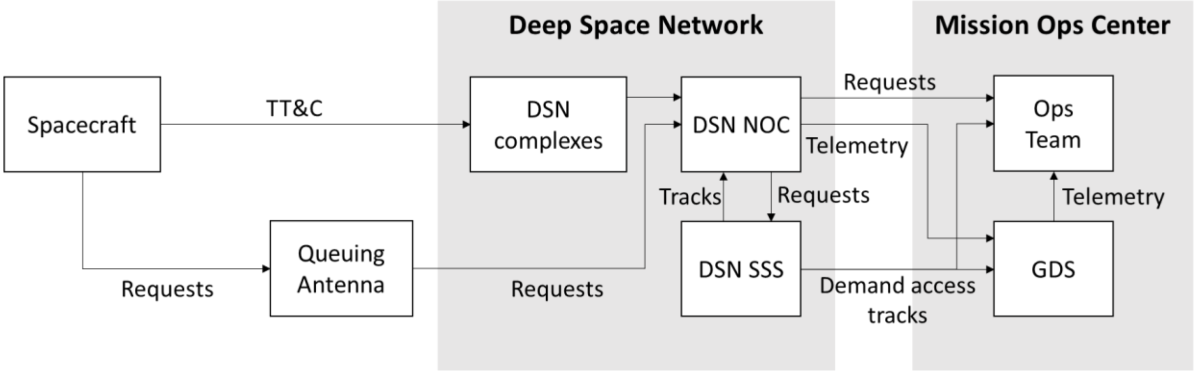

Figure 3. High-Level Components of a Demand Access System. TT&C stands for telemetry, tracking and commanding, DSN NOC stands for DSN Network Operations Center, DSN SSS stands for DSN service scheduling software, GDS stands for ground data system

2.1 Concept of Operations

Figure 4 provides a high-level depiction of the concept of operations (CONOPS) for the proposed DSN demand access service. It is divided into 9 panels with numbering on the top right corner of each text box, and depict how the sequence of events unfolds on three different time scales: DAS planning, spacecraft request, and demand access track execution.

Panels 1 through 4 happen during the mission planning phase, well ahead of the day of operations. Missions wanting to utilize the DAS notify the DSN during the the service negotiation phase. The number of missions using the DAS is then used by a new scheduling algorithm to reserve optimized blocks of time on the DSN schedule (green blocks in panel 2) [8]. These blocks mark certain antennas as "busy" for specific periods of time, but leave them unassigned and available for future use. In essence, these blocks become free resources that can be allocated in near-real time based on the requests arriving in any given day of operations. Next, missions using the DAS provide information about how often they would like to be polled. This information is used, together with the DSN schedule and a queuing antenna scheduling algorithm, to generate the schedule of the queuing antenna (panel 3). This schedule indicates, as a function of time, the sequence of spacecraft the queuing antenna will point to to look for incoming requests. It takes into account traditional factors such as line-of-sight visibility between spacecraft and the queuing antenna, as well as the position of the reserved blocks in the DSN schedule. Once the DSN schedule (with reserved blocks) and the queuing antenna schedule have been generated, they are both provided to projects for spacecraft sequencing purposes.[§] Mission planers use this information to generate sequences that indicate the periods of time during which the queuing antenna will be pointed towards the spacecraft, thus effectively providing time windows for the spacecraft to place its requests. These sequences are then uplinked to the spacecraft (panel 4) using a normally scheduled contact using the normal DSN commanding service (see blue blocks in **Figure 2.b**), where they are stored for future use in a manner akin to today's operation.

Panels 5 and 6 depict the events that unfold on any given day of operations when a spacecraft "decides" that for ground support is needed (the blue exclamation mark is used to graphically indicate that event). First, the spacecraft awaits until the start of the next polling contact with the queuing antenna. At that point, the request is sent via the downlink and is received by the queuing antenna. From there, it is forwarded to the DSN NOC and disseminated to both the MOC and the DSN scheduling system. As previously stated, the receipt of a request can elicit several responses depending on the operational policies of each mission; some of them can be automated, while some others require manual intervention from the MOC. For the purposes of Panel 5, however, we assume that receipt of a request is handled in an automated manner by the DSN SSS. In particular, a just-in-time allocation engine assigns a block of time reserved in the DSN schedule to the spacecraft placing the incoming request, this new allocation is notified to the MOC and, if

---

[§] According to the DSN Service Catalog [11], the DSN Planning and Scheduling Office publishes the mid-range schedule at least eight weeks in advance of tracking activities. This schedule is then used by projects for spacecraft sequencing.

needed, to the GDS. The timeliness of this just-in-time allocation process is variable, and based on our current prototypes we expect it to happen between 30 minutes to 2 hours prior to the start of each reserved block. As a way of comparison, normal pre-scheduled tracks are allocated several months ahead of time.

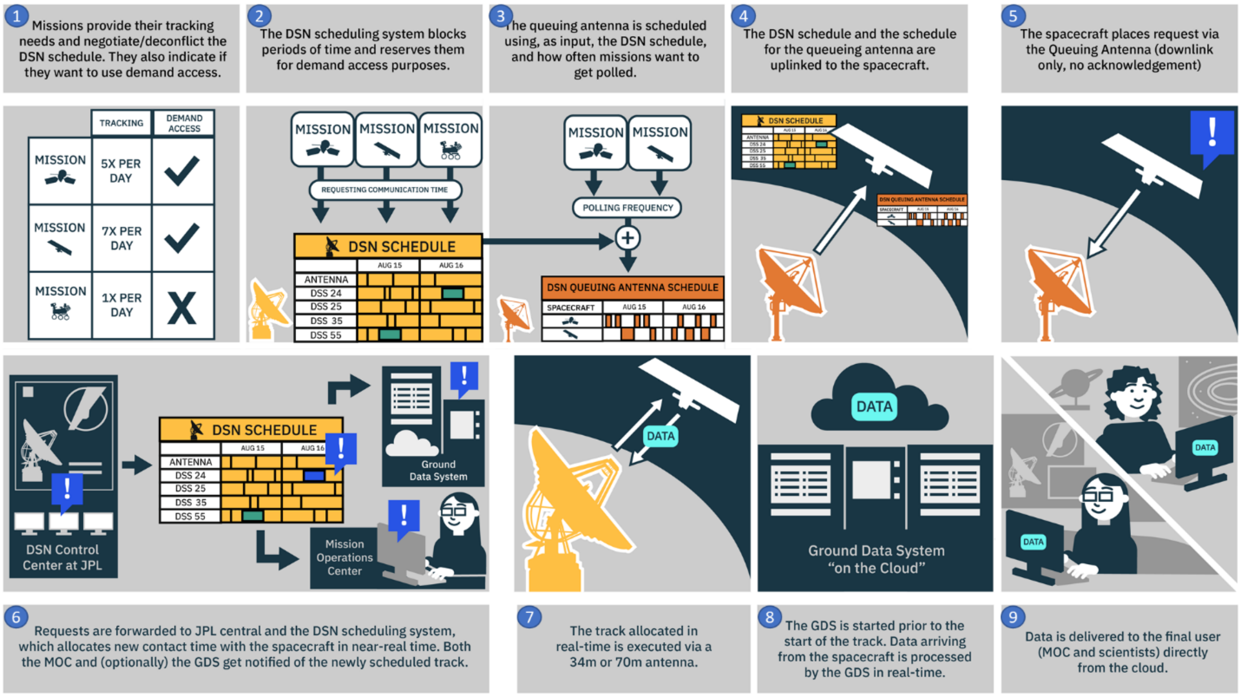

Figure 4. DAS Concept of operations. The orange antenna represents the queuing antenna, while the yellow antenna represents either a 34m or a 70m DSN antenna. The blue exclamation mark indicates and "event" on board the spacecraft that triggeres the rest of the Demand Access Service.

Once the spacecraft has sent a request to Earth, it receives no acknowledgement or confirmation that the request was received or that DSN and GDS resources will be available. Therefore, the spacecraft *assumes* that the reserved block immediately following its request is a valid opportunity to establish a TT&C link, as if the DSN had indeed allocated it for spacecraft use. When that reserved block starts, it attempts to contact Earth and send telemetry. If the link can be established, the spacecraft sends/receives data to/from Earth using normal DSN data services (see panels 7 through 9) in a manner no different than todays operation. This completes this interaction between the DSN and the spacecraft in demand access mode, after which the spacecraft resumes its "normal activities", sending requests via the queuing antenna whenever needed. If, on the other hand, the link cannot be established, likely because the block was assigned to another mission, then the spacecraft stops transmission, retains the data in its on board storage, and waits for the next reserved block, which it assumes has been assigned as a backup contact opportunity. When this second track starts, the spacecraft attempts once again to establish a link with the DSN and, upon success, data transfer is enacted using normal DSN data services. Finally, in the unlikely event that neither the primary nor the backup blocks are available to the spacecraft after placing a request, then the spaceraft resends its original request once again via the queuing antenna, and the process repeats until a demand access track is successfully allocated, or the next periodic but sporadic scheduled with the spacecraft starts.

2.2 Functionality of the DSN Demand Access Service

The functionality required to execute demand access operations can be grouped into four categories. Together, they for the core of the DSN Demand Access Service:

- **DAS Management and Scheduling Functionality**: It includes all functions related to planning, coordinating, and managing the DSN Demand Access Service, including the different scheduling mechanisms required for the system to operate.
- **DAS Testing Functionality**: It includes all functions required to perform DAS-related tests during assembly, test and launch operations (ATLO).

- **DAS Request Functionality**: It includes all functions related to triggering, transmitting, receiving and disseminating requests from spacecraft.
- **DAS Data Fuctionality**: It includes all functions related to executing demand access TT&C tracks with the spacecraft, processing the raw data, and managing data on board the vehicle.

Table 1 decomposes these four high-level functions into several second-level functions and maps them to different elements presented in **Figure 3**. In turn, these functions are used in the following section to describe the implementation of the DSN DAS, as currently conceived and prototyped.

2.3 Implemention of the Demand Access System

This section describes how the different functional capabilities described in **Table 1** can be implemented. In general, we focus here on describing the high-level implementation of the system, with further detailed provided in Section 3 and in previous publications enumerated throughout this section.

2.3.1 DAS Service Management

The DSN demand access service requires the mission and the DSN to exchange enough technical information to be able to establish RF links with the spacecraft. This information includes predicts of the spacecraft trajectory, as well configuration files that are used as inputs to the DSN downlink, uplink and ranging subsystems. This information is collected and processed using the DSN Service Preparation Subsystem (SPS), which output antenna tracking files, downlink and uplink Doppler-compensated frequency profiles, and other configuration files needed for the different DSN subsystems.

For the purposes of the DAS, no modifications to the DSN SPS will be needed. As explained later, requests are encoded in a format already supported by the DSN. Therefore, the current DSN receivers and their associated configuration files are already capable of capturing and managing the necessary information to receive spacecraft requests.

2.3.2 DAS Scheduling Services

The demand access systes requires three scheduling mechanisms to (1) reserve blocks of time on the DSN antennas, (2) schedule the queuing antenna, and (3) perform just-in-time allocation of reserved blocks to received spacecraft requests. Details on these algorithms have already been published in the literature – see [9], [10] and [6] – and are thus omitted here for brevity. That being said, we do provide a succinct discussion on two important topics related to the operation of the DAS, namely the number of blocks of time reserved for DAS, and how to deal with contention should it occur.

A key element to efficiently operate the DSN DAS deals with estimating the number of blocks of time that must be reserved for just-in-time assignment. Evidently, if too few blocks are reserved, several missions will compete for them leading to a starvation problem in which many of them will not get the tracking time they requested. Alternatively, if too many bloks are reserved, then it is very likely that some of them will be unusued because no requests will have been received to occupy them, thus resuling in "wasted" DSN time. From the fundamental laws of queuing theory, we know that for the DAS to be stable and avoid the starvation problem the average number of spacecraft requests per unit of time (e.g., a day or week) *must* be smaller than the average number of blocks reserved in that same unit of time.[**] However, determining the optimal number of reserved blocks in any given week is complicated, and requires careful analysis of the exact loading conditions at that point in time. Therefore, as part of our prototyping effort, we developed a simulation engine that can be used to perform rapid loading and what-if analysis given snapshots from the DSN operational schedule and a pool of DAS "customers". DSN schedulers could use it, ahead of time, to predict the "right" number of blocks to reserve in any given week, thus *proactively* minimizing the chances of wasting time on DSN 34m or 70m antennas.

[**] Technically, a nominal request indicates that no demand access track is needed. Therefore, stability requires that the average number of reserved blocks be greater than the average number of non-nominal requests.

Table 1. Functional Decomposition for the DAS

Category	Function	Description	Allocated To
DAS Management and Scheduling	DSN service management	Functions related to gathering, storing and processing required information to establish the RF link (e.g., trajectory predicts, space link profiles, etc.) and scheduling DSN tracks.	DSN
	DAS scheduling: Block reservation	Functions related to scheduling time on the DSN 34m and 70m antennas for the purpose of executing demand access tracks that are allocated in near real-time.	SSS
	DAS scheduling: Queuing antenna	Functions related to scheduling an antenna for the purpose of receiving requests from a spacecraft.	SSS
	DAS scheduling: Just-in-time block allocation	Functions related to allocating time of DSN 34m or 70m antenna in response to a request received via the queuing antenna.	SSS
	Ground support notification	Functions related to notifying the MOC and GDS when a new demand access track with a DSN 34m or 70m antenna has been allocated.	SSS, GDS
DAS Testing[††]	DAS compatibility testing	Functions related to performing compatibility tests to ensure that requests generated by the spacecraft can be received and disseminated.	DSN
	GDS compatibility testing	(Optional) Functions related to testing the on-demand GDS provisioning system in response to DAS events.	GDS
DAS Request	Request triggering	Functions executed on board the spacecraft to identify the need for placing a request.	Spacecraft (FSW)
	Request signalling	Functions related to turning a request into an RF signal for the purpose of sending it via a space link.	Spacecraft (radio)
	Request reception	Functions related to receiving the RF signal and decoding the spacecraft request from it.	Queuing Antenna
	Request dissemination	Functions related to disseminating the request received by the queuing station to the MOC and the DSN SSS.	DSN NOC
DAS Data	Execution of DSN data services	Functions related to establishing a link between the spacecraft and a DSN 34m or 70m antenna for the purpose of sending commands, telemetry and, optionally, obtaining tracking observables.	Spacecraft, DSN complexes
	Execution of GDS data processing	Functions related to processing the commanding and telemetry stream exchanged during a demand access track.	MOC, GDS
	On-board data management	Functions related to persistently storing the data on-board the spacecraft as it awaits for DSN time to be allocated.	Spacecraft (FSW)

[††] In addition to DAS testing, projects also need to perform testing of their traditional TT&C link. This is not listed explicitly in **Table 1** since such tests are no different than what would be required from a mission requesting traditional DSN data services.

Avoiding wasting DSN antenna time in the DAS can also be done *reactively*, once we know the number of requests that have been received and prior to the just-in-time allocation process. In that sense, consider a case where N spacecraft are operating in demand access mode and the queuing antenna receives $0 \leq M \leq N$ requests. Then, three cases are possible:

- No requests were received (M=0): The reserved block of time must be repurposed to avoid "wasting" time on a DSN 34m antenna. Several backup alternatives can be used including (1) performing maintenance activities, (2) using the block for radio science, radio astronomy or radar science, or (3) tracking a spacecraft that is inview of the reserved antenna and that is known to transmit information continuously (e.g., Voyager spacecraft).
- Exactly one request was received (M=1): The just-in-time allocation simply allocates the reserved block to the spacecraft that has placed the request. This is the nominal case and wastes no DSN time.
- More than one request was received (M>1): In that case, no DSN time is wasted, but the just-in-time engine must deal with a contention problem. Several options are possible to select which spacecraft obtains DSN resources including selection based on fairness considerations (e.g., round-robin mechanism), based on levels of urgency and/or priority, or splitting the reserved block in time among the different missions. Note that the latter option has limits on the number of splits that can occur in any given block due to pre-calibration and post-calibration times needed before/after each split.

2.3.3 DAS Testing

In this subsection we describe additional testing that must be performed prior to launching a spacecraft in order to operate in demand access mode. To facilitate the discussion, **Figure 5** provides depicts the architecture of the DAS with interfaces between the project and the DSN color-coded depending on whether they are new (red), inherited from the current DSN data services and GDS (green), or internal to the DSN (blue). Green interfaces must be tested whether a mission operates in demand access mode or not. Therefore, activities related to testing them are not described here since they are "business-as-usual" from the point of view of projects. Similarly, internal DSN interfaces must be tested and validated by the DSN prior to operationally deploying the DSN DAS, but need not be retested every time a new mission wants to use it.

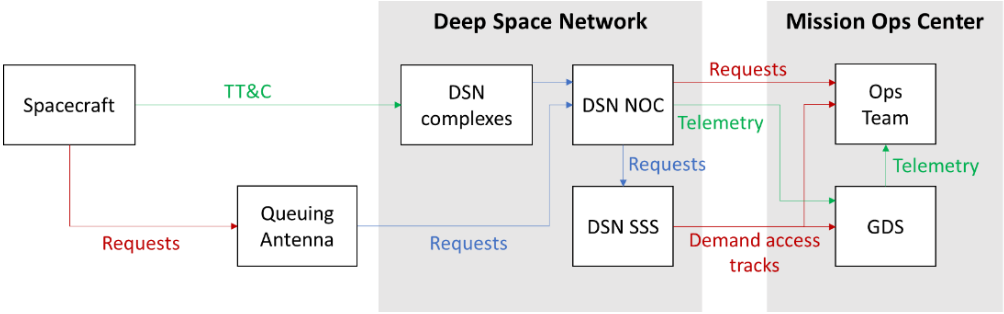

Figure 5. DAS Interface Categorization

From the point of view of missions, three additional interfaces must be tested using the DSN DAS: The RF link between the spacecraft and queuing antenna, with the waveform selected to encode requests (see **Section 2.3.5**); the interface between the DSN NOC and the MOC which, as explained in **Section 2.3.7**, is built upon currently available capabilities but requires updates to DSN and MOC configuration files to operate correctly; and, finally, the new interface allowing the SSS and GDS to interface directly for the purposes of being notified when demand access tracks are allocated. Note that this interface will be built around the technologies already used by SSS such as standard REST APIs, or other high-level messaging services. Thus it is not expected to be technically difficult to instantiate and will likely not require significant testing procedures prior to becoming operational.

Out of these three interfaces, probably the one that requires the most attention is the RF link between the spacecraft and the queuing antenna. Luckily, because the selected waveform is already used by the DSN beacon service, the test equipment already available at the DSN can be used to validate the request mechanism of the demand access service. In particular, DSN's Development and Test Facility (DTF-21) currently has four receivers capable of tone detection.

Therefore, we expect that missions would conduct their DAS-related tests while performing all other tests at DTF-21, nominally 18 months prior to mission launch [11].

2.3.4 DAS Request Triggering Mechanisms

We assume that either the spacecraft flight software (FSW) or the spacecraft fault management system will be responsible for triggering a request. The specific trigger points are unique to each mission, its concept of operations, and its hardware and software. However, in general several categories can be identified:

a) Requests sent in response to an engineering event/fault. Examples of such events are varied and include notifying the ground that (1) the data recorder is full, (2) the momentum wheels are saturated, (3) an instrument is in an off-nominal state (e.g., a cryocooled instrument has not reached its operating temperature), (3) a thrust activity was shorter than expected, etc.

b) Requests sent in response to a scientific event of interest, which assumes FSW has a science data summarization mechanism that allows such events to be identified on board the spacecraft (e.g., the Magnetospheric Multiscale mission has demonstrated this capability).

c) Requests sent in response to an event that must communicated to another asset for a coordinated observation, a mode of operations that is currently used for detection and observation of Gamma Ray Bursts, but that also applies to other scientific areas of interest such as detection of characterization of hydrogen on the lunar surface.

The level of effort required to implement the request triggering mechansims depends on the type of trigger. For category (a), our conversations with flight software engineers to date indicate that low engineering effort is needed for missions using state-of-the art flight software libraries. For example, spacecraft using NASA's Core Flight Software (cFS) can take advantage of its native Limit Checker application [12] to trigger a request whenever a certain telemetry point exceeds a thresdhold. When that happens, cFS' Limit Checker creates an event record in the telemetry buffer and, optionally, initiates a Relative Time Sequence (RTS) that can include the necessary commands to configure the radio for transmission of a ground request. Similar capabilities are available in other flight software libraries such as JPL's F-Prime, which includes a "Monitor" component that can be used in manner analogous to cFS' Limit Checker [13].

On the other hand, requests generated based on categories (b) and (c) are tailored to the particular science activities and CONOPs of a given mission. That being said, emerging autonomy frameworks such as JPL's SYNOPSIS are good candidate technologies that can be used for science data analysis, summarization and event triggering. They are particularly suitable if built on top of the aforementioned flight software libraries (e.g., SYNOPSIS is compatible with both cFS and F-Prime), hence being able to leverage common tools and command sequences to trigger the request.

2.3.5 DAS Request Format

Normal TT&C communications between a spacecraft and a ground station involve phase modulating bits onto an carrier. However, for the purposes of signalling a request, we assume that the spacecraft will use a simpler waveform, namely a squarewave subcarrier[‡‡] modulated onto a suppressed carrier [6].[§§] With this waveform, the request transmitted is encoded directly in the frequency of the subcarrier and thus no bits need to be transmitted. This choice is motivated by two factors:

- First, detection of square-wave subcarriers can be performed using algorithms based on power-spectral density estimation. Therefore, large amounts of coherent and incoherent integration time can be performed (up to ~1000s of seconds), which allows the queuing antenna to provide service deep into the solar system even if its diameter is smaller than 34 meters [6].
- Second, providing a limited set of subcarriers helps prevent the problem of data tantalization. In other words, mission designers and operators know that, by design of the DAS, requests from spacecraft will only provide insight into a limited set of states for the sole purpose of determining the level of ground support required by

---

[‡‡] A sine-wave subcarrier could also be used. However, it is common at JPL to use squarewave subcarriers in the downlink direction, hence our assumption.

[§§] This signaling format is currently used in the DSN Beacon Tone Service. Therefore, it is sometimes referred to as a "tone", "beacon", or "beacon tone", which we treat as synonyms to DAS requests for the rest of the discussion.

the DSN. Further insights into the full state of the spacecraft will come later, when the actual demand access TT&C track is executed.

In theory, the number of subcarrier frequencies that can be used to transmit different types of requests is infinite. In reality, however, several factors reduce this "alphabet" including limitations in the reception algorithm, Doppler effects, errors in the spacecraft trajectory predicts, etc. In that sense, based on the capabilities of the current DSN receivers, 4 tones per carrier are currently supported, with projects being able to configure tones on multiple carriers to extend the number of requests if needed.[***] Additionally, the availability of DSN tones in radios qualified for deep space operations has also been assessed while prototyping the DSN DAS. In particular, both JPL's IRIS and UST radios, as well as APL's Frontier radio are capable of producing DSN tones. In fact, the latter used it operationally while in cruise to Pluto, while availability in the IRIS radio was tested during this prototyping effort.

2.3.6 DAS Request Reception

Several algorithms have been developed in the literature to receive a DSN tone in the presence of Additive White Gaussian Noise (AWGN). For example, detailed analysis for different optimal criteria using both the log-likelihood ratio and maximum-likelihood ratio tests is provided in [14] and [15]. Nevertheless, due to their computational complexity, the DSN receiver uses a more conventional non-coherent algorithm based on estimating the received signal's power spectral density (PSD) – see [16] for details. This algorithm first uses the Fast Fourier Transform (FFT) to estimate the signal spectra, and then performs non-coherent averaging of FFT outputs over consecutive periods of time. Furthermore, Doppler post-compensation using spacecraft trajectory predicts is used to ensure all tone energy falls within one bin of the discrete PSD estimate.

The theoretical performance of this non-coherent tone reception algorithm is compared to that of a traditional BPSK downlink in Reference [17]. It is shown that with sufficient integration time, a tone may be detected up to a tone power to noise spectral density (P_t/N_0) ranging from -5 to 0 dB-Hz, a ~15 dB improvement over what would be achievable if we were to use coherent modulation schemes to encode the spacecraft request. In reality, however, the DSN service catalog currently advertises a minimum P_t/N_0 of 5 dB-Hz for reliable tone detection as part of the DSN Beacon Tone Service.[†††] Since the proposed DSN DAS inherits the tone reception capability directly from the DSN Beacon Tone Service, this value provides a good estimate for the minimum signal to noise ration that must be received for the DSN to successfully receive a spacecraft request.

2.3.7 DAS Request Dissemination

Once the request has been received by the queuing antenna, it must be disseminated to the DSN NOC, the DSN SSS and the MOC. This is performed using standard interfaces currently available at the DSN. In particular, the DSN receiver generates a standard CCSDS Standard Formatted Data Unit (SFDU) that contains a tone value number that maps to a subcarrier frequency based on configuration files generated by the DSN during the service negotiation phase. Additionally, the SFDU also includes metadata relating to the tone detection such as the system noise temperature, the power to noise ratio, the configured integration time, or the center frequency, among others [18].

The tone SFDU generated by the DSN receiver is forwarded to the DSN NOC at JPL using the ground lines connecting JPL with the different DSN sites. From there, it is made available to the MOC via a traditional CCSDS Space Link Extension (SLE) interface [19]. Several standard SLE services can be used deliver tone SFDUs from the MOC, including the Return All Frames (RAF) and Return Channel Services (RCF). These are configured ahead of time by DSN personel at JPL, during the service negotiation phase, in a manner similar to all other DSN data services.

Finally, the interface between the DSN NOC and the DSN SSS is internal to the DSN and has currently not been implemented yet. However, the SSS block allocation algorithm has been prototyped to have standard REST-based interface for the purpose of being notified whenever a request arrives.

2.3.8 Execution of DSN Data Services

The execution of a demand access track that has been allocated just-in-time based on requests arriving from spacecraft is, in most ways, no different than any other track currently supported by the DSN. Once the track is allocated in SSS,

---

[***] JPL has demonstrated the use of a much larger alphabet of tones during critical events such as Mars Entry Descent and Landing [24]. The receivers used for these critical events are current not baselined for the DAS implementation, but they could be repurposed within some engineering effort should a mission require it.

[†††] This assumes a non-coherent integration time of 1000 seconds, i.e., the queuing antenna would need to track a spacecraft for more than 15 minutes to receive a request.

SPS will automatically generate all metadata and configuration files necessary for executing the track. At the appropriate time, these will be retrieved by the DSN upink, downlink and tracking subsystems to establish the link with the spacecraft, send commands, receive telemetry, and, if necessary, collect tracking observables.

There are two main constraints in the currently prototyped DAS for executing the DSN data services. First, blocks of tracks can only be assigned to one spacecraft. Indeed, additional machinery would need to built into the system to allow for split tracks or other ways of dealing with contention events. Second, demand tracks involving MUPA, MSPA, OMSPA sessions have been left out of scope since they require coordination between multiple projects.

2.3.9 Execution of GDS Data Processing

Once a demand access track is scheduled by the DAS just-in-time allocation engine, computing capabilities to process the incoming telemetry stream are needed. In general, these can be provisioned in two manners:

- **Dedicated GDS**: Computational resources are deployed during the mission's development phase and are available exclusively and continuously during the mission's lifecycle. The mission is responsible for provisioning, maintaining and paying for these resources, which are connected to the DSN NOC prior to launch.
- **On-demand GDS**: Computational resources are provisioned on-demand whenever the DAS assignes a track to the spacecraft. These resources are nominally allocated on cloud-based infrastructure, shared by many missions, using already available software libraries and tools such as the Advanced Multi-Mission Operations System (AMMOS) SmallSat Toolkit on Amazon Web Services (AWS).

For on-demand GDS operations, a GDS provisioning system is needed (see **Figure 6**). It acts as coordinating element, interfacing the DSN SSS with the cloud-based GDS solution, and spinning up/down computational resources on the cloud as needed. For example, in Section 3 we show how such a provisioning system can be built using the REST interfaces already available in SSS, together with scripting capabilities provided by cloud providers such as AWS.

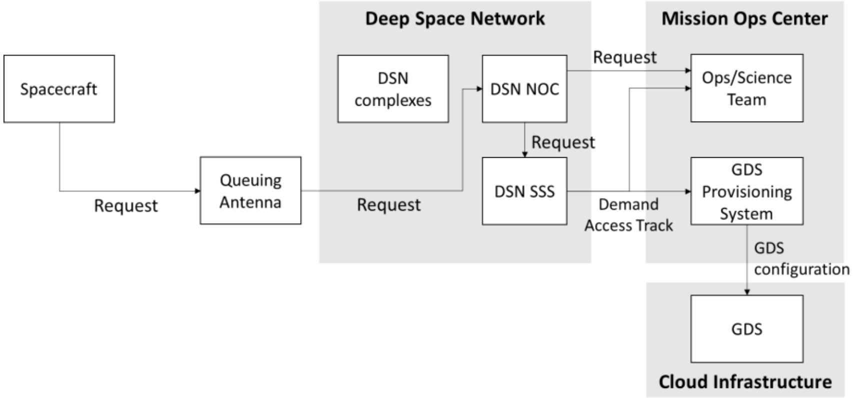

Figure 6. On-demand GDS Architecture: Request Processing

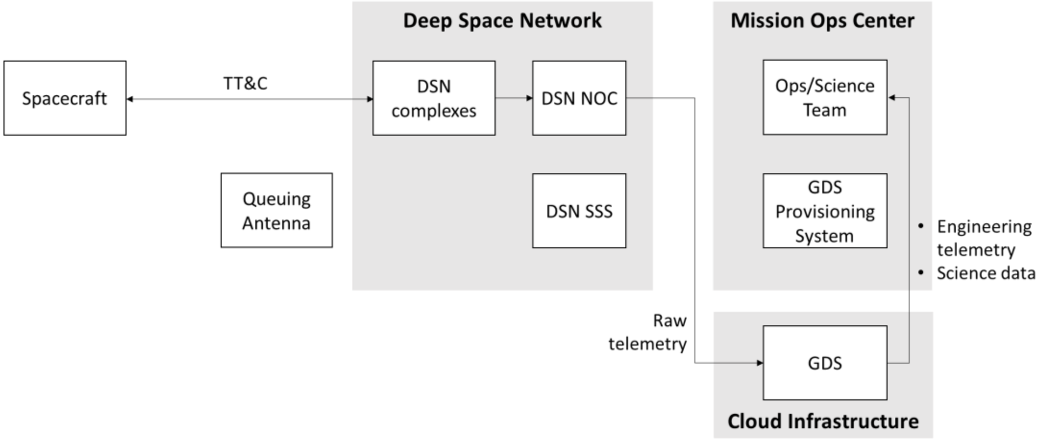

Figure 7. On-demand GDS Architecture: Telemetry Processing

2.3.10 On-Board Data Management

Management of data on-board the spacecraft requires functionality to ensure no data is lost in the event that a demand access track with a 34m or 70m cannot be allocated because some other mission also requested DSN time. There are several ways to implement this functionality (e.g., do not transmit any information unless a two-way link between the spacecraft and Earth has been established), one of which is leveraging Delay Tolerant Networking (DTN) as standardized way of (1) handling persistent storage of information when dealing with uncertain deep space links,[‡‡‡] and (2) managing transmission times when operating in demand access mode. Note, however, that DTN is in no way required for a mission to operate in demand access mode using the proposed DSN DAS. Note also that further research should be conducted to properly assess the suitability of DTN for on board operations with the DSN DAS.

In DTN, data is stored on board a spacecraft in the form of bundles, i.e., self-contained units of information with source and destination addresses akin to IP packets that are defined in the Bundle Protocol (BP) [20]. Bundles are managed by a bundle agent, typically an application residing in flight software, which “tags” them with retention constraints to indicate how many steps in the processing chain have been completed. Importantly, a bundle can only be deleted from on board storage once all retention constraints have been removed, an event that only happens when the bundle agent has been informed of successful bundle transmission by the selected converge layer (i.e., the radio).[§§§] If, on the other hand, the convergence layer informs the bundle agent that transmission has failed (because the DSN did not allocate a demand access track as requested by the spacecraft), then a DTN implementation tailored for demand access will choose to retransmit the bundle at a later time, when a new communication opportunity arises (i.e., in the next DSN reserved block).

DTN also provides facilities for automating data transmission over multiple consecutive transmit windows via a contact plan. In particular, a contact plan is a list of contact opportunities between two DTN nodes (the spacecraft and Earth, for the purposes of this discussion), where each contact opportunity has associated start and end times, as well as a nominal transmission data rate. For the purposes of the demand access operations, this contact plan would be built on Earth by the mission operators and contains (1) scheduled DSN contacts, as depicted in blue in **Figure 2.b**, and (2) unassigned but reserved blocks of time, as shown in panel 2 of **Figure 3**. Together, these contact opportunities provide the BP agent with a full list of all the contacts that *may* be used for transmitting/receiving TT&C data. Note that some of these contacts are fully certain, while others are conditional on being granted in near real-time by the DSN SSS system.

Finally, as previously mentioned, the spacecraft also need a contact plan for the queuing antenna so that request can be sent to the ground. However, this would likely be managed independently from the rest of the DTN system and handled by normal sequencing activities.

2.4 Related Concepts

This section briefly summarizes DSN services that can be built using subsets of the technical capabilities needed for a demand access service. Rather than detailed expositions of how they work, we focus here on technical and operational commonalities with the DSN DAS as currently described, and highlight some differences to avoid confusion.

2.4.1 The DSN Beacon Tone Service

The DSN Beacon Tone Service is an already operational DSN service that provides a way for the MOC to monitor the high-level state of a spacecraft according to a set of tones generated by the spacecraft [11]. Tracks for the DSN Beacon Tone Service are scheduled via the DSN SSS like any other data service, and must be negotiated and allocated in the DSN schedule weeks if not months in advance. Operationally, they are executed using traditional DSN 34m and 70m antennas, albeit the required pre-calibration and post-calibration times are reduced.

There are two key differences between the DSN Beacon Tone Service as currently offered by the DSN and the proposed DSN DAS. First, the Beacon Tone Service relies on the large 34m and 70m antenna for its operation, thus not offloading these type of tracks to smaller antennas for the purpose of saving DSN antenna time. Second, the Beacon Tone Service is pre-scheduled and there is no accommodation for tones to be received in response to a spacecraft event. That being said, the request signalling in the proposed DSN DAS has been inherited from the DSN Beacon

---

[‡‡‡] In demand access mode, the spacecraft does not know with full certainty if DSN support is available after placing a request. Therefore, from its point of view demand access tracks are uncertain.

[§§§] Note that this notification can come from the converge layer protocol, or from a command send from the ground indicating successful reception of the data.

Tone Service, a choice that facilitates the development of the demand access service by reusing technical capabilities already available in the DSN such as reception algorithms and testing equipment.

2.4.2 Red/Green Tone Service

The Red/Green Tone Service also provides a way for the MOC to monitor the high-level state of a spacecraft according to a set of tones generated by the spacecraft, but does so using the queuing antenna rather than DSN 34m and 70m antennas. It is currently not operational. In essence, it can be considered a basic form of the DSN DAS in which only the request mechanism is implemented. Tones received by the queuing antenna are forwarded directly to the MOC, which responds to them using already scheduled tracks with the spacecraft. In other words, there is no need for SSS to reserve time on the DSN 34m and 70m antennas, nor there is a need for a just-in-time block allocation that matches these reserved time to incoming requests. Similarly, integration between the Red/Green tone service and the GDS is not needed since new demand acess tracks are never scheduled.

2.4.3 Ground-Triggered Demand Access Service

The Ground-Triggered Demand Access Service allocates time on DSN 34m or 70m antennas in near real-time, with a timeliness comparable to the proposed DSN DAS. In its most basic instantiation, the allocation engine selects between two users: A spacecraft that is known to be transmitting at that time, and a science target that is selected based on prior knowledge provided by external entities. As an example, consider the case of the Goldstone Solar System Radar[****] (GSSR) and the Voyager spacecraft. GSSR operators may want to schedule tracks that are contingent on new asteroids being discovered by other observatories (among other factors), a concept of operations only possible with a ground-triggered DAS. Indeed, if a new asteroid is available for observation, then a track can be allocated in near real-time to the GSSR. Alternatively, the track can be turned into a traditional telemetry pass with Voyager, and additional data that would otherwise not be collected is obtained.

Observe that from the point of view of testing the DSN DAS, the Red/Green Tone Service and the Ground-Triggered Demand Access Service are complimentary and allow us to mature different technical capabilities. In particular, while the Red/Green Tone Services allows us to prototype and test the request mechanism and queuing antenna, the Ground-Triggered DAS allows us to test and validate the allocation engine that matches reserved blocks of time to customers in real-time (despite the fact that the criteria for allocation may not be the same).

3. Prototyping of the DSN Demand Access System

This section describe the efforts undertaken during Fiscal Year 2021 and 2022 to prototype different parts of the DSN DAS. We group these tests in three categories, depending on whether they are related to the request mechanism via the queuing antenna, the development of the scheduling algorithms required for executing demand access tracks, or tasks related to prototyping an on-demand GDS that is responsive to tracks allocated in near real-time.

3.1 Testing of the Request Mechanism and the Queuing Antenna

Tests related to the request mechanism and the queuing antenna were performed at MSU, using their 21m antenna (see **Figure 8**). MSU and JPL have collaborated over the last years to turn this antenna into a DSN associated station, now known as DSS-17. It is equipped with DSN-rated transmit and receive electronics, as well as a frequency and timing distribution system based on a Microwave Amplification by Stimulated Emission of Radiation (MASER). The antenna itself, as well as the RF frontends (including the low noise amplifiers), are different from the other DSN antennas, and the system's monitor and control is also less automated.

[****] The GSSR uses the 70m antenna at the Goldstone Complex, in CA, together with a specially designed transmit subsystem, to perform planetary radar observations.

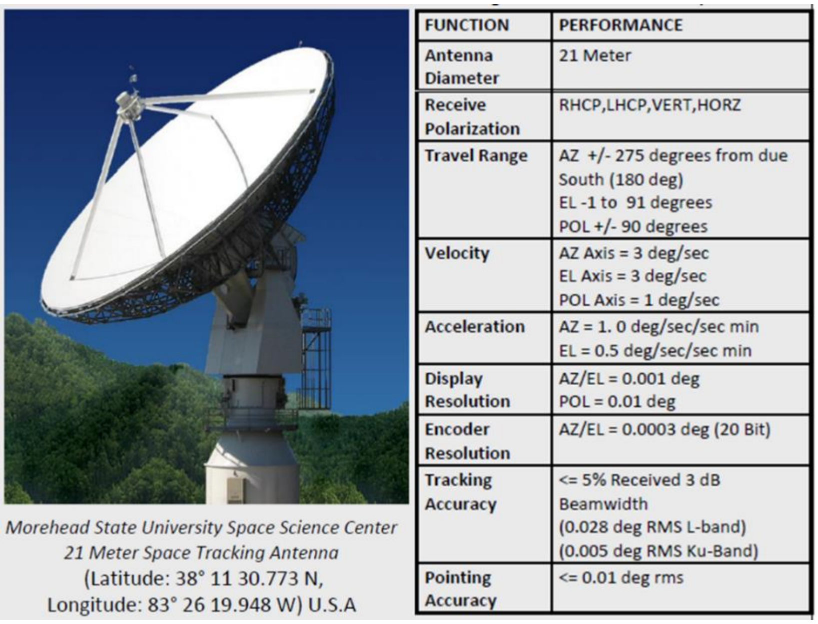

Morehead State University Space Science Center
21 Meter Space Tracking Antenna
(Latitude: 38° 11 30.773 N,
Longitude: 83° 26 19.948 W) U.S.A

FUNCTION	PERFORMANCE
Antenna Diameter	21 Meter
Receive Polarization	RHCP,LHCP,VERT,HORZ
Travel Range	AZ +/- 275 degrees from due South (180 deg) EL -1 to 91 degrees POL +/- 90 degrees
Velocity	AZ Axis = 3 deg/sec EL Axis = 3 deg/sec POL Axis = 1 deg/sec
Acceleration	AZ = 1. 0 deg/sec/sec min EL = 0.5 deg/sec/sec min
Display Resolution	AZ/EL = 0.001 deg POL = 0.01 deg
Encoder Resolution	AZ/EL = 0.0003 deg (20 Bit)
Tracking Accuracy	<= 5% Received 3 dB Beamwidth (0.028 deg RMS L-band) (0.005 deg RMS Ku-Band)
Pointing Accuracy	<= 0.01 deg rms

Figure 8. DSS-17 at MSU (adapted from [9])

Using DSS-17 as a queuing antenna has several advantages. First, the antenna is now scheduled via the DSN SSS and ingests products generated by the DSN SPS. Second, although not collocated with the DSN Goldstone Complex, DSS-17 covers a similar fraction of the sky, with some additional overlap with the DSN complex in Madrid, Spain. And third, although the antenna has started providing operational support to certain lunar spacecraft (e.g., CAPSTONE), the amount of time on the antenna that can be devoted to testing is significantly larger, thus alleviating time pressure concerns while developing the DSN DAS.

Three types of tests were conducted at MSU with DSS-17:

- **Demonstration tests**: Their goal is to provide the first empirical evidence that the equipment at DSS-17 can successfully perform a necessary functionality to act as a queuing antenna.
- **Performance tests**: Their goals is to quantify figures of merit that are key to understanding the level of service provided by DSS-17 if used as a queuing antenna.
- **Operational tests**: Their goal is to exercise the operational procedures required to use DSS-17 as a queuing antenna, and quantify to resulting operational load on the MSU team.

Next, we provide a succinct description of the tests performed in each category.

3.1.1 Demonstration Tests: Tone Generation, Reception and Dissemination

To demonstrate the ability of DSS-17 to operate as a queuing antenna, the MSU team developed a beacon emulator, i.e., an RF frontend coupled with an Arduino microcontroller that can be used to generate a square-wave subcarrier modulated onto an S- or X-band subcarrier. The RF signal generated by the beacon emulator was then connected to a horn antenna, which was directly pointed at DSS-17 as shown in **Figure 9**. This provided an over-the-air test setup that could be used to validate the ability of DSS-17 to detect DSN tones, as well as perform sensitivity and operational tests.

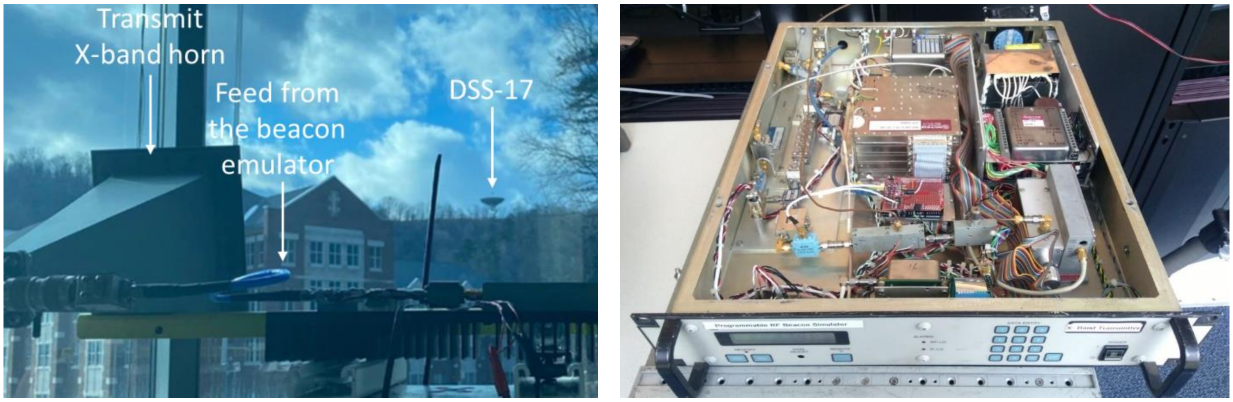

Figure 9. Beacon Emulator Connected to a Horn Antenna

The initial DSS-17 tests were conducted using the DSN receiver graphical user interface (GUI), available directly from the receiver console. For example, **Figure 10** shows the results of two tests in which the beacon emulator was configured to send a subcarrier at 3.8 kHz and 7.5 kHz respectively. Observe that the DSN receiver correctly detects the transmitted tone and reports it in tone reception GUI.

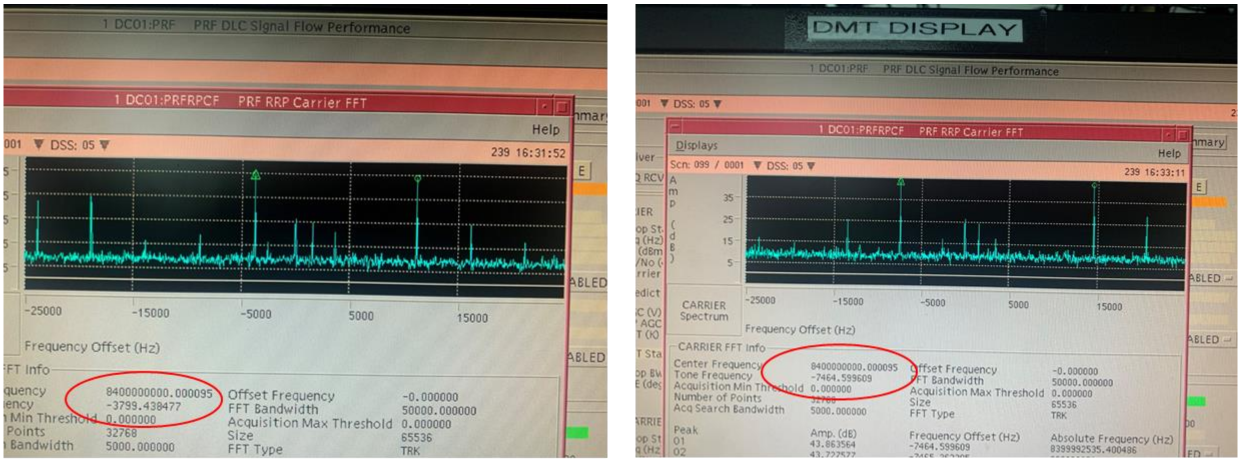

FFT Info
quency 8400000000.000095
iency -3799.438477
Min Threshold 0.000000
Points 32768
Bandwidth 5000.000000
Offset Frequency -0.000000
FFT Bandwidth 50000.000000
Acquisition Max Threshold 0.000000
Size 65536
FFT Type TRK

CARRIER FFT Info
Center Frequency 8400000000.000095
Tone Frequency -7464.599609
Acquisition Min Threshold 0.000000
Number of Points 32768
Acq Search Bandwidth 5000.000000
Offset Frequency -0.000000
FFT Bandwidth 50000.000000
Acquisition Max Threshold 0.000000
Size 65536
FFT Type TRK

Peak	Amp. (dB)	Frequency Offset (Hz)	Absolute Frequency (Hz)
01	43.863564	-7464.599609	8399992535.400486
02	43.722577		

Figure 10. Detection of a DSN Tone at DSS-17

In a similar manner, we also tested the ability of DSS-17 to receive tones generated by an actual space flight radio. To do so, we provided MSU with an JPL IRIS radio and they integrated it with a FlatSat.[††††] Then, the RF output of the IRIS radio was connected via a cable to the input of DSS-17's receive subsystem, and the IRIS radio was commanded to transmit a tone using via software from the FlatSat. Results of this test, not shown here, were similar to those of **Figure 10**, with successful detection of the tone frequency.

Finally, the ability of DSS-17 to disseminate received tones to the DSN NOC at JPL was also verified while performing tests with the IRIS radio. In particular, a total of 61,000 tone detection messages were generated by DSS-17 overnight and forwarded to the JPL NOC, where reception was later confirmed. Dissemination of tone reception messages from the DSN NOC to the LIC MOC was also successfully verified, in a later test, while forwarding to SSS was not demonstrated.

3.1.2 Performance Tests: DSS-17 Sensitivity vs. Polling Track Duration

Most of the test described in Section 3.1.1 were conducted under benign channel conditions, with plenty of marging on the received signal-to-noise ratio. To validate the breaking point of DSS-17, we conducted a sensitivity test using

---

[††††] MSU designed and built the Lunar IceCube (LIC), launched on board the ARTEMIS-1 mission. Therefore, we were able to leverage the LIC's FlatSat to test the DSN DAS request mechanism.

the beacon emulator and a variable RF attenuator (see **Figure 11**). Detection of the subcarrier was measured in the DSN receiver's GUI as a function of the integration time, which we varied from 1 to 100 seconds.

Figure 11. Setup for the Sensitivity Test

Results of this test are provided in **Figure 12** and indicate that there is a quasi-linear relationship between tone detectability and integration time, when both metrics are expressed in logarithmic scale. Note that this data is based on the fact that DSS-17 has a typical system noise temperature of 80K, more than twice as high as all other DSN antennas, and the fact that the tone detectability threshold is calculated using [16] and ranges from 6.8 dB-Hz at 1 second integration time to 1.02 dB-Hz at 100 seconds. It is expressed in units of dBm/Hz, i.e., it is a measure of the tone spectral density that must be exceeded for the receiver to indicate that a tone has been successfully received.

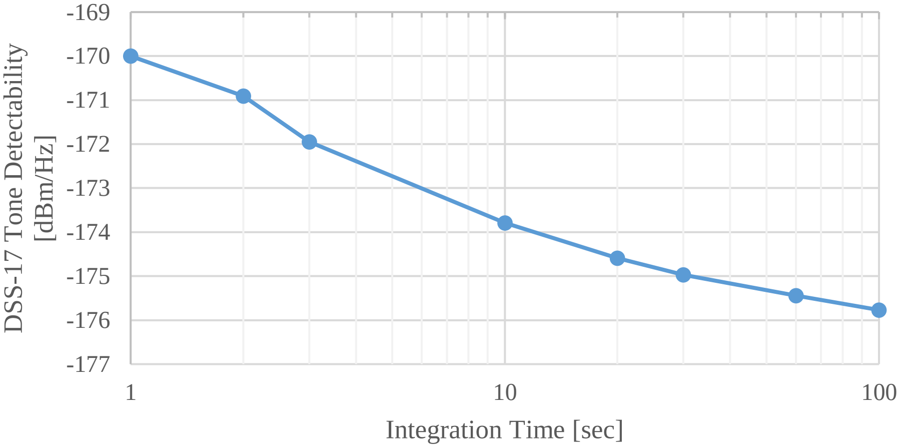

Figure 12. Tone Detectability at DSS-17

3.1.3 *Operational Tests: Operator Load during Queuing Antenna Operations*

The number of spacecraft that can be polled per unit of time depends on several technical and operational factors. For example, in [6] we report typical pointing and receiver configuration times for DSS-17 and use them to estimate the number of spacecraft that can be polled per hour as a function of the average polling track duration. Here, we aim at documenting the operational procedures that must be followed to perform queuing antenna operations at DSS-17 and quantify the operators load, measured in track prepration and execution time per spacecraft polled.

To conduct this test, we simulated the execution of four consecutive polling tracks. Using trajectory predicts from different spacecraft supported by the DSN, we identified four spacecraft that would be in view of DSS-17 during a span of 3 hours. We then configured the beacon emulator to transmit a 17 kHz subcarrier at X-band and mapped that subcarrier to a different tone for each of the four spacecraft (i.e., 17 kHz mapped to tone 1 for spacecraft 1, tone 2 for spacecraft 2, etc.). We also documented the set of steps required for the antenna operators at MSU to prepare and conduct the queuing antenna tracking activities.

Results indicate that the limiting factor when conducting queuing antenna operations at DSS-17 is the setup time per pass, which we currently estimate at ~10 minutes per spacecraft, with moderate economies of scale (a single spacecraft requires 12 minutes of setup time, while 7.5 minutes was measured as an average across the four simulated spacecraft). These times far outweigh the equipment setup times required for executing the polling track, which are on the order of a minute, and are driven by the fact that DSS-17 does not have the standard monitor and control software that allows for seamless and automated interface with DSN systems like SPS. As an example, the antenna pointing file generated by SPS is not directly compatible with DSS-17's pointing subsystem and must be translated using a script, a step that must be performed manually for each polling track. Finally, during the execution of the polling tracks, an operator has to be on console at MSU to update DSN receiver configuration manually every time a new spacecraft is tracked. This, again, is due to limited monitor and control capabilities.

3.2 Testing and Demonstration of the DAS Scheduling Algorithms

We conducted two demonstrations related to demand access scheduling: (1) Demonstration of the just-in-time block allocation; and (2) demonstration of the queuing antenna scheduling algorithm. All tests were conducted in the DSN scheduling system test environment, using existing interfaces to SSS, so as to replicate the system used for day-to-day DSN operations as much as possible. Demonstration of the block reservation mechanism was not performed since previous studies had already achieved this step, as documented in Reference [8].

Demonstration of the scheduling and assignment of demand access blocks was performed using as input a snapshot of the real negotiated schedule for the week of July 19 2021. We assumed that both Themis B and Themis C operated in demand access mode and that LRO tracks scheduled for that week were reserved for the DAS (i.e., as if LRO did not exist). Then, we simulated the process of allocating multiple reserved blocks in the SSS test environment over a span of multiple days, based on requests from Themis B and Themis C that were manually inputed in the system.

Figure 13. Allocation of Demand Acces Tracks in SSS

The same baseline scenario was then used to demonstrate how polling tracks would be scheduled using an automated queueing antenna scheduling algorithm. The placement of these polling tracks is based on a value function that takes into account the number of missions using the DAS, the antennas available to schedule demand access tracks and the periods of time when they are reserved to do so, geometric and visibility constraints, queuing antenna load constraints, and, optionally, a probability function that indicates how likely it is that a spacecraft generates a request. Note that constraints related to the queuing antenna load were accounted for in case queuing operations were performed with an antenna that also performs tracks with normally scheduled spacecraft, as would be the case with DSS-17. As an example, **Figure 14** shows the expected latency (red line) experienced by the Themis B and Themis C spacecraft, calculated over 1000 random simulation runs, as a function of the antenna busy percentage, which we define as the percentage of allocatble antenna time that is unavailable because of commitments other than queuing antenna operations. Here, latency is defined as a measure of how long the Themis B and C have to wait after a request before a demand access tracks is executed. It is reported normalized by the "optimal latency", which is calculated assuming a polling track time exists at the time a request is placed, and there is no contention in the system. As expected, the system performance degrades as the load of the queuing antenna increases (latency ratio equal to 1 is the best-case scenario) because it is less likely that the queuing antenna algorithm can place polling tracks at their optimal time.

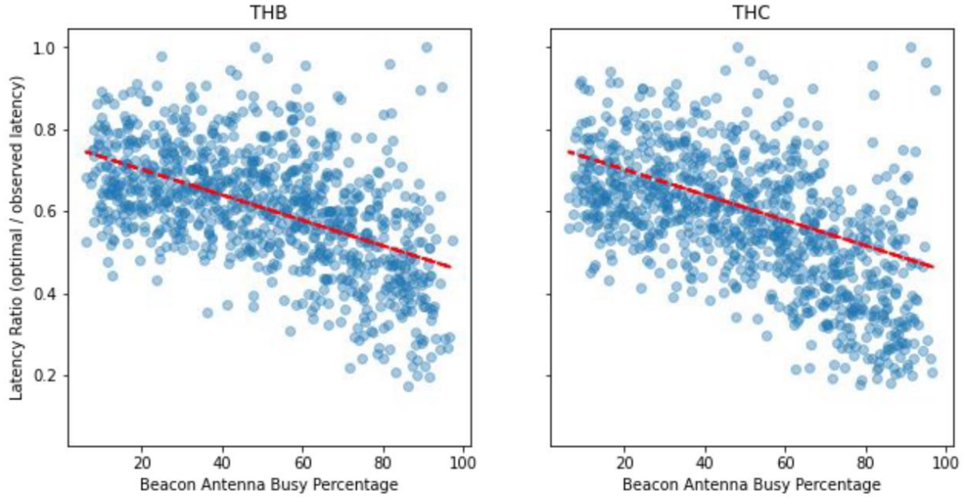

Figure 14. Latency Experienced with the DAS as a Function of Queuing Antenna (DSS-17) Load

3.3 Testing of the On-Demand Ground Data System

We tested the ability to control a GDS running on the cloud, in an automated manner, using the DSN scheduling system as the input to indicate when computational resources are allocated. To facitate the task, we worked with the LIC team at MSU to adapt the spacecraft's GDS for our tests, and used recorded telemetry generated by the spacecraft's FlatSat to ensure our emulation setup was as close as possible to reality.

The end-to-end architecture of this test is shown in **Figure 15**. The spacecraft was emulated by a server at the LIC's MOC which stores several files with CCSDS Space Packets [21] containing spacecraft telemetry. These files were streamed, one Space Packet at a time, into the Interplanetary Overlay Network (ION), JPL's implementation of the DTN protocol stack, which encapsulated them in bundles and routed the information to the next hop, another DTN node running at DSS-17.[‡‡‡‡] From there, data was forwarded to the Protocol Technology Lab at JPL,[§§§§] which finally forwarded the data stream to a cloud-based instance of AMMOS Instrument Toolkit, a telemetry processing software also developed at JPL and capable of parsing and displaying the telemetry contained in the Space Packets [22].

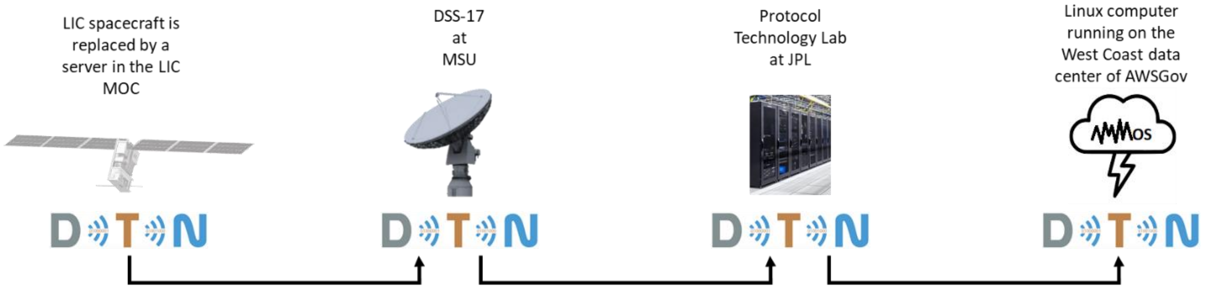

Figure 15. End-to-End System for Testing the On-Demand GDS

[‡‡‡‡] Because we partnered with MSU for all our tests, the testing setup had both the MOC spacecraft and the queuing antenna collocated in the same physical space, the Space Science Center at MSU. However, our tests did not assume that this was necessary for system operation.

[§§§§] Due to cybersecurity restrictions, a direct connection between the DTN nodes at MSU and the GDS instance running on the AWS GovCloud could not be established. Therefore, the Protocol Technology Lab was used as an intermediate node, a design choice enabled by DTN and that would have been significantly more challenging to implement using traditional SLE-based interfaces.

Copyright 2023. All rights reserved. Published by the Mohammed Bin Rashid Centre (MBRSC) on behalf of SpaceOps, with permission and released to the MBRSC to publish in all forms.

To manage the cloud-based GDS instance, we first manually created an AWS instance with AIT and ION installed, configured both application as services so that they would run upon system start up, and saved an image of this instance on the cloud for later use. Then, we created a script that used SSS' REST interface together with the AWS Python library [23] to query the DSN schedule and start/stop new instances automatically by loading the pre-saved imaged. Finally, to monitor the system and ensure telemetry was flowing from the spacecraft to the cloud, we deployed an Elastic-based monitoring system that automatically parses and indexes the logs generated by ION and provides summary visualization metrics in a web interface.

Using this setup, we conducted several tests to demonstrate functional validation of the system, including:

- Demonstrating that the GDS instance could be spin up and down on the cloud using information extracted from activities in SSS.
- Processing of the telemetry was performed in the cloud GDS instance, with results displayed in real-time on a web-based GUI accessible from a computer within the JPL network.
- Testing the DTN store-and-forward mechanism and the DTN monitoring system by breaking some of the network connections and observing drops in the telemetry stream that were later recovered once the connections were brought back online.

4. Conclusion

This paper describes the systems engineering and technology development effort conducted at JPL during Fiscal Year 2021 and 2022 to define and prototype a demand access system for deep space operations. This work is motivated by (1) the expected growth in DSN tracking time over the next coming decades, which is driven by the advent of deep space smallsats, renewed interest in lunar exploration, and a set of new missions to explore the outer solar system; and (2) the development of autonomous spacecraft that no longer operate “blindly” based on a set of pre-planned activities and their associated sequences.

The primary goal of a demand access system is to let spacecraft request unplanned and/or unpredicted ground support whenever an on board event motivates it. This contrasts with the current way of performing operations, which is based on long-lead planning cycles, sequencing activities and scheduled DSN support, and leverages advances in autonomy and automation both in space and on our ground data systems. It has been conceived for deep space operations and, consequently, does not rely on acknowledgements that incur in costly round-trip light time delays. It has also been designed to take advantage of a new type of antenna, which we termed queuing antenna, that can be used to periodically “poll” spacecraft operating in demand access mode whenever they point to Earth and, in the process, offload some of the tracking time from the already heavily subscribed DSN 34m and 70m antennas.

The demand access system is built around three core capabilities, a queuing antenna to receive spacecraft requests, a set of scheduling algorithms to first reserve DSN antennas time and allocate it to mission in near real-time based on their requests, and an on-demand GDS that runs on the cloud and is responsive to allocation of new demand access tracks. A description of the different architectural elements of the demand access system was provided, together with details on the technical implementation prototyped over the last two years.

Several subsystem-level demonstrations were performed, including the request mechanism and the queuing antenna, the scheduling algorithms needed to operate the DAS, and the on-demand GDS on the cloud. Together, these subsystem-level tests provide the first (partial) prototype of a demand access system for deep space operations, and provide a proof-of-concept to build upon for operationalizing the system. Future work includes performing system-level tests in which the three described capabilities are interfaced seamlessly with each other, and emulation of the end-to-end concept of operations as described in this paper. This can be accomplished initially using ground analogs, much in the same manner that we emulated the spacecraft for testing the ground data system, but ultimately a flight demonstration is of interest. Originally, we hoped to perform such a demonstration with Morehead Lunar IceCube in 2023, but we are currently searching for new candidate missions that are interested in further pursuing this concept.

Acknowledgments

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, and at Morehead State University, under contract with the Jet Propulsion Laboratory (80NM0018D0004).

Copyright 2023. All rights reserved. Published by the Mohammed Bin Rashid Centre (MBRSC) on behalf of SpaceOps, with permission and released to the MBRSC to publish in all forms.

References

[1] Jet Propulsion Laboratory, California Institute of Technology, "DSN Telecommunications Link Design Handbook," Jet Propulsion Laboratory, Pasadena, 2000.

[2] D. J. Scheeres, "Janus: A NASA SIMPLEx Mission to Explore Two NEO Binary Asteroids," in *7th IAA Planetary Defense Conference*, 2021.

[3] R. J. Lillis, "The ESCAPADE mission to Mars: Escape and plasma acceleration and dynamics explorers," in *AGU Fall Meeting Abstracts*, 2018.

[4] J. Kasper, J. Lazio, J. Romero-Wold, J. Lux and T. Nielsen, "The sun radio interferometer space experiment (SunRISE) mission concept," in *2020 IEEE Aerospace Conference*, 2022.

[5] J. Schoolcraft, A. Klesh and T. Werne, "MarCO: interplanetary mission development on a CubeSat scale," *Space Operations: Contributions from the Global Community,* pp. 221-231, 2017.

[6] M. Sanchez Net, "Enabling a Larger Deep Space Mission Suite: A Deep Space Network Queuing Antenna for Demand Access," in *IEEE Aerospace Conference*, Big Sky, MT, 2022.

[7] M. McKinnon, "ngVLA: The Next Generation Very Large Array," in *Bulletin of the American Astronomical Society*, 2019.

[8] T. M. Hackett, M. Johnston and S. G. Bilen, "Spacecraft Block Scheduling for NASA's Deep Space Network," in *2018 SpaceOps Conference*, 2018.

[9] G. R. Eaglin, "21m Space Tracking Antenna - Deep Space Station 17," Morehead State University.

[10] N. Dhamani, M. D. Johnston and G. Lucena, "A Demand Access Paradigm for NASA's Deep Space Network," in *12th Internatoinal Workshop on Planning and Scheduling for Space (IWPSS)*, 2021.

[11] Jet Propulsion Laboratory, California Institute of Technology, "Deep Space Network Service Catalog," Jet Propulsion Laboratory, Pasadena, 2022.

[12] NASA, "Core Flight System (cFS) Limit Checker (LC)," [Online]. Available: https://github.com/nasa/LC. [Accessed 29 December 2022].

[13] Jet Propulsion Laboratory, California Institute of Technology, "F Prime," [Online]. Available: https://nasa.github.io/fprime/. [Accessed 29 December 2022].

[14] M. K. Simon, M. Shihabi and T. Moon, "Optimum Detection of Tones Transmitted by a Spacecraft," *The Interplanetary Progress Network Report,* Vols. 42-123, 1995.

[15] M. K. Simon, V. Vilnrotter, A. Mileant and S. Hinedi, "Optimum Strategies for Monitoring the Operational Status of a Spacecraft," *The Interplanetary Network Progress Report,* Vols. 42-127, 1996.

[16] G. Lanyi and R. Kahn, "Tone Detection Via Incoherent Averaging of Fourier Transforms to Support the Automated Spacecraft-Monitoring Concept," *The Interplanetary Network Progress Report,* Vols. 42-129, 1997.

[17] M. K. Sue, T. Peng and E. J. Wyatt, "A New Approach in Spacecraft Monitoring for Efficient Use of the Deep Space Network," *The Interplanetary Network Progress Report,* Vols. 42-130, 1997.

[18] Deep Space Network, "0162-Telecom Beacon Telemetry Standard Formatted Data Unit (SFDU) Interface.," Jet Propulsion Laboratory, California Institute of Technology, Pasadena, DSN No. 820-013, 0162-Telecom, Rev. A.

[19] The Consultative Committee for Space Data Systems, "Cross Support Concept - Part 1: Space Link Extension Services," CCSDS 910.3-G-3, 2006.

[20] Internet Engineering Task Force, "Bundle Protocol Version 7," Internet Engineering Task Force, 2022.

[21] The Consultative Committe for Space Data Systems, "Space Packet Protocol," 2020.

[22] NASA, "AMMOS Catalog," 2022.

[23] Amazon Web Services, "Python on AWS," [Online]. Available: https://aws.amazon.com/developer/language/python/. [Accessed 29 December 2022].

[24] M. Soriano, S. Finley, D. Fort, B. Schratz, P. Ilott, R. Mukai, P. Estabrook, K. Oudrhiri, D. Kahan and E. Satorius, "Direct-to-Earth Communications with Mars Science Laboratory during Entry, Descent and Landing," in *2013 IEEE Aerospace Conference*, 2013.