Scheduling Within a Demand Access Paradigm for NASA's Deep Space Network

Nihal Dhamani, Mark D. Johnston, Girly Lucena
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Dr., Pasadena CA 91109
nihal.n.dhamani, mark.d.johnston, girly.d.lucena @ jpl.nasa.gov

Abstract

NASA's Deep Space Network (DSN) is the primary resource for communications and navigation for interplanetary space missions, for both NASA and partner agencies. As part of an investigation into improved efficiency and responsiveness, we have been exploring and prototyping the infusion of a "demand access" model into the DSN scheduling process. Today, DSN is fully pre-scheduled in advance, and many users rely on a stable schedule to plan their own spacecraft activities, weeks in advance of execution. However, a new class of missions is emerging that may not be scheduled as far in advance, and may be event-driven in coming across science targets at unpredictable times. These users could take advantage of an on-demand mechanism to download data. In this paper, we describe a prototype of the demand access process which consists of reserving blocks of shared antenna time (demand access tracks) and utilizing special "beacon" tracks to request access to those blocks. Furthermore, we dive into a proposed scheduling solution for the beacon tracks using an aggregated value function in order to minimize data latency and track starvation.

Introduction

NASA's Deep Space Network (DSN) consists of three communications complexes, located in Goldstone, California; Madrid, Spain; and Canberra, Australia. Each complex contains one 70-meter antenna and three or four 34-meter antennas. These ground antennas are responsible for communications and navigation support for a wide range of scientific space missions, from those in highly elliptical earth orbits, to some beyond the solar system. The placement of the three DSN complexes allows at least one of them to be in view of any distant spacecraft at all times.

The current DSN scheduling process (Johnston and Lad 2018) starts roughly six months ahead of execution when missions enter detailed requirements for their tracking. These requirements are submitted into the Service Scheduling Software (S3) (Johnston et al. 2014), where a variety of scheduling strategies are run to integrate the requirements into a single initial schedule. Next, a human scheduler, called the Builder of Proposal (BOP) makes further changes to the schedule, often eliminating hundreds of conflicts through a time consuming and labor intensive process lasting about a week. This is followed by a negotiation phase, also about a week in duration, where the schedule is released to all mission representatives who must collaboratively negotiate any changes or updates to the schedule. Any further changes to the schedule need to be mutually agreed upon by all of the representatives of missions affected by the change.

Following the deconfliction and negotiation phases, the schedule is baselined and usually remains stable for about 22 weeks (4-5 months) prior to execution. This allows missions to use the schedule from the DSN as an input to their internal spacecraft planning and sequencing processes, which can be labor intensive and time consuming. For deep space missions with long light-travel times, sequences are pre-loaded on-board weeks ahead of time, and late changes can be difficult. Adding onto the already complicated DSN scheduling problem is the fact that in recent years, the increased number of missions and the increased data return from missions has led to high oversubscription of DSN resources, leading to a very limited availability of DSN antenna time.

Some newer mission concepts and proposals would benefit from a more dynamic scheduling process. Lower cost missions that may have smaller operations teams, such as cubesats, may be power-limited and have some constraints that may not be accurately modeled months in advance. Additionally, some science missions are proposed that would respond to ephemeral phenomena such as outbursts on comets or asteroids, or astrophysical transients such as flares or bursts. For these kinds of missions, predicting the times when a future download would be needed is not possible, and yet their allocation of DSN time would not be large, likely only a few tracks per week. Time would be wasted if the mission had nothing to report, and conversely, if there was something to report but the next scheduled track was a week away, there would be a significant delay introduced in getting data to the ground so that follow-up observations could be scheduled. As such, the traditional static pre-scheduling method described above is not favorable for smallsats and event-driven missions as it lacks adaptability in response to changing events. The proposed solution would be to incorporate an on-demand

Copyright ©2022, California Institute of Technology. Government sponsorship acknowledged.

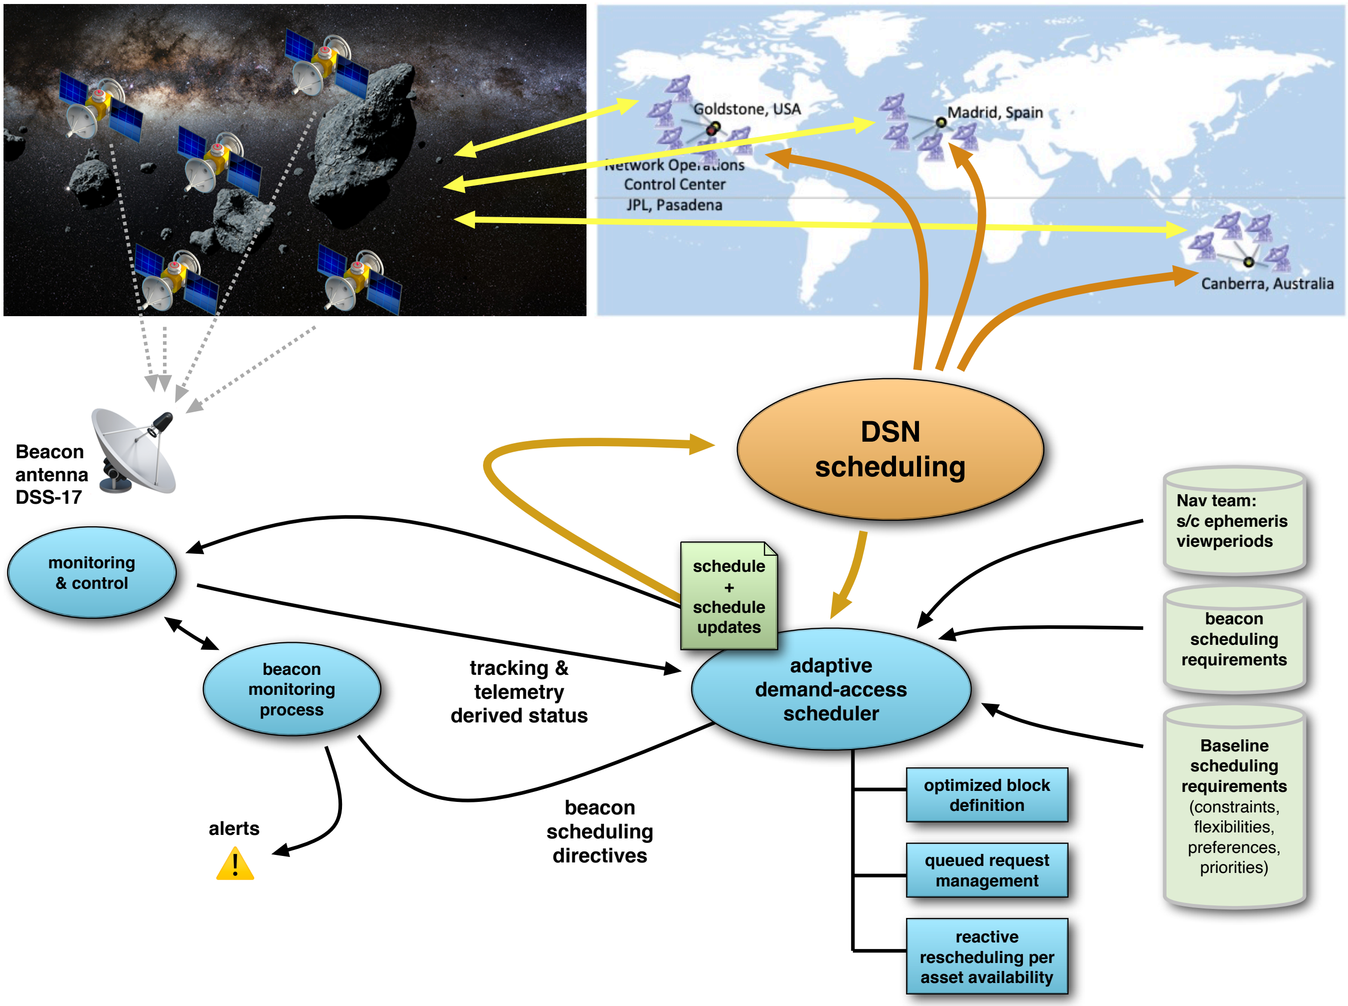

Figure 1: An overview of the demand access concept for DSN. A fleet of autonomous spacecraft, for example exploring the asteroids, sends beacon tones requesting contacts if needed. Pre-allocated optimized tracks are in the DSN schedule and allocated to a specific mission just in time.

communications approach, where spacecraft could request tracking time and be scheduled on an as-needed basis, referred to as “demand access” (Johnston and Wyatt 2017; Dhamani, Johnston, and Lucena 2021). The concept of demand access and responding to changing events is not novel and many systems have been developed to accomplish these tasks. For instance, the Demand Access Network Scheduler (DANS) was designed to be capable of automatically rescheduling antenna and subsystems in the event of changing track requests or equipment outages (Chien, Lam, and Vu 1997). Additionally, NASA’s Space Network also defines its own Demand Access Service (DAS) for some usage of the Tracking and Data Relay Satellite System (TDRSS), where spacecraft orbiting near the Earth can request tracking time through a handshake protocol (Gitlin, Kearns, and Horne 2002). Similarly, the Space Mobile Network also utilizes the concepts of user-initiated services (UIS) to request links to communications resources (Israel et al. 2018). These services are non-preemptive and provide for late allocation of available capacity. While there has been previous work done in the domain surrounding demand access scheduling, the increasing oversubscription of DSN resources and human-driven scheduling process of the DSN present unique challenges. Ultimately, any demand access paradigm for the DSN must be able fit within the semi-manual scheduling framework and needs to coexist with stable long-term allocations that many missions require in order to develop their on-board command sequences. The demand access work described in (Dhamani, Johnston, and Lucena 2021) and extended upon in this paper, expands upon the traditional demand access models by introducing a way to roughly model anticipated demand, generate optimized tracks to accommodate the anticipated demand, and reallocate holding time in the schedule to meet an anticipated demand when a request comes in, all while existing within the regular DSN scheduling process. For the DSN, the cost of preemption is high in terms of disruption to deep space mission operations, and so pre-allocation for anticipated demand provides a buffer against late schedule changes.

Demand Access Overview

While the notion of "demand access" for space communications networks has been available for some time for Earth orbiters and geosynchronous relay satellites as discussed in the introduction, incorporating the demand access model into the DSN presents some unique challenges that must be met to make such an approach work, most notably:

- existing users who participate in the static DSN scheduling process rely on a stable schedule, and so the infusion of demand access must not disrupt these users
- must be able to exist within the context of highly oversubscribed DSN resources
- a key objective of demand access is to utilize DSN resources more efficiently, and to service DSN users in a more timely manner – so approaches that do not accomplish these objectives will not be viable

The demand access approach outlined in (Dhamani, Johnston, and Lucena 2021), consists of the following steps:

- Users input mission requirements, specifying expected tracking parameters, durations, and frequency.
- The demand access system aggregates and optimizes mission requirements into larger generic "pseudo-spacecraft" blocks which contend for track time just like any other mission in the DSN scheduling process.
- In near real-time, the demand access system dynamically allocates the generic blocks to individual spacecraft depending on the requests received via the beacon (also called "queuing") antenna.

The resulting process of demand access track generation and subsequent real-time allocation is fully automatic – given the set of missions and their expected demand access requirements, it generates and uploads DSN scheduling requirements to reserve blocks of time that are optimized for sharing among demand access users, while accommodating their overall expected tracking needs. This introduces the concept of a "demand access" track, i.e, generic tracks that can be assigned to any mission supported by that track. These demand access tracks would be included in the normal deconfliction and negotiation process and would be present in the schedule as generic, i.e. not allocated to a specific mission until the near real-time dynamic allocation process. Figure 1 visualizes the addition of this demand access process on top of the regular DSN scheduling architecture.

Beacon Tracks

Beacon Tracks Overview

As we have discussed, DSN antenna time is a scarce resource. As such, the concept for demand access incorporates a secondary antenna at another site with a smaller antenna diameter and shorter integration time that can detect a "beacon" tone from a spacecraft (Wyatt et al. 1998). Beacon tones can have a small range of values, usually 4 to 8, that can be used to indicate that the spacecraft is healthy or not, and that it has data to download and with what level of urgency. The tone is a one-way signal – it is not feasible to return an acknowledgement. Beacon tones are already operational, most notably on the New Horizon's spacecraft which utilized beacon tones on the way to Pluto (Kusnierkiewicz et al. 2005).

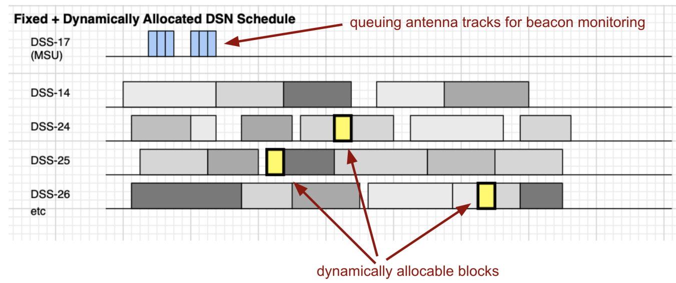

Figure 2: An overview of how the beacon tracks fit within the overall demand access schedule. The beacon tracks are shown in blue while the demand access tracks are shown in yellow. Note that in this diagram, the beacon antenna (DSS-17) is shown to have no other track reservations or avoidance times.

.

In the case of demand access, beacon tones can be useful for spacecraft to request track time via a demand access track. Once a tone is received that indicates a need for contact, the request is put in the queue for allocation from a set of pre-specified demand access tracks in the schedule. The beacon tone requests are sent on a smaller antenna called either the "beacon antenna" or the "queuing antenna". For this study, the Morehead State University 21-meter antenna is considered the beacon antenna since it has been outfitted with DSN-compatible S- and X-band equipment.

While beacon tones combined with the beacon antenna provide a useful way to communicate need for demand access track time, the spacecraft also needs to know the times at which it is allowed to send beacon tones. As such, spacecraft must maintain a set of precomputed track times onboard during which the spacecraft can communicate with the beacon antenna through beacon tones. This set of on-board tracks is referred to as "beacon tracks". Beacon tracks give the spacecraft the opportunity to indicate desire (or not) to utilize the upcoming demand access track. Figure 2 showcases how the yellow demand access tracks are scheduled within the regular DSN schedule and also shows the blue beacon tracks which are scheduled on the beacon antenna.

The concept of beacon tracks begs the question of how should these tracks be scheduled? Since the scheduling of beacon tracks exists within the larger framework of the demand access paradigm, these tracks must be scheduled in a way that promotes overarching goals of utilizing DSN resources in an efficient manner. As such, the objectives of scheduling these beacon tracks can be summarized as the following:

- service with minimal latency
- no starvation – ensure the right spacecraft are polled in time to make timely and fair demand access allocations
- maintain efficiency – don't waste or under-utilize demand access blocks on DSN antennas

Beacon Track Scheduling

Before we dive into a proposed solution of how to best schedule these beacon tracks, we must outline a few assumptions to our scenario. In our scenario, we are scheduling on a rolling weekly basis where each week is composed of three different types of tracks for each mission (see Figure 3):

1. **F**: Fixed (dedicated) pre-scheduled 2-way tracks: for upload (command sequences, schedule changes, etc.) and download (science data/housekeeping/engineering)
2. **B**: Beacon: (dedicated) pre-scheduled opportunities for spacecraft to send a beacon tone to the beacon antenna indicating desire for a subsequent demand access track
3. **D**: Demand access: dynamically allocable tracks based on queued requests. From the spacecraft point of view, these tracks are potential, not actual, until allocated. Each demand access track has a decision time trigger (1 hour) before the start of the track by when it must be allocated – to allow time to update the schedule, generate and send requests to the beacon antenna, etc.

Additionally, there are some high-level considerations for when to schedule beacon tracks. Ideally, there needs to be one beacon track for each mission that a demand access track supports. Some considerations for scheduling each of these beacon tracks are:

- Between fixed tracks
- Avoid beacon antenna outages or other prior commitments
- Avoid spacecraft-determined avoidance times
- In view of beacon antenna
- Maximize probability of having science ready to communicate. For some missions this could uniform; for others, probability could peak up after flyby, favorable lighting conditions, or other situations
- Prefer as late as possible, but earlier than demand access trigger cut-off time

The proposed solution to scheduling a beacon track would be to construct an aggregated value function $V(t)$ which includes the considerations described above. In this way, the beacon track would be scheduled at a spot with the maximum value. Since the scheduling of a beacon track would update the times during which the beacon antenna would be busy, the value function would also have to be updated after each beacon track iteration. As such, the value function would be generated for each mission and for each beacon track until all tracks have been scheduled. Should no valid times exist for placement of a beacon track according to the value function, the scheduling of that specific beacon track would be skipped. To provide more detail:

- Each spacecraft can generate a beacon track value function $V_s(t)$. Higher value = better beacon track placement
- At avoidance times t_a where a track cannot be scheduled, $V_s(t_a) = 0$
- Can only be scheduled at times, t_v, when spacecraft is in view
- Includes an event probability timeline, E_s, indicating the probability of having science ready to communicate
- Includes heuristic timeline, H_s, to schedule the tracks as late as possible but before the trigger cutoff times.
- The goal is to schedule a beacon track in each interval I_j, where each is bounded at the front by either an F (fixed) or D (demand access) track for spacecraft S and at the end by a D (demand access) track minus the D_{tr} (trigger time)
- t_a will be updated as beacon tracks are scheduled as to not introduce conflicts
- For a fleet $S = 1 \ldots N$ spacecraft, we can generate a set of assignments and score it based on sum of the values $V_s(t_b)$ where t_b are the scheduled beacon times

As such, we can use each of these timelines to come up with a general formula, shown in equation 1, for constructing the aggregated value function for each spacecraft. Additionally, we can schedule the beacon track for each spacecraft and demand access track according to the algorithm 1. Once the beacon tracks are scheduled, the demand access process can continue as described previously.

$$V_s(t) = t_a \cdot t_v \cdot E_s \cdot H_s \tag{1}$$

Algorithm 1 Pseudocode for Beacon Track Scheduling

```
prev_t ← 0
b_t ← beacon track duration
score_map ← initialize
for each spacecraft S do
    for each demand access track D do
        score_map ← reset
        for t in range [prev_t, D_tr] do
            score ← get_score(t, t + b_t)
            score_map[t] ← score
        end for
        prev_t ← D_endtime
        schedule_beacon(max(score_map))
    end for
end for
```

Simulations

In order to validate and better understand the limits of our approach, we must test the beacon track scheduling algorithm under a variety of different environments. Testing the value function algorithm under different conditions of antenna busyness, varying track times, and different demand access groupings will help us evaluate the robustness and

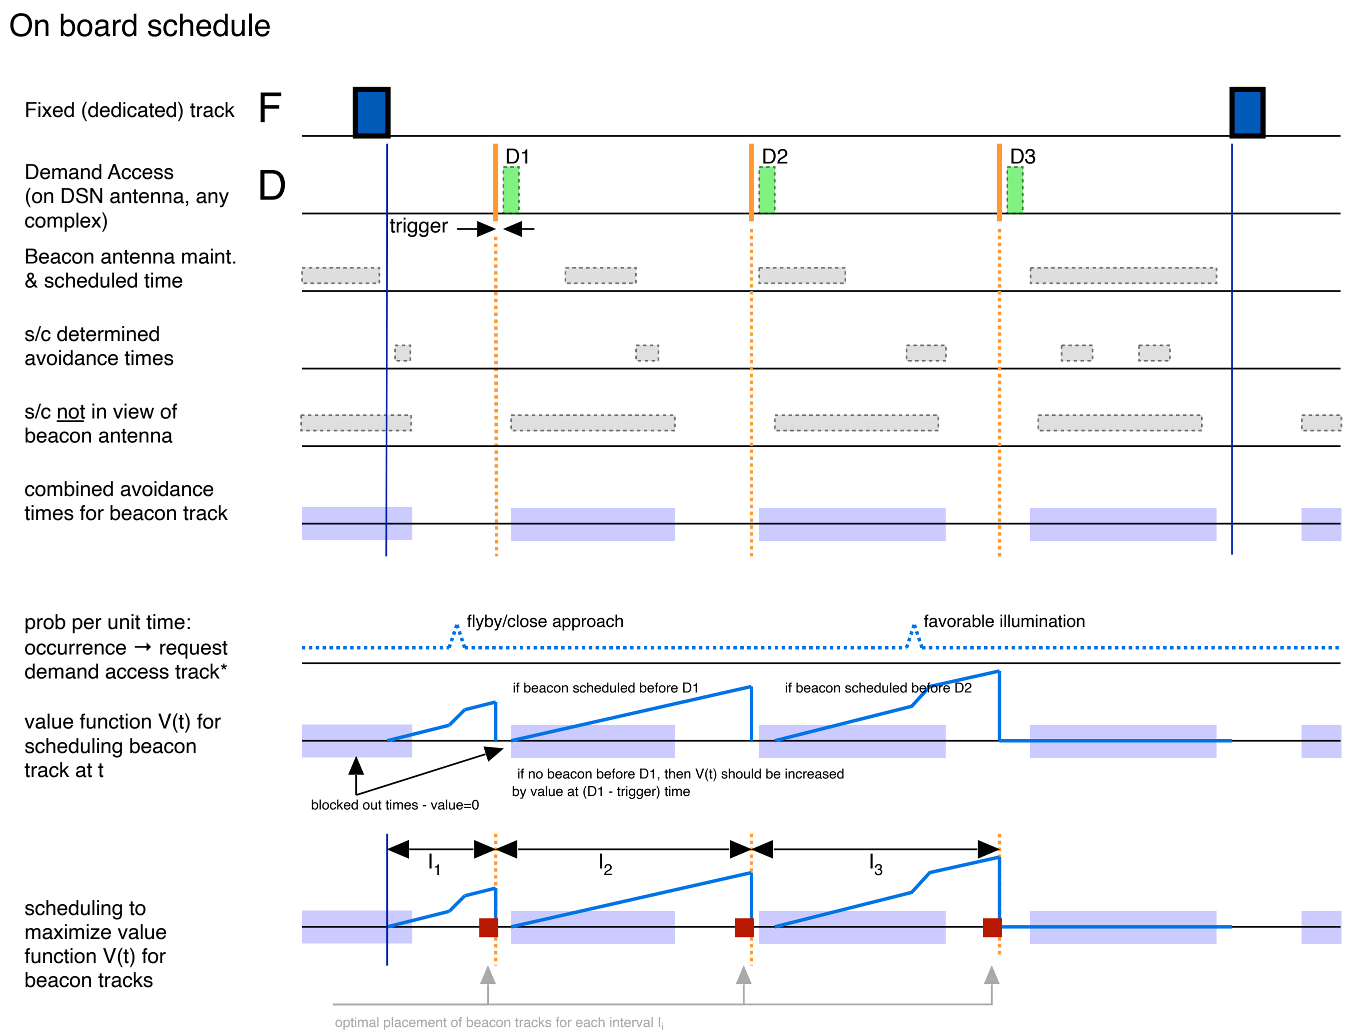

Figure 3: A visualization of the multiple components that end up forming the aggregated value function. Times before each demand access track where the value function is the local maximum provides us the best time for scheduling the beacon tracks.

limitations of our approach. However, since we are operating under the demand access paradigm which has no real operating schedules associated with it, we must rely on simulated schedules and attempt scheduling beacon tracks on simulated schedules with scheduled demand access tracks and busy avoidance times. In order to accurately assess the performance of the beacon track scheduling approach, these simulated tracks must resemble actual DSN weekly schedules.

As such, a simulation engine was developed to not only simulate the DSN schedules but also simulate the demand access tracks according to actual viewperiods, on-board spacecraft events, and scheduling of beacon tracks. Additionally, since the engine relies on random generation of schedules, it is highly configurable in order to tune the randomness parameters. The full capabilities of the simulation engine is as follows:

- Support different spacecraft, antennas, and time ranges.
- Generating avoidance times for each antenna (busyness) – with a random amount of busyness but spread out over whole week.
- Generating set of fixed tracks for each spacecraft across the set of given antennas.
- Generating random event probability distributions for each spacecraft to be utilized in calculating the value function.
- Ability to specify and group multiple missions into a demand access group.
- Ability to calculate antenna viewperiods for both single spacecraft and and combined viewperiods for demand access groups.
- Generating and scheduling demand access tracks of varying durations and separations for each group.
- Highly configurable. Varying model parameters to introduce randomness.

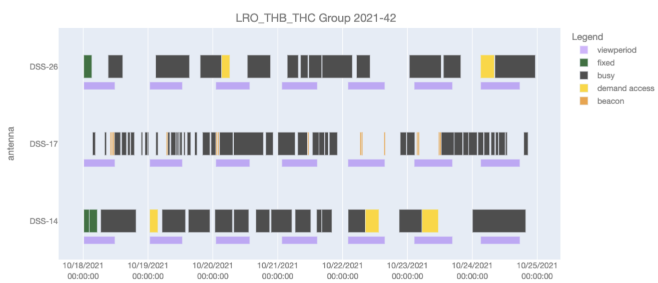

Figure 4: An example of a simulated schedule using the simulation engine for a demand access group consisting of Lunar Reconnaissance Orbiter (LRO) and THEMIS-B/C (THB/THC) lunar missions. The visualization includes both demand access tracks and scheduled beacon tracks. Additionally, the viewperiod availability for the demand access group is also shown for each antenna.

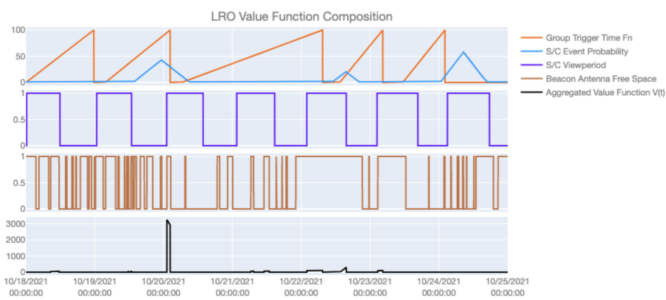

Figure 5: An example of a simulated aggregated value fuction calculation using the simulation engine for the Lunar Reconnaissance Orbiter (LRO) mission.

- Ability to run multiple times and to collect and store key metrics after each run.

These capabilities result in the simulation engine to be able to generate and schedule beacon tracks for each spacecraft according to the aggregated value function method under a variety of conditions – all while storing key metrics to assess performance for each run. Figure 4 visualizes of one such run of the simulation engine while Figure 5 shows one snapshot of the simulated value function for the LRO mission.

Metrics Collected

Initial metrics were collected with the simulation consisting of 3 lunar spacecraft (LRO, THB, THC) and using 2 DSN antennas (DSS-14, DSS-26) with the beacon antenna being DSS-17 at Morehead State. The simulation was run 1000 times to test beacon scheduling algorithm under varying conditions while collecting data for two main metrics: data latency and track starvation.

Latency is the measure of how long the spacecraft has to wait after an onboard event to send data back from either a fixed or demand access track. It is measured at each time step (every minute) and averaged over the course of the whole week. To expand, we compare both optimal and observed latencies. *Optimal latency* is calculated by assuming a usable beacon track exists at the trigger time of a demand access track. This represents the latency under ideal conditions when scheduling a beacon track (the beacon track is scheduled at the place that would minimize latency). Alternatively, *Observed latency* is calculated with the beacon tracks scheduled with the value function. This is the latency that is actually measured when we schedule beacon tracks according to our value function approach. In order to compare, we utilize the metric *latency ratio* which is the optimal/observed latency where *latency ratio* $<= 1$ and 1 is optimal. The latency ratio is essentially a measure of how optimally we were able to schedule a beacon track. Since we expect a lot of contention, the latency ratio is expected to be mostly less than 1 while also degrading with increasing antenna busyness. Figure 6 shows the metrics output of the simulation runs. We can see that the average latency ratio hovers around 0.7 until we hit a beacon antenna busy percentage of 60%.

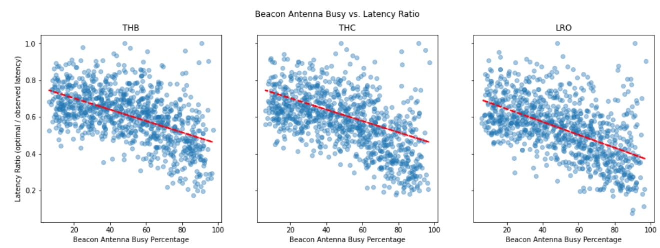

Figure 6: Beacon Busy Percentage vs. Latency Ratio of each mission in LRO, THB, THC group. Consisting of 1000 simulation runs. Beacon antenna busy percentage refers to the percentage of the allocable time (where the spacecraft is in view of the antenna) that is unavailable because of other commitments.

Track starvation: Additionally, we collected metrics on track starvation, i.e. the average number of beacon tracks that remain unscheduled at every busy percentage value. Since we want to efficiently utilize each demand access track and give each mission the ability to send a beacon tone to utilize a demand access track, the number of beacon tracks left unscheduled gives us insight into how likely it is that missions and demand access tracks will be starved if the beacon antenna is busy with other commitments (busyness). Figure 7 shows the collected metrics for track starvation. Once again we see that around the 60% busyness point, we see in a large increase in unscheduled beacon tracks.

Conclusions and Future Work

This study resulted in preliminary research on optimal scheduling of beacon tracks using an aggregated value function to determine best times. Additionally, the study resulted in the development of a highly configurable simulation engine where antennas, missions, and random generation can all be changed with a simple configuration file and has the

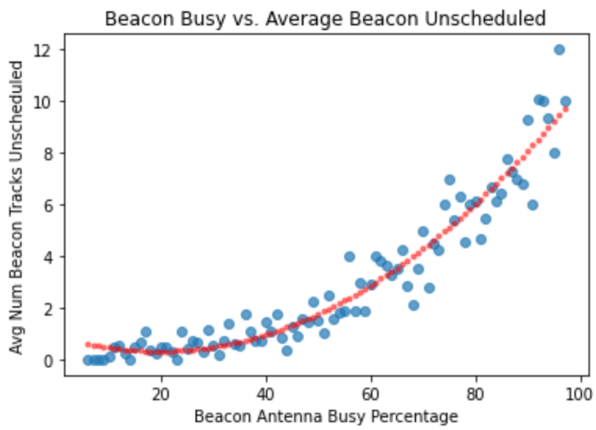

Figure 7: Beacon Busy Percentage vs. Avg Number of Beacon Tracks Unscheduled. Consisting of 1000 simulation runs. In this scenario, the maximum number of beacon tracks to be scheduled is 15.

ability to run thousands of simulations. Additionally, after analysing collected metrics, we saw that beacon antenna busyness greater than 60% results in much higher latency and higher amounts of starvation, giving insight into some limits of demand access on a single, shared antenna.

As we are in the very early steps of formalizing this approach, lots of future work has been planned. Firstly, we would like to repeat simulations with multiple demand access groups and more antennas to test this approach with as many different missions and scenarios as possible. Additionally, we would like to spend some more time tuning randomness parameters to produce more realistic DSN schedules. Furthermore, there is opportunity to further optimize the beacon scheduling algorithm as currently, each demand access track is treated independently – but could add a memory, e.g.: number of times mission beacon track left unscheduled should result in higher priority for next demand access track. As we progress further, we would also like to integrate the simulation engine with more real sources of data (actual DSN planned schedules) that can be augmented to produce the necessary schedules we need. And finally, we would like to explore more metrics that we can collect in order to better assess the performance of our scheduling approach.

Acknowledgements: This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

Chien, S.; Lam, R.; and Vu, Q. 1997. Resource scheduling for a network of communications antennas. In *IEEE Aerospace Conference (IEEE-Aero 1997)*.

Dhamani, N.; Johnston, M.; and Lucena, G. 2021. A demand access paradigm for nasa's deep space network. In *International Workshop on Planning and Scheduling for Space (IWPSS)*.

Gitlin, T.; Kearns, W.; and Horne, W. D. 2002. The NASA Space Network Demand Access System (DAS). In *SpaceOps 2002*.

Israel, D.; Roberts, C. J.; Morgenstern, R.; Gao, J. L.; and Tai, W. 2018. Space mobile network concepts for missions beyond low earth orbit.

Johnston, M. D., and Lad, J. 2018. Integrated Planning and Scheduling for NASA's Deep Space Network – from Forecasting to Real-time. In *SpaceOps*.

Johnston, M. D., and Wyatt, E. J. 2017. AI and Autonomy Initiatives for NASA's Deep Space Network (DSN). In *IJCAI AI in Space Workshop*.

Johnston, M.; Tran, D.; Arroyo, B.; Sorensen, S.; Tay, P.; Carruth, J.; Coffman, A.; and Wallace, M. 2014. Automated Scheduling for NASA's Deep Space Network. *AI Magazine* 35:7–25.

Kusnierkiewicz, D. Y.; Hersman, C. B.; Guo, Y.; Kubota, S.; and McDevitt, J. 2005. A description of the Pluto-bound New Horizons spacecraft. *Acta Astronautica* 57(2):135–144.

Wyatt, E. J.; Hotz, H.; Sherwood, R.; Szijarto, J.; and Sue, M. 1998. Beacon monitor operations on the deep space one mission. Technical Report 98-0712, JPL.