User Preference Optimization for Oversubscribed Scheduling of NASA's Deep Space Network

Mark D. Johnston*

*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena CA 91109
mark.d.johnston @ 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. With three complexes spread roughly evenly around the globe, the DSN provides services to dozens of active missions. Growth in mission demand, both in number of spacecraft and in data return, has led to increased loading levels on the network, and projected demand has exceeded network capacity for quite some time. The DSN scheduling process involves peer-to-peer collaborative negotiation, which consumes significant time and resources in order to reach a baseline version of the schedule, and then to manage and agree to changes. The delays inherent in this process are exacerbated by the high level of oversubscription experienced by the DSN: it is not unusual for the scheduling process to start with 20-40% more requested time can be accommodated on the available antennas. The other NASA networks make use of a static priority list to address a similar problem: missions are ranked in priority order, and the schedule is populated by priority from highest to lowest. Such a mechanism would not work for DSN due to the heterogeneity of the mission set, and to the time-varying mission requirements with mission phase. This paper reports on a new paradigm for DSN scheduling that addresses the key problems inherent in the current process. The main characteristics of the new approach are to use loading-based limits on requested time, and user preferences as the basis for optimization criteria.

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 (Fig. 1). 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. In future years, DSN will also support human missions to the moon and beyond. The placement of the three DSN complexes allows at least one of them to be in view of any distant spacecraft at all times (Imbriale, 2003).

The current DSN scheduling process is lengthy (with a lead time of around four months) and labor intensive. It relies on peer-to-peer negotiation for changes, with frequent proposals and counter-proposals, and so it is a major challenge to add a large number of new missions without impacting the current mission users. This calls for new approaches to scheduling that minimize the impact of adding new missions, while accommodating existing ones. At the same time, it is important to satisfy mission tracking and telecommunications requirements to the greatest extent possible, on an ongoing basis, week after week.

We have investigated a novel approach to this problem, characterized by the following:

1. Use advance (>6 months) planning/loading requirements, provided by users, to calculate overall anticipated loading, and then derive limits on the time allowed for submission later in the detailed scheduling process.
2. Require users to associate one of a small number (< 10) preference levels with each of their scheduling requests, thus providing a relative ranking of their own submissions, and then optimize the generated schedule to satisfy as large a fraction of higher preference requests as possible across the entire mission set.

We have prototyped a system based on this approach and run an initial series of experiments to investigate the quality of the schedules generated, as well as the degree to which user preferences could be met in practice. Results to date are encouraging, and it is planned to conduct more extensive trials in the near future. Moving to this paradigm could enable the DSN to greatly reduce the time and effort to build and manage schedules, while still allowing for unexpected late changes that are not uncommon in this domain.

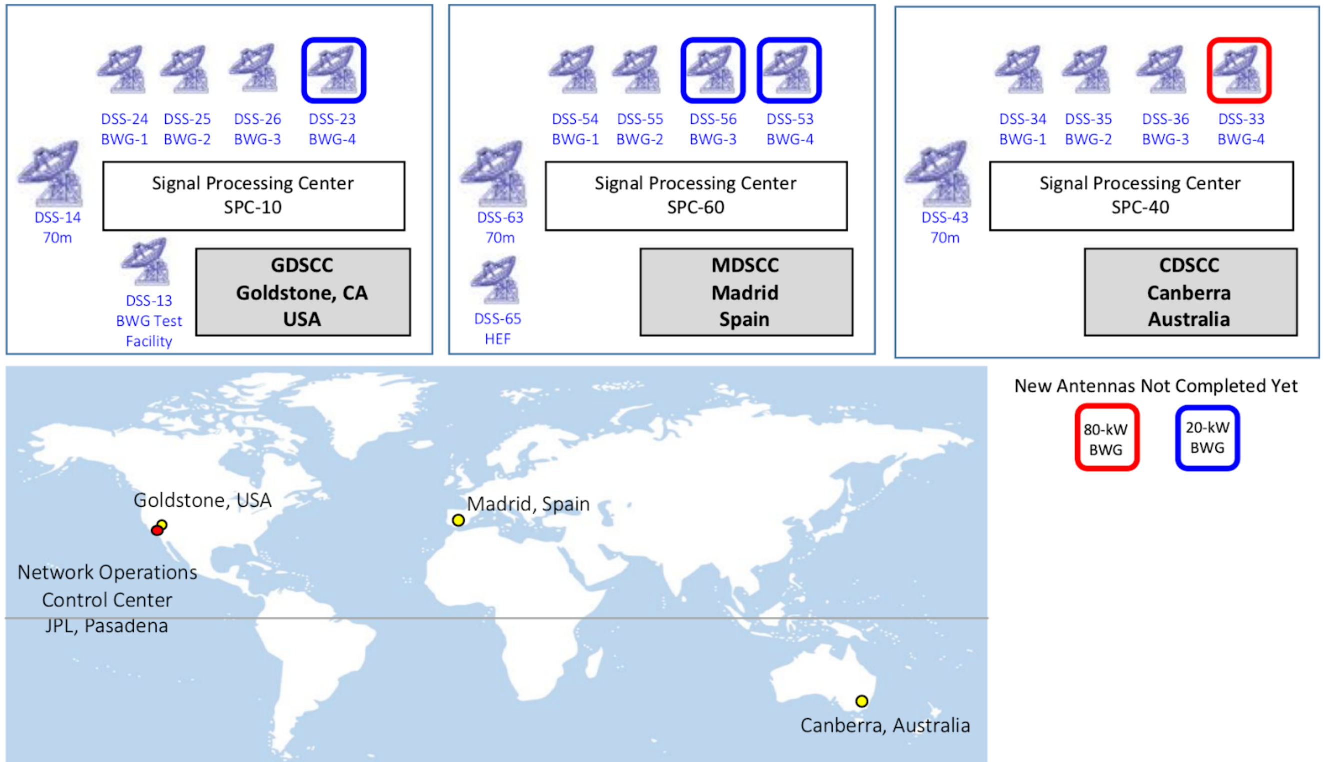

Fig 1. An overview diagram of the DSN antennas at the three complexes in Goldstone, CA, Madrid, Spain, and Canberra, Australia. As of 2019 there are 12 active antennas, expected to grow to 15 by 2024 with the construction of new antennas and the decommissioning of older ones.

In the following, we first briefly describe the DSN scheduling process, and then some of the factors that come into play as more missions are included in the DSN processes. This is followed by discussion of the new approach, and then by conclusions and directions for future work.

DSN Scheduling Process Overview

Fig. 2 shows a block diagram of the DSN planning and scheduling systems, indicating the key mission interfaces and data flows. On the longest time scale, the Loading Analysis and Planning Software (LAPS) provides for long-range planning and forecasting, potentially as much as a decade into the future, taking into account anticipated mission requirements and planned DSN asset capabilities. LAPS users include analysts and mission planners, as well as others involved in long-term planning of DSN activities. At this stage, the missions provide long-term ephemeris information along with their expected utilization of DSN resources.

The Service Scheduling Software (SSS, or S^3) maintains user-provided detailed scheduling requirements and expands them into specific communications and navigation passes, taking into account resource constraints (antennas and equipment) as well as a wide variety of DSN operational constraints and rules. S^3 supports the scheduling negotiation process, followed by a consensus-based change process with workflows for approval of all schedule modifications by authorized mission representatives. More accurate ephemeris information, along with specific communications and navigation requirements, are provided by the missions during this phase. At the conclusion of this phase, missions receive negotiated DSN allocations that they use to plan and sequence their onboard activities.

The Service Preparation Subsystem (SPS) merges the schedule with mission-provided sequence and ephemeris information to generate detailed DSN ground system sequences of events and a variety of predict calculations to support each scheduled activity. These are distributed to the DSN complexes where they are used by the Network Monitoring and Control subsystem to handle the execution of each pass.

DSN scheduling differs from the other NASA networks in large part due to the operating characteristics of most of its supported missions. Deep space missions typically do extensive advance planning and scheduling, to the point of

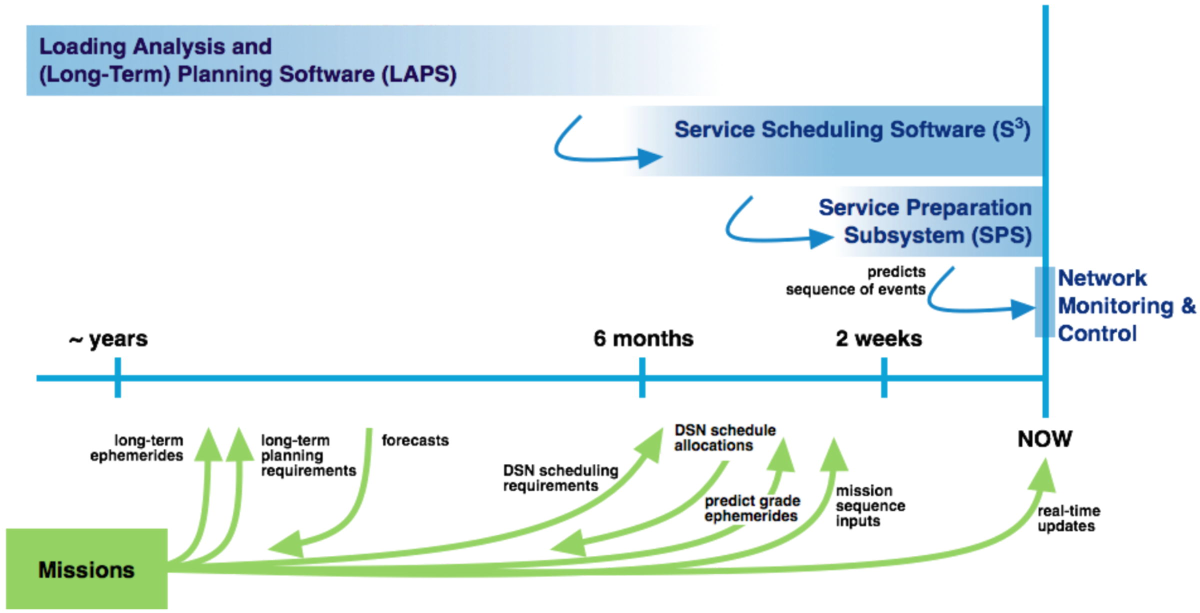

Fig 2. Schematic of the flow of information through the end-to-end planning, scheduling, and execution monitoring and control systems (top), highlighting interactions with mission users (bottom)

building detailed command sequences that are uploaded weeks or months ahead of execution, reflecting a range of mission phases, science events, and engineering activities. Additionally, there are extensive checks on these plans and sequences, as an error can be catastrophic. Long light travel times preclude extensive real-time interaction. As a result, the DSN schedule is baselined months ahead of time, with changes occurring only when agreed to by all involved.

The DSN scheduling process (Johnston et al., 2014) operates on a rolling weekly basis: as the deadline for a week approaches (roughly four months before the start of the week), mission scheduling representatives enter the requirements for that mission into the Service Scheduling Software, or S^3 (Johnston et al., 2012). Once all inputs for a week are in, they are integrated into a single schedule and the DSN Scheduling Engine (DSE) (Johnston et al., 2010), is run to deconflict as much as possible, given any specified flexibilities in the input requirements from each mission. In practice, little flexibility is allowed by users in their initial specifications, and the net oversubscription level means that many conflicts necessarily remain in the schedule.

Once the scheduling engine has been run, and conflicts reduced automatically as much as possible, a human scheduler called "Builder of Proposal", or BOP, starts to work on the schedule and makes further changes based on experience and background knowledge of each mission's requirements. These changes include: deleting some activities, shortening tracks below their specified minimums, splitting tracks designated as unsplittable, placing the (now shorter) segments into gaps in the schedule, and moving tracks to different antennas. This a time-consuming and labor-intensive process, requiring a great deal of familiarity with the entire DSN mission set and their typical requirement patterns and unstated flexibilities. The BOP can generally eliminate hundreds of conflicts, but at the end there usually remain 10-20 conflicting activities. At the conclusion of the BOP phase, the week is released to the full set of mission scheduling representatives to negotiate the remaining conflicts and to make any adjustments to changes made by the BOP. In this phase, individual mission representatives collaboratively negotiate peer-to-peer to reach a state where all users are agreed (Carruth et al., 2010). In this process, one user will propose a set of changes, to which all affected users must concur before it becomes the new baseline. If any user disagrees with the changes, it falls on him or her to counter-propose an alternative, with a justification (where just undoing a previous proposal is not allowed!). This process continues until the deadline is reached, at which point conflicts are either cleared or (rarely) waived, and the schedule is considered baselined and published. From the completion of the automated scheduling run to the baseline conflict-free schedule is typically 2-3 weeks. The overall duration of this process means that multiple weeks are being worked on in parallel, and about 18 weeks are in the pipeline in normal operations, with about 10 or more weeks negotiated and stable.

Priority Scheduling

DSN has been asked to look into developing a priority scheduling scheme for some time. The motivation for this includes a desire to reduce the time and effort to prepare and publish the baseline schedules, to reduce the effort of the BOP, and to improve consistency in the number of negotiated weeks that are available for missions to plan their detailed activities. It's also been suggested that DSN use a priority system for consistency with the other NASA networks, specifically the Near Earth Network (NEN) and Space Network (SN). Both of these networks use a very similar priority scheme as part of their routine scheduling processes. Note that DSN has adopted a *resource-based* prioritization policy for scarce 70m time, in that missions that can only use the 70m antennas are preferentially given time on them, with other users filling in as time is available. However this is very different from a *mission-based* prioritization.

Both NEN and SN use a "two-dimensional" priority scheme: the first and overriding dimension is often called "event" or "absolute" priority. It reflects the overall importance of an activity in some absolute way, as a relatively small number of categories and subcategories into which activities can be slotted. The second dimension is a strict mission ranking from *1* to *N*, such that more highly ranked missions have their scheduling requests satisfied first. The absolute priority list includes such categories as human spaceflight, launch and landing services, critical operations, etc. and goes down to normal operations and various kinds of network tests. The scheduling process works to satisfy user requests in order down the absolute priority list, then the mission priority list at the same absolute level, etc., until all requests are satisfied or there are no further placement possibilities.

NEN and SN have a process for developing and approving their priority lists, run by Goddard Space Flight Center and with concurrence by NASA HQ. These lists are updated about once per year, mostly to reflect the changing mission mix as new missions begin and old ones end. There are several important differences between the scheduling processes for NEN/SN and DSN:

- NEN and SN are entirely scheduled by a *central authority* based at White Sands, New Mexico: there is no mission-to-mission negotiation, and the full schedule is not published to the full user community, only that subset directly required for each mission
- NEN and SN both run their processes on a *very short timescale*: requirements are due two weeks before a week starts execution, and the confirmed schedule is published one week before. In the week before execution, mission priorities are not used and time is allocated on a first-come first-served basis

DSN does have an absolute priority list that specifies types of activities that are more important, as recognized by all participants in the process. These include support for human spaceflight, spacecraft emergencies and survival activities, major and unique scientific events, etc. all of which rank higher than normal DSN non-time-critical science activities. That said, only a small percentage of each week's time falls into these higher priority categories, with the exception of antenna maintenance. Launches and planetary orbit insertions and landings are examples of higher priority activities when they occur, as well as in-flight maneuvers. In the past year, fewer than 1% of over 24,000 scheduled DSN activities were at support levels higher than routine.

What DSN does not have is a ranked mission list: there has never been a process by which agreement could be reached on which *missions* are higher priority than others. Instead, DSN uses a peer-to-peer collaborative process to negotiate time and changes: even a small extended mission can engage in negotiation with a recently launched flagship with justification to make mutually agreeable schedule changes. There are other reasons why a ranked mission list is problematic for the DSN:

- Some missions require only a few hours per week to meet their mission goals, while others seek and can use nearly full time coverage. A strict mission ranking could *leave some missions completely out of the schedule*, which would undoubtedly be viewed as unacceptable
- Many DSN users have *requirements that change frequently* with mission phase or with planned science activities, week to week or day to day – not reflected well by a static priority list
- Given the *variable mix of activities and mission phase updates*, a mission priority list would have to change so frequently as to be essentially useless

All of these factors have played into the evolution of the current process and away from the static mission priority list adopted by the other NASA networks.

Previous investigations of priority schemes for DSN have made suggestions for addressing some of these concerns (Shouraboura et al., 2016), which will be considered after examining the role of oversubscription in the next section.

Oversubscription

The DSN is routinely oversubscribed by a variable amount, depending on the current mission set and their activities. An illustration of this is shown in Fig. 3, where the oversubscription level is plotted vs. time for a period of

30 weeks in 2018. The ramp up is due primarily to three factors: the arrival of the InSight mission at Mars in late November 2018, and the arrival of the two asteroid missions at their targets, OSIRIS-REx and Hayabusa2, also in the fall of 2018.

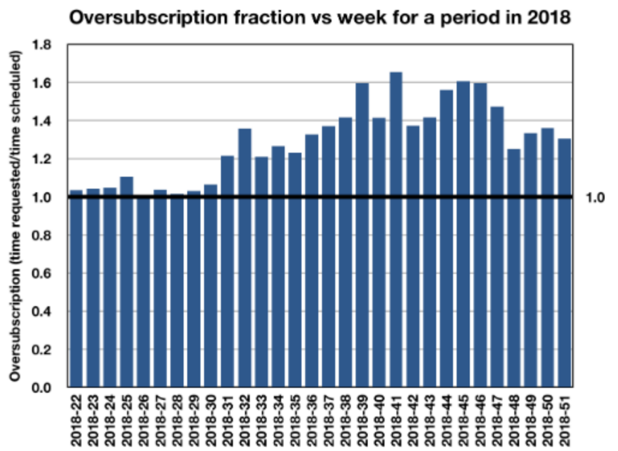

F*ig. 3. Oversubscription vs time for a period in 2018. The chart shows the ratio of requested time to actually executed time for a range of weeks, illustrating a peak overage of about 60%.*

For example, when a mission reaches its destination planet, it can go from minimal DSN usage to nearly continuous coverage for some period, sometimes months or years. When oversubscription reaches a level of as little as 10% it is equivalent to having one additional antenna's worth of demand to remove before a feasible schedule can be developed and published. At oversubscription levels of 60%, the scheduling engine algorithms, which are prohibited from dropping requirements, can do little more than try to spread out the overage. Most of the burden of dealing with the oversubscription falls on the BOP, who has to make wholesale reductions to come up with a feasible proposed schedule for negotiation. Negotiation itself is more arduous because missions try to accomodate the reductions and adjust the schedule to their best advantage. In spite of this, consensus is reached that the final schedule is sufficiently fair that the process of escalating disagreements is virtually never invoked.

User Preference Optimization

Figure 4 illustrates the overall lifecycle and timeline of a week in the DSN scheduling process. We first describe the key phases, and then comment on the introduction of process and software changes to address oversubscription.

1. Users must enter their scheduling requirements at a deadline, roughly 5 months ahead of start of execution of the week.
2. The requirements are integrated and conflict reduction algorithms are run, followed by human expert conflict reduction, in the stage designated "Builder of Proposal", or BOP.
3. When complete, the BOP releases the week to negotiation as a "Negotiation Workspace", or NWS.
4. When negotiation is complete, the baselined conflict-free schedule can still be changed by mutual agreement in a "Proposal Workspace", or PWS.

The actual concurred schedule is stored as the "Master" schedule, which is synchronized to the DSN complexes to control the execution of each tracking pass, engineering, science, and maintenance activity.

The concept being developed for infusing priorities, limits, and optimization is called "User Preference Optimization". This name is meant to indicate that "priorities" that pit *one mission against another* are not

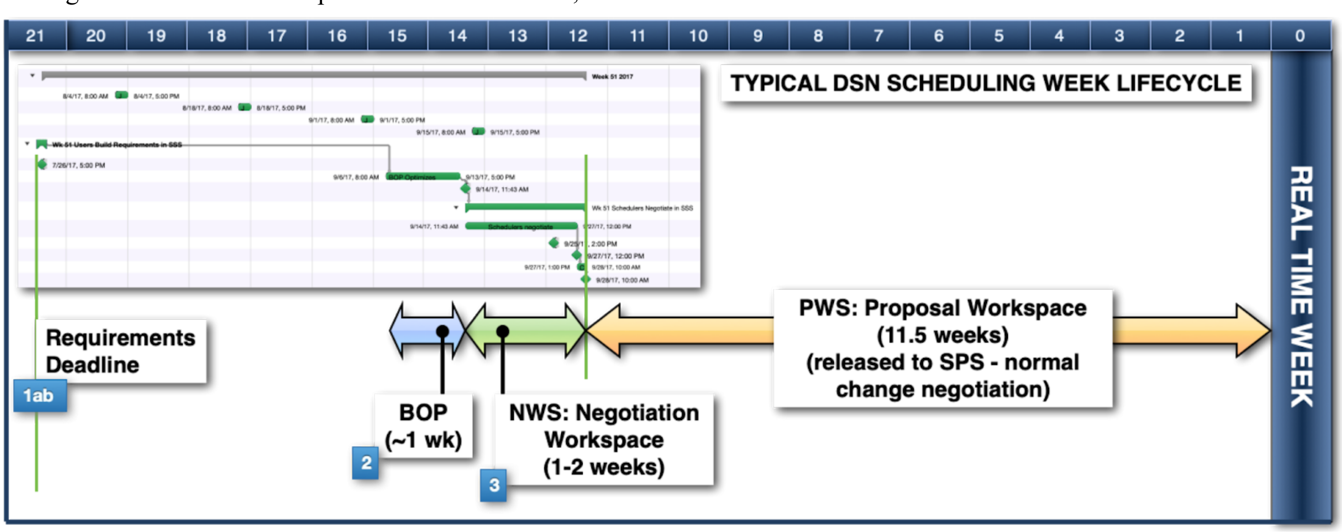

Fig 4. Lifecycle of a week in the DSN scheduling process, from requirements entry to execution. The main stages, as described in the text, including reducing conflicts with human assistance ("Builder of Proposal", or BOP), and negotiation of any remaining conflicts in order to baseline the schedule. The overall timeline is typically ~5 months in duration.

involved, rather that individual mission user preferences among their *own* activities are driving the process. The primary objective shifts from eliminating conflicts, to maximizing the weighted degree of satisfaction of user requirements, while eliminating conflicts by a combination of making use of flexibilities (stated and implicit), limiting input to a near feasible level, and dropping low preference requests when they cannot be made to fit without causing conflicts. These enter into the lifecycle illustrated in Fig. 4 as follows:

1a. Limit each week's submitted requirements to a cap calculated to be supportable by anticipated network resources (possibly allowing for a small margin to account for uncertainty). The cap would be calculated by the Loading Analysis & Planning S/W (LAPS), based on projected planning requirements from all missions, along with antenna engineering activities and downtime, and commissioning and decommissioning of antennas. This will require current long-range planning requirements to be provided by each mission (most do this already), and may require deconfliction if contention levels are too high when projected into the future.

1b. Require each mission to provide a relative prioritization (preference level) for their own inputs (only). This would be done by each user assigning every request to one of a small number of "levels" or "tiers" (<10). The significance of this is that it has users explicitly indicate what is most important to them, while avoiding any attempt to compare importance of one mission relative to another. The specification of user preference levels would be such that all requests at the same level are equally important, while any request in a higher level is more important than any request at a lower level. User preference levels would not replace the current DSN absolute priority scheme, but it would be expected that critical activities such as launches, orbit insertions, landings, etc. would all be given very high levels of user preference, corresponding to their treatment as DSN critical events.

2. At BOP time, run user preference optimization algorithms that search for feasible (conflict free) schedules that maximize the scheduled preference level for all missions. All provided flexibility will be used to fit as much as possible into the schedule, but requests that cannot fit will be dropped at this point (though may be negotiated in later by the affected missions). This flexibility can including shrinking requested tracking duration, splitting tracks across antennas or complexes, adjusting min/max gaps between tracks, etc. While the BOP is expected to still intervene as a human expert to make explicit adjustments in some cases, it is a goal to reduce the BOP level of effort by a factor of 5 over current levels. Following the BOP schedule release, users would negotiate as they do today. However, the time allowed to conduct negotiation could be reduced to ensure the pipeline of negotiated weeks remains consistently far enough ahead of realtime.

3. During, and post negotiation, only absolute priorities would be considered and consensus-based changes would be the norm. This is how the schedule is maintained today, in that late changes for spacecraft emergencies or hardware outages are accommodated, but changes impacting a mission must first be concurred by the affected participants. This is also how the other NASA networks handle changes after the schedule is released, except that changes are made by the central authority rather than peer-to-peer.

Status and Plans

A prototype system has been implemented to begin assessing this proposed approach. LAPS is currently being used to generate loading and forecast usage analyses for extended future time periods, and agreement with actual observed loading is very good. An interface is being defined to automate the incorporation of LAPS results into the midrange scheduling system, SSS, so users are aware of the limits they must adhere to. The result is expected to be that users will only enter a total demand that is supportable by DSN resources.

The second part of the paradigm change requires users to assign preference levels to each request, to indicate the relative degree of importance. The incentive to provide this information is the knowledge that lower preference requests will be omitted if they don't fit. This means less work for a user than if a high-preference request was omitted, since it would have to be negotiated back into the schedule, potentially requiring concurrence from numerous other missions for changes that frequently ripple to affect more and more missions.

Several algorithms are under investigation for the pre-BOP schedule deconfliction, including:

- conflict-directed hill-climbing search, using overall user preference level satisfaction cumulative to level i as an optimization metric
- squeaky-wheel optimization as a means to dynamically prioritize requests (Barbulescu et al., 2006b; Joslin and Clements, 1999; Lewellen et al., 2017)

Results to date are promising, but more experiments need to be done before significant conclusions can be drawn. One of the experiments being designed is to compare schedules generated with extensive manual editing vs. those prepared algorithmically, to identify areas where additional work is needed.

We are also looking into the potential value of multi-objective optimization (Johnston, 2006) as a means to evaluate tradeoffs among missions in the course of generating the schedule.

Conclusions

In this paper we have described a new paradigm for DSN scheduling that addresses several key objectives:

- Reduction in BOP manual schedule editing effort to reduce conflicts prior to schedule negotiation
- Reduction in negotiation effort and time to clear conflicts and baseline the schedule
- More uniformity in the number of negotiated weeks, to support extended (multi-month) mission planning and sequencing cycles
- Consistency with the other NASA SCaN networks (the Near-Earth Network and Space Network) in using priorities as a means to generate and deconflict schedules from a disparate set of user requirements

The approach we have described above has several compelling aspects:

- Users are asked for minimal additional information over what is already provided for scheduling
- The thorny issue of assigning relative priorities to missions is avoided
- The oversubscription of DSN resources is managed at the initial submission gate, thus speeding all downstream processes by not having to address what to reduce – this decision being left in the hands of the users who are best able to judge relative priorities of their own inputs

While there remains significant work to go to validate and fully implement this approach, it should provide a quantitative reduction in effort and cost in the DSN scheduling process.

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

Barbulescu, L., Howe, A., Whitley, D., 2006a. AFSCN scheduling: How the problem and solution have evolved. Math. Comput. Model. 43, 1023–1037.

Barbulescu, L., Howe, A.E., Whitley, L.D., Roberts, M., 2006b. Understanding algorithm performance on an oversubscribed scheduling application. J. Artif. Intell. Res. 27, 577–615.

Bretthorst, G.L., 1990. An Introduction of Parameter Estimation Using Bayesian Probability, in: Fougere, P. (Ed.), . Kluwer Academic Publishers, Dordrecht, pp. 53–79.

Carruth, J., Johnston, M.D., Coffman, A., Wallace, M., Arroyo, B., Malhotra, S., 2010. A Collaborative Scheduling Environment for NASA’s Deep Space Network. Presented at the SpaceOps 2010, Huntsville, AL.

Imbriale, W.A., 2003. Large Antennas of the Deep Space Network. Wiley.

Johnston, M., Tran, D., Arroyo, B., Sorensen, S., Tay, P., Carruth, J., Coffman, A., Wallace, M., 2014. Automated Scheduling for NASA’s Deep Space Network. AI Mag. 35, 7–25.

Johnston, M.D., 2006. Multi-Objective Scheduling for NASA’s Deep Space Network Array, in: Proceedings. Presented at the International Workshop on Planning and Scheduling for Space (IWPSS-06), Baltimore, MD.

Johnston, M.D., Tran, D., Arroyo, B., Call, J., Mercado, M., 2010. Request-Driven Schedule Automation for the Deep Space Network. Presented at the SpaceOps 2010, Huntsville, AL.

Johnston, M.D., Tran, D., Arroyo, B., Sorensen, S., Tay, P., Carruth, J., Coffman, A., Wallace, M., 2012. Automating Mid- and Long-Range Scheduling for NASA’s Deep Space Network. Presented at the SpaceOps 2012, Stockholm, Sweden.

Joslin, D.E., Clements, D.P., 1999. Squeaky Wheel Optimization. J. AI Res. 10, 353–373.

Lewellen, G., Davies, C., Byon, A., Knight, R., Shao, E., Tran, D., Trowbridge, M., 2017. A Hybrid Traveling Salesman Problem - Squeaky Wheel Optimization Planner for Earth Observational Scheduling, in: Proceedings. Presented at the 10th International Workshop on Planning and Scheduling for Space, Pittsburgh, PA.

Shouraboura, C., Johnston, M.D., Tran, D., 2016. Prioritization and Oversubscribed Scheduling for NASA’s Deep Space Network, in: Proceedings. Presented at the ICAPS SPARK Workshop, London.