Appears in Proc. AI in the Oceans and Space Workshop, International Joint Conference on Artificial Intelligence, Melbourne, Australia, August 2017.

AI and Autonomy Initiatives for NASA's Deep Space Network (DSN)

Mark D. Johnston and E. Jay Wyatt
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove, Pasadena CA 91109 USA
{firstname.lastname}@jpl.nasa.gov

Abstract

NASA's Deep Space Network (DSN) consists of thirteen large (34- and 70-meter) antennas that are used to communicate with approximately 40 NASA and partner spacecraft, all at great distance from the earth (generally at Lunar distances and beyond). The DSN has a long history — over 50 years — and has evolved with cutting edge, often custom, telecommunications equipment and associated software systems. In recent years, and in preparation for an increasing future demand, there has been an effort to invest in initiatives that will result in significant cost savings in the future. These efforts are building on, or augmenting, the recent deployment of "Follow-the-Sun" operations (day shift remote operational control of the entire network from each of the three antenna complexes in turn) — which is being deployed in 2017. This paper focuses on *Adaptive Demand Access*: in a paradigm shift from completely pre-planned operations, this concept calls for spacecraft to signal their intent (or not) for near-future contacts, in case they have science results of interest, or have experienced an anomaly. This would take advantage of a beacon tone transmission, which can be detected using smaller antennas. When a connection request is received, the DSN ground systems would adaptively accommodate the request, inserting the contact into the plan as soon as possible, subject to constraints and priorities. The demand access concept incorporates onboard data analysis and science data processing, so that beacon tones can be generated with maximum information. This area is representative of several where infusing AI technologies can lead to improved effectiveness of the DSN as the network readies for support of expanded Mars exploration efforts in the 2020's and beyond.

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. 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).

Figure 1: The DSN 70-meter antenna at the Goldstone, California, Deep Space Communications Complex

All NASA planetary and deep space missions, as well as many international missions, communicate to Earth through the DSN. In some cases, mission's closer to Earth also use the DSN, some routinely, others on an occasional basis. The capabilities of the DSN make it a highly capable scientific facility in its own right, so it is used for radio astronomy (including very long baseline interferometry) as well as radio science investigations. At present, there are about

---

45 regular distinct users of DSN, who together schedule about 500 activities per week on 13 antennas. Over the next few decades, utilization of the DSN is expected to grow significantly, with more missions operating, higher data rates and link complexities, and the possibility of manned mission support. In addition, the total number of antennas will grow to 18 by the mid 2020s, while at the same time there is pressure to reduce ongoing costs yet maintain an around-the-clock operational capability.

Presently, each of the DSN complexes is staffed 24x7 with local personnel who manage the antenna/spacecraft links. The individuals directly responsible for this are designated Link Control Operators, or LCOs. In general, each LCO manages up to two links at a time. Future plans for increased automation are presently in progress, under the general term *"Follow the Sun" Operations (FtSO)* (Johnston et al., 2016)*,* which includes the following two fundamental shifts in operational paradigm:

- *Remote Operations (RO)* — at each complex during their local day shift, each complex will operate not only their local assets, but also all the assets of the other two complexes as well, via remote control.
- *Three Links per Operator (3LPO)* — the number of links a LCO will manage will increase from two (today) to three.

These changes represent a major paradigm shift and will require numerous software changes to improve DSN automation, as well as WAN upgrades to increase bandwidth and reliability of complex-to-complex communications. The benefit will be a significant savings in operations costs while continuing to provide high-quality support to DSN users.

Future DSN Automation Drivers

The major factors affecting future DSN automation are:

- *Cost savings*: operations and development budgets are under constant pressure for reduction, at the same time that aging software and hardware tends to increase cost. A careful program of productivity and efficiency improvements is in place to invest in upgrades that will reduce future costs.
- *More missions*: plans for future Moon, Mars, and asteroid exploration indicate to a larger number of missions, with a corresponding increase in complexity and capability. Current plans call for at least four new Mars missions launching in 2020, joining a fleet of 7 already there. Ka2 band communications will lead to higher data rates but increase sensitivity to weather, and allow for new levels of simultaneous uplink and downlink service.
- *Low cost missions*: not only are more complex missions anticipated, but the largest growth may be in the number of simpler missions in the Smallsat or Cubesat family. These are much cheaper to build and operate from the mission perspective, but present challenges to the DSN in that their communications and navigation requirements may place a substantial load on the network.

At the same time that these factors are increasing the drive to increase automation and reduce costs, missions are expecting no lowering of DSN standards for service and data delivery. Typical contact success rates are in the high 90% range of what is scheduled. For critical events, such as orbit insertions or planetary landings, DSN support and reliability is a determining factor in mission success. This is expected to remain at high levels as the network evolves.

In this paper we will concentrate on one of the initiatives in the advanced development area that shows promise in addressing a number of these drivers. This is a new operations scenario called *Adaptive Demand Access* for reasons that will become clear in the following. We will also briefly discuss other automation initiatives where AI and automation shows the potential to improve DSN operations.

Adaptive Demand Access

The standard operations scenario for most DSN users is dictated by the following factors:

- *light travel time*: planetary exploration invariably required round trip light travel times of minutes, to hours, which means that pre-sequenced command loads must be prepared and uploaded to spacecraft in advance. "Real-time" control is not an option.
- *limited onboard computational resources*: hardened processors take years to reach capabilities that are common on today's workstations, which means that on-board programs enabling higher levels of autonomy are slow to progress to flight TRLs. While there are some examples noted below, in general all operations are fully pre-programmed in detail and constraint checked to the degree possible on the ground.
- *effort to sequence and re-sequence*: related to the point above, there is a great deal of human effort put into creating, optimizing, and checking command sequences before they are uplinked. The cost of a mistake may be loss of mission, and all possible steps are taken to minimize this risk.

The consequence of this is a standard operations model in which DSN allocations are generated months in advance and then baselined, and late changes are hard and expensive to make. The result is a DSN schedule that is essentially frozen about 3-4 months ahead of execution, and changed only in very minor ways by mutual concurrence of all affected missions. Of course, there can be upsets due to asset failure or spacecraft emergencies, but these are relatively rare and the response is to return to the pre-planned schedule as soon as possible.

This standard operations model is illustrated in Figure 2. Missions submit their scheduling requirements about 20

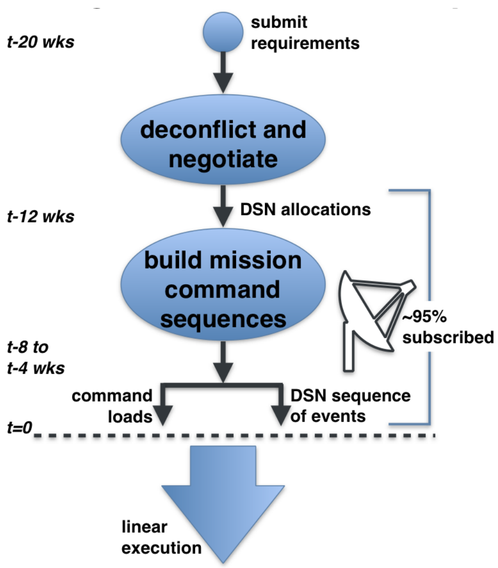

Figure 2. Standard operations scenario for DSN users

weeks before execution, which are deconflicted by the DSN scheduling system, Service Scheduling Software (SSS or S^3)(Johnston et al., 2014), and by a peer-to-peer collaborative negotiation process (Carruth et al., 2010). Baseline DSN allocations are published, and missions then develop their command sequences working around the detailed allocations that have been negotiated. These sequences are prepared for upload to the spacecraft, and are also used to generate sequences of events for DSN ground equipment. Typically, the network is fully booked with fixed activities, to a utilization level of 95% or more. Making any late changes is difficult because affected missions must concur, and must potentially rework their command sequences and spacecraft plans, which is expensive and time consuming.

Demand Access

The notion of demand access is based on a mechanism for spacecraft to signal their intent for a more extensive contact to download telemetry or data. Such signals could come from onboard processes that are monitoring spacecraft health and safety, or from processes that are analyzing collected science data and determining whether interaction with the ground is desirable or necessary.

Demand access is a mode that has long been used on the other NASA networks, in particular the Space Network, which operates the Tracking and Data Relay Satellite System (TDRSS), where it has been facilitated by higher levels of coverage, shared relay resources, and the negligible light travel times between ground and spacecraft (Gitlin et al., 2002). For deep space missions, contacting the ground to inform the network of the need for further communication is the challenging part of a demand access system. The leading candidate concept is so-called ‘beacon mode’ contacts to provide an easy to detect signal that can convey limited information about the need for follow up contacts.

Beacon Mode

Beacon mode (Sherwood et al., 2000; Sue et al., 1997; Wyatt et al., 1998) is a DSN service that provides for the receipt and detection of “tones” generated around the carrier frequency by phase-modulating the RF carrier with a square-wave subcarrier with a 90° modulation angle, where the carrier is completely suppressed. The signal structure is illustrated in Figure 3 for the case of four tones on a single carrier; additional tones can be signaled on different carrier frequencies. Tone center frequencies are separated by 5 kHz.

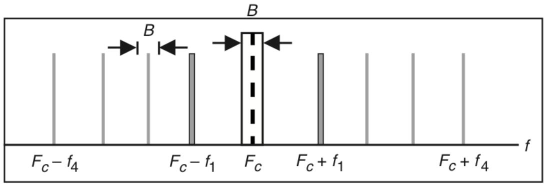

Beacon signal structure

- **B is the frequency uncertainty (2 kHz)**
- **F_c is the carrier frequency**
- **F_i is the i^{th} subcarrier frequency (i = 1…4)**

Figure 3. Frequency structure of a set of beacon tones.

Tones can be used to indicate various meanings, and have generally been encoded to indicate anomaly or interesting situations. The first operational use of beacon tones was on the Deep Space One mission in 1998, where four tones indicated nominal, interesting, important, or urgent; a fifth case, no tone, covers the case where beacon mode is not operating when expected, possibly because the spacecraft is not Earth-pointed, or is experiencing an anomaly that prevents the tone from being generated. The increasing significant of the non-nominal tones indicated the timescale on which a full DSN tracking pass was requested to be scheduled, to follow up.

A second mission on which beacon mode was successfully used was the New Horizons probe to Pluto (Bowman, 2010). New Horizons spent 7.5 years out of its 9.5 year cruise phase to Pluto in a ‘hibernation’ mode in which contact was limited to one 90 minute beacon pass per week, and a monthly telemetry contact. This represented a savings of 80% in the DSN antenna time that would be required when compared with a low-level cruise mode with normal engineering telemetry passes.

It has been proposed that future cubesat missions consider the use of beacon mode along these same lines (Wyatt et al., 2016). While extended cruise phase and anomaly detection and signaling is the most commonly considered use case, the extension to handle opportunistic science presents very interesting possibilities, especially as a sufficiently large number of missions are participating in this operations mode to justify special support for it.

Beacon Mode Operations Scenarios

Consider a set of missions that are equipped to use beacon mode to signal either anomalies or interesting science results that might require some degree of immediate attention. Following the normal DSN process, with a 20 week scheduling horizon, is out of the question, but disrupting the fixed and baselined schedule over the next few days from an event of interest is also highly undesirable.

An alternative is to insert optimized placeholder "reserved" time intervals in the schedule that can be allocated to missions that signal near-term needs for contacts. The process for including these times in the schedule can take into account the probability that any particular mission will need them, as well as the regularly scheduled time to survey for beacon mode passes on one of the smaller antennas. For a given mission set, the optimal set of reserved intervals will minimize the time that missions declare interest in using but ends up not available, and the excess reserved time that ends up unused. For missions in the vicinity of the high-interest planetary exploration targets, there is often a significant amount of viewperiod overlap, so a single reserved interval, if carefully placed, could service one of a set of missions in roughly the same part of the sky.

Such an operations scenario would also lend itself to mixing with the standard operations model where missions require a fixed set of schedule allocations baselined long in advance. Figure 4 illustrates this variant of the standard model: at the 20 week point, fixed mission requirements are submitted as well as statistical expectation rates for missions using beacon mode, along with expected beacon mode and follow-up requirements (e.g. "daily beacon passes with 20% probability of a 4 hour pass within 24 hours, and a 10% probability of a 6 hour pass within 48 hours"). As part of the schedule generation and deconfliction process, reserved time will be included in the schedule that remains unallocated to any specific mission. At the same time, fixed allocations are included that can be used by legacy (deterministic) missions to conduct their normal sequencing processes. As the schedule enters execution, beacon contact missions lay claim to the reserved unallocated time. Such time could also be used for emergency scenarios as well, minimizing the impact on other missions. The net result is less disruption to the primary schedule, while flexibility to accommodate missions that are signaling interesting scientific results.

On Board Autonomous Science

The active use of beacon mode in a demand access context for other than response to anomalies requires that onboard autonomous science processing advance to a degree that the risk of inclusion in mission standard operations is seen as acceptable. This includes missions that are specifically studying transient or infrequently detected phenomena. Onboard science data processing, including classification and re-targeting, has made significant advances. Some examples that are representative of the transient event detection that could make use of a beacon/followup paradigm include the onboard autonomous science operation of EO-1 (Chien et al., 2005) as well as the onboard targeting functionality provided by the AEGIS system on the Opportunity and Curiosity Mars rovers (Castano et al., 2003; Estlin et al., 2007).

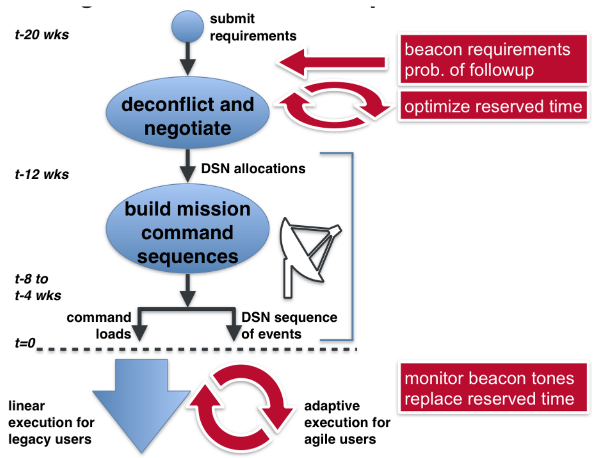

Figure 4. Variation (in red) of the standard process (Figure 3) to incorporate non-deterministic demand access mission requirements.

Other AI Infusions into DSN Operations

In this section, we briefly discuss two other infusion efforts to bring improved automation levels to the DSN.

Complexity-Based Link Assignment

In addition to the increased FtSO workload of managing three times the number of antennas, both local and remote, another driving factor has added to the LCO's demand for attention during link operations: increased link complexity. DSN has long supported Multiple Spacecraft per Antenna (MSPA) operations for missions at Mars and in certain other situations, but until 2016, these were limited to a maximum of *two* simultaneous downlinks, and a single uplink (2-MSPA). In 2016, software upgrades were made to support up to *four* simultaneous downlinks (4-MSPA), and this capability has been heavily used to increase the downlink time available for the Mars mission set. However, 4-MSPA links are significantly more complicated than 2MSPA links, due to the increase in allocated receivers and the impact on the console displays to monitor all the spacecraft activities involved in such links.

To address these concerns, DSN has developed a new software system to analyze and assign links to operators, taking into account workload distribution and link complexity, as well as special operational rules (Tran and Johnston, 2015). The system has a number of unique capabilities: it is distributed, with operating nodes at each of the DSCCs for local operator information and assignment management, and an integrated read-only node at JPL. It aggregates scheduled activities into links, taking into account MSPA groupings as well as other multiple spacecraft activities that should be assigned to a single LCO. Each link has a complexity profile computed based primarily on subsystem usage (downlink, uplink, ranging) with special consideration for multiple antenna activities such as Delta-Differential One-Way Ranging (DDOR). Complexity profiles are based on distinct modeling of the setup, in-track, and teardown phases of each tracking activity. In addition, the system analyzes detailed sequences of events so that the complexity impacts of planetary occultation and data rate changes can also be incorporated into the overall complexity timeline for each link.

The software provides for link operators to be assigned to shifts based on each DSCC's local rules and rosters, and for each operator to be constrained to a maximum number of links supported at one time (nominally this is two), and for a maximum complexity threshold level. Special rules are defined to prevent multiply assigning 2MSPA links to the same operator, and to prevent any other concurrent assignments to operators with 3- and 4-MSPAs. In addition, activities with the highest level of support are assigned by policy to be run locally at each complex, whether in local day shift or not.

The software incorporates a unique assignment search algorithm, per FtSO shift, that finds a set of assignments that do not violate constraints and rules noted above, and which distribute the link assignments as evenly as possible across the available LCOs for all antennas. Assignments can be modified interactively, or "locked" in place and worked around. There is also an intelligent "assistant" mode which scores each possible assignment and checks for constraint violations, then gives the user the opportunity to assign a link to any of the available operators, or to undo any existing assignment. In case conflicts are generated by over-assigning any operator, or violating the rules, the automatic algorithm will search for alternative placements that attempt to resolve them, if possible. Because the system is continuously synchronized with the latest published DSN schedule, for multiple weeks into the future, potential staffing issues can be identified and worked well ahead of time. For example, this could lead to an identified need for overtime personnel, or other changes in how a particularly complicated set of activities can be best managed.

The new software system is being deployed concurrently with the transition to FtSO at each of the DSN complexes.

Complex Event Processing

Complex Event Processing, or CEP, is a technique to bring data streams together representing current, historical, and derived information, subject to real-time analysis, learning, and inference. DSN has been studying the use of CEP in the context of real-time operations for several years (Johnston et al., 2016), and is preparing to roll out a framework in time to support the second phase of Follow-the-Sun Operations.

The CEP infrastructure provides a mechanism to augment DSN's existing capabilities. Data from antenna pointing, frequency, round trip light time predicts, configuration tables, and SOEs are fed into a CEP engine. Streams of realtime data including several hundred monitor data streams and logs are also fed in. Algorithmic rules are run on the data streams and events are output through streams of middleware messages, logs, and TCP streams. Client applications can be connected to these output streams. Operators can then be alerted to conditions that can be watched over by the CEP system.

CEP is still in the early phases of analysis and pre-deployment, but holds great potential for diagnosis and response to anomalous conditions that impact various elements of the DSN.

Conclusions

The DSN Follow-the-Sun paradigm change represents a major change in the way the network is operated. Among the benefits will be lower operations costs, greater resilience, and the ability to support a larger network and user base. Many challenges have been identified in migrating the network from its current state into a future more highly automated configuration. The DSN passed its 50 year anniversary recently, and many parts of the network have legacy roots or residuals that complicate efforts at modernization. However, the initiatives described in this paper are all steps in the direction of improved cost-effective operations, for an increasingly large and complex mission set.

Acknowledgments: 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

Bowman, A.F., 2010. Spacecraft Hibernation: Concept vs. Reality A Mission Operations Manager's Perspective, in: SpaceOps 2010 Conference Delivering on the Dream Hosted by NASA Marshall Space Flight Center and Organized by AIAA. p. 2161.

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.

Castano, R., Anderson, R.C., Estlin, T., DeCoste, D., Fisher, F., Gaines, D., Mazzoni, D., Judd, M., 2003. Rover traverse science for increased mission science return.

Chien, S., Sherwood, R., Tran, D., Cichy, B., Rabideau, G., Castano, R., Davies, A., Mandl, D., Frye, S., Trout, B., Shulman, S., Boyer, D., 2005. Using Autonomy Flight Software to Improve Science Return on Earth Observing One. J. Aerosp. Comput. Inf. Commun. 2, 196–216.

Estlin, T., Gaines, D., Chouinard, C., Castano, R., Bornstein, B., Judd, M., Nesnas, I., Anderson, R., 2007. Increased Mars Rover Autonomy using AI Planning, Scheduling and Execution, in: Proceedings 2007 IEEE International Conference on Robotics and Automation. Presented at the Proceedings 2007 IEEE International Conference on Robotics and Automation, pp. 4911–4918. doi:10.1109/ROBOT.2007.364236

Gitlin, T., Kearns, W., Horne, W.D., 2002. The NASA Space Network Demand Access System (DAS). Presented at the SpaceOps 2002, Houston TX.

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., Levesque, M., Malhotra, S., Tran, D., Verma, R., Zendejas, S., 2016. NASA Deep Space Network: Automation Improvements in the Follow-the-Sun Era. Presented at the IJCAI Workshop on AI in Space, Buenos Aires, Argentina.

Sherwood, R., Wyatt, J., Hotz, H., Schlutsmeyer, A., Sue, M., 2000. Lessons learned during implementation and early operations of the DS1 beacon monitor experiment, in: Reducing the Cost of Spacecraft Ground Systems and Operations. Springer, pp. 385–393.

Sue, M., Peng, T., Wyatt, J., 1997. A New Approach in Spacecraft Monitoring for Efficient Use of the Deep Space Network (Telecommunications and Data Acquisition No. 42–130). JPL.

Tran, D., Johnston, M.D., 2015. Automated Operator Link Assignment Scheduling for NASA's Deep Space Network. Presented at the IWPSS 2015, Buenos Aires, Argentina.

Wyatt, E.J., Abraham, D., Johnston, M.D., Bowman, A., Malphrus, B., 2016. Emerging Techniques for Deep Space Cubesat Operations, in: Proceedings of ICubesat 2016. Presented at the iCubesat 2016, Cambridge, England.

Wyatt, E.J., Hotz, H., Sherwood, R., Szijarto, J., Sue, M., 1998. Beacon monitor operations on the deep space one mission (No. 98-0712). JPL.