Commentary: Automated Scheduling for NASA's Deep Space Network

Mark D. Johnston and Daniel Tran

Commentary by: Bruno Teixeira de Sousa
ESOC /ESA, OPS-OPV (Venus Express Flight Control Team)
Robert Bosch Strasse 5, 64293 Darmstadt, Germany, bruno.sousa@esa,int

Paper Summary

This paper deals with the development and deployment of a new component (S^3) of the DSN network scheduling system, which enhances the capability to automate some of the scheduling tasks, in particular in the area of user requirement interpretation and conflict solving, and provides the users with a collaborative environment to negotiate and sort conflicts.

Problem domain

The DSN scheduling problem boils down to 35 missions "fighting" for time (in average 5 daily hours) in one or more of the available assets (13 antennas in 3 sites), to secure (some times redundant) access to their space platforms, which are constraint by visibility windows and equipment performance. This results in an average network load close to 90%, and a good number of conflicts to deal with and solve.

A more traditional approach whereby missions demand specific activities at specific times results in a higher likelihood of conflict. A mixed solution, whereby users express also equipment preferences or timing constraints, allows for more flexibility in the determination of a more robust and satisfying solution, and minimizes the occurrence of conflict and the number of cases that require actual negotiation, thus making the whole scheduling exercise more efficient and speedy.

Solution

The following steps were taken to introduce the aforementioned component:

- Interpretation of user requests, which are specified via a dedicated language and which includes: specification of desired service configuration; timing constraints (whereby tracks may be allowed to be reduced, extended, split, or time intervals that exclude tracks are specified); track relationships (separation between successive tracks); and prioritization schemes when applicable.
- Introduction of a DSN Scheduling Engine (based on ASPEN framework–see Chien et al, 2000) capable of interpreting the user requests, modelling the DSN domain and finding and resolving scheduling conflicts – the paper deals at length with the details of this module, including its distributed architecture for higher availability of the system, the specific approach to the MPSA scenario (whereby multiple spacecrafts converging in a location in outer space, like Mars, are tracked by a single antenna), the scheduling strategies adopted and the conflict and violation repair methods.
- Integration of the aforementioned component in the existing overall DSN system, including merging of databases.
- Introduction of a web based environment, where users can view, edit and negotiate schedule modifications and conflict repair solution (these last two bullets are deal with in a separate paper, see Carruth et al 2010)

Commentary

I believe it is clear, to the user community at large (of which I include myself) that the times when building a planetary probe meant building a set of antennas for it are clearly over. There are too many satellites out there; it just doesn't make any economical and technical sense to proliferate specific ground equipment, with high maintenance costs and short life spans. A network of re-usable, multi-purpose, long term assets makes sense, but poses an immense management problem. How to allocate what to who and when?

The actual multitude of users makes traditional approaches, whereby we fight for the slot, inefficient and in most cases nerving. It is clear that we need to introduce flexibility and automatisms, and users need to learn to express constraints and preferences, rather than specific wishes. Obviously it should be acknowledged that there are critical situations (launches, planet insertions, landings, etc.) where a specific slot on a specific asset (due to geography/performance) are strictly necessary, but these are typically the exceptional cases, not the routine ones.

The presented system goes absolutely in the right direction. And we at ESA have started already, in the scope of the MEX mission (uses both DSN and ESTRACK interchangeably), interacting directly with it to improve our station utilization profile, by telling the system in advance which periods we do not require tracking and thus reduce the need for stations de-overlapping further down the planning process.

At ESA we have also put in place equivalent mechanisms (EMS – see Damiani et al, 2006) by which missions establish periodical service agreements with the track provider (ESTRACK) to allow flexible conflict resolution and have a forum to view edit and negotiate track changes, but we are still going through the (sometimes painful) process of introducing the paradigm change within the user community and streamlining the associated processes and tools.

I believe we users need to incorporate the request vs. slot principle much earlier at mission operations concept development, to prevent that we design mission profiles that foresee the regular fixed slots on fixed asset as baseline. I'll postulate that when you try to change the paradigm with such missions what you get are requests that are so constraint that they basically indicate a specific slot! But there are obviously specific project drivers that determine the mission's higher or lower flexibility to plan tracks (the orbit phasing and the science opportunity profiles, whether there is a movable antenna, the performance of the radio equipment, etc.), and we all follow the law that the more flexibility we require, the more mass we'll carry, which in turn reflects in cost – the main driver for the mission's approval and respective implementation.

But, as most of us progressively adapt, what we users dearly miss from the network providers is, on one hand, a movement towards more harmonisation of the **interfaces** to the different tracking scheduling systems **across agencies**, and harmonization of the **scheduling processes** in terms of steps and cycles supported and their timings. Missions that are forced, by cooperative political arrangements, to use fixed percentages of this and that network, are in a position that they need to support several planning iterations to come to a balanced and acceptable track profile that fits their needs to access the spacecraft with the service levels agreed at project level. And on the other hand we also miss a harmonisation of the actual catalogue of tracking services provided by different agencies, and the mechanisms whereby we can parameterise them (choosing bit rates; requesting usage profiles such as suppression of uplink, usage of ranging; recording of specific radiometric parameters, etc.).

References

Carruth, J., M. D. Johnston, A. Coffman, M. Wallace, B. Arroyo and S. Malhotra (2010). A Collaborative Scheduling Environment for NASA's Deep Space Network. SpaceOps 2010. Huntsville, AL, AIAA.

Chien, S., G. Rabideau, R. Knight, R. Sherwood, B. Engelhardt, D. Mutz, T. Estlin, B. Smith, F. Fisher, T. Barrett, G. Stebbins and D. Tran (2000). ASPEN - Automating Space Mission Operations using Automated Planning and Scheduling. SpaceOps 2000, Toulouse, France.

Damiani S., Dreihahn H., Noll J., Niezette M., and Calzolari G.P. (2006) - Automated Allocation of ESA Ground Station Network Services. IWPSS 2006, Baltimore, USA.