Forecasting Telecommunications Support Boundaries for Satellite Constellations Using Deterministic Scheduling

David Joesting, CSOC/Honeywell
Gus Larsson, CSOC/Honeywell
7515 Mission Drive
Lanham, MD 20706
E-mail: david.joesting@csoconline.com
gus.larsson@csoconline.com

Abstract

The NPAS is an advanced mission planning tool set utilized by CSOC personnel to forecast space and ground network loading in NASA's telecommunications commitment process for new LEO/NEO customers. This is accomplished by maintaining baseline network user telecommunication models, incorporating new customer support requirements, and generating near operationally valid planning schedules. Subsequent analyses are conducted to assure knowledge of support variations due to changes in customer requirements, available network resources, and inter-satellite orbital phasing. Recent studies have involved constellations of "inexpensive" satellites that collectively gather patterns of information through their planned dispersion. Results to date have been favorable with some recommendations for improvement in storage capacity made. Expectations are for an increase in multi-satellite proposals for both scientific investigation and commercial exploration.

Overview

The Network Planning and Analysis System (NPAS) is a National Aeronautics and Space Administration (NASA) developed resource used in advanced mission planning which forecasts upper and lower limits of expected low earth orbit (LEO) or near earth orbit (NEO) satellite telecommunications tracking support through sensitivity variations in deterministic resource scheduling. Collectively LEO and NEO orbits at their closest approach to the Earth's surface vary from several hundred miles to several Earth diameters in distance.

As a part of the Space Operations Management Office (SOMO) Consolidated Space Operations Contract (CSOC), the NPAS team's evaluations are an integral part of the commitment process for accepting new customers on both NASA's Ground Network (GN) antennas and Space Network (SN) relay antennas through aperture and data flow assessments. Incorporating new telecommunications requirements supplied by the requesting customer into baseline models of committed/potential customers, the team constructs composite simulation models that reflect the expected load on a given network for the support phases in question. Individual customer supports can be defined as multiple services and, if necessary, broken down to the satellite instrument level to allow a complete space-to-ground assessment. Evaluations may cover multiple phases of customer support: launch and early operations, nominal operations, special in-orbit maneuvers, end-of-life, etc., as requested by the customer.

Other than antenna loading/aperture availability, additional constraints are considered as necessary. These include data flow and latency requirements that are a function of ground site data traffic. To reflect the reality of NASA schedule generation, establishment of a new project's priority by the appropriate NASA resource management is often one of the first steps in any analysis. Based on the complexity of the support requirements, this may include a priority scheme that is multi-leveled.

Resultant assessments are used by the appropriate level of NASA/CSOC management in reviewing and approving the Project Service Level Agreement (PSLA) which documents the agreed level of support over the service life of the customer.

Methodology

Due to the number of unknowns, completing a meaningful advanced mission planning evaluation is seldom straightforward. The performing analyst typically must often read beyond the supplied telecommunications requirements and ask for more detailed information concerning support duration ranges, timing or patterns of support, antenna or site priorities, radio frequency incompatibilities, etc. A formal analysis is then performed using analytic simulation models that assist in defining expected bounds of spacecraft tracking support for assumed network configurations. Inadequate network resources, particularly for constrained sets of ground antennas, may necessitate recommendations to search for alternatives, particularly commercial providers, or modification of the required support.

The team's collective experience has shown that regardless of the care taken in planning a network, unforeseeable events or factors will ultimately affect the makeup of the network or modify the currently projected users' requirements. Any attempt to forecast more than eight years for the NASA Space Network or five years for the Ground Network with any certainty are efforts in futility. Nevertheless, planning in distant timeframes should be done by applying the range of possible network resources against extremes in the user spacecraft population so as to identify support boundary situations.

In many cases, the object of a planning analysis is first to define a minimal set of supporting antennas and their earth based location (either on the ground or in geostationary orbit.) The second step is to identify the maximum set of antennas that may be required under a worst-case scenario. Worst-case here can mean several things based on the customer's requirements, but most typically reflects a lack of available resources on the requested primary support antennas due to commitments to higher priority users.

The shortcomings of forecasting the future must be appreciated and reflected in analyses by incorporating possible tracking support bounds as appropriate. For example, given any set of missions with uncontrolled orbits, over time their inter-spacecraft phasing will drift to unknown deviations. The importance of considering this orbital variability, although not linear, will vary directly with the number of supported spacecraft but inversely to the number of the supporting antenna apertures.

For any customer phase to be analyzed, it is important to be able to model each user satellite's telecommunication science/telemetry, command, and tracking support requirements to a degree that allows generation of a near operationally valid planning schedule. In addition, a certain level of support that is close to optimal is highly desirable, although due to total uncertainties it would be an unnecessary overkill if excessive resources are required to attain more than a three to five percent variation from that value.

Baseline requirements and resource models need to be maintained for rapid responses. Experience has shown that customer requirements sometimes change unexpectedly based on either changing engineering needs or science goals. Projects are slow to update written requirements, thus direct contact with knowledgeable user project team members on a periodic basis is important for accurate baseline models. In keeping with the possibility of future changes in support, network sensitivity to possible requirement changes should be evaluated for specific support levels, if not for the customer, then for management insight.

Modeling Parameters and Constraints

User requirements are translated into NPAS model parameters with the intent of identifying constraint situations between network resource users and allowing priority considerations (if any) to resolve such conflicts. Requirements for a single spacecraft may themselves be quite extensive when all tracking, telemetry and control (TT&C) support constraints are considered. The NPAS can model user support down to the individual on-board instrument level and can accurately simulate all up and down link dataflows. Most constraints revolve around the need to satisfy the science and telemetry needs of a spacecraft over any given interval, such as a day. Many spacecraft have the capability to perform command uploading at the same time (and for same or shorter duration) as telemetry downloading. The NPAS model supports complex service assignments of almost any configuration through the use of prototype events (Levine and Joesting 1997).

With regard to individual TT&C supports, many spacecraft have some flexibility in their support duration. For example, they may allow supports as small as 8 minutes but would prefer supports that are 15 minutes long. If shorter supports are all that is available, then the number and proximity of such supports would be higher and closer together, respectively, in order to assure the same total amount of dataflow. Other users are not as flexible in their support duration, and as such, may find larger variations in available network support under different situations.

Other modeling parameters that may be defined are quite extensive. This includes minimum separation requirements between services of the same or different spacecraft, special sunlight or darkness requirements on the spacecraft of its Earth sub-point, user defined ground or political boundaries, spacecraft antenna pointing limitations, and mutual interference between spacecraft.

Special Forecasting Considerations

One important uncertainty to consider in any customer feasibility assessment is the orbital phasing differences between users of the same network resources. Unless a project controls its satellite through powered maneuvers on a regular basis, LEO/NEO orbits will drift over time due to many factors such as solar radiation pressure, atmospheric drag, etc., such that their actual position cannot be estimated more than a month or two into the future. Such changes in phasing between any two, or more, users can affect their ability to be supported on the same resource. Analyses have shown that this effect is more significant for ground based antennas than for space based relay systems due to the shorter view periods and lack of overlapping coverage; however, for low priority SN users, this can still affect the expected bounds of support.

The CSOC team handles the coverage variability uncertainty by using a Monte Carlo simulation with random variations applied to those orbital parameters that are most likely to change. These orbital parameters include the mean anomaly, the argument of perigee, and the right ascension of the ascending node. Each mission is analyzed to determine those parameters that are most likely to vary, and appropriate bounds are established for random variation in the simulation. The total number of trials in any individual Monte Carlo simulation is a function of the number of varied orbital parameters. However, at a minimum, one hundred deterministic planning schedules normally are generated for each phase model in order to establish the range of possible schedule satisfaction and patterns of support sensitivity to the varied parameters. Support commitment analyses should include consequences of worst-case phasing situations and occurrence timing among network users.

The key to generating "valid planning profiles" is a set of scheduling functions that closely match the operational characteristics of the user community. Classic forecasting tools have used modified greedy algorithms to approximate the network load. This does not in itself guarantee a near operationally valid schedule for any user. For example, regardless of network reliability, some users actively manage their support schedules, anticipating a loss of service on any subsequent support. For those users who actively manage the science download of their spacecraft's on-board recorder, it is often the case that they do not wish to have the recorder exceed more than 50% of its capacity at any given time.

The NPAS team has built an integrated collection of scheduling algorithms using heuristics for optimization to allow complex, meaningful planning schedules to be generated for user spacecraft as required, down to the individual instrument level if necessary. Our methods do not guarantee an optimal schedule as our goal is to forecast for each variation modeled a representative support profile that is within 3% of an operational schedule. Over time the tools have produced spacecraft planning schedules that compare favorably to realized support (Stern, Levine, and Pitt 1994).

Due to user community characteristic differences in radio frequency, data format, etc., on NASA's GN, the amount of effective utilization that can be forecast on these resources is significantly less that what the Air Force can achieve on its similar capability RTS Network. A maximum utilization reaching only about 50% (including antenna turn-around time) is all that can be committed to for any given NASA X-/S-band GN antenna due this lack of uniformity. Of course, more load can be added to the sites during the actual schedule period when that time comes in the future. But to analyze a user's requirements at a medium to low priority one or more years in advance requires the effective utilization not to be exceeded if those users are to actually receive their stated support.

Actual loading on the SN is limited by the inherent problem of scheduling and resolving conflicts when users generate their schedules on their own and submit them for inclusion on a user-by-user priority basis. For the resultant actual schedule, conflict resolution is in many cases a manual process with tradeoffs between users made in an unpredictable fashion. The NPAS algorithms have are designed to simulate the individual user's initial requests. In addition when requirements are not met due to priority conflicts, special automatic conflict resolution/priority override options are available to assure that at least minimum levels of support are met. The need to model these overrides is seldom obvious to the modeling analyst at first and typically requires an "iterative build" for any particular new customer's phase model. Not only do schedules generated by NPAS have the ability to accurately forecast network loading, comparisons of individual NPAS "plans" to actual SN schedules have shown similar patterns of support (Simons and Larsson 1994). It must be noted that it is impossible to forecast the actual schedule of any user, such as the Hubble Space Telescope, without knowing the actual targets of observation/science requirements over the period in question. These are unknown even to the project more than several weeks in advance.

Satellite Constellation Studies

Recent studies have been devoted to forecasting the load incurred by constellations of numerous "inexpensive" satellites. These types of projects involve deploying a large number of spacecraft that each consist of minimal hardware and thus limits how much data can be stored on-board. Such a scenario normally means that data will need to be downloaded to a receiving antenna almost whenever one is in range of a particular spacecraft. This can be difficult, however, when the number of satellites in the constellation is sufficiently large as compared to available antenna resources. The objective is to produce a schedule such that minimizes the data loss experienced over all spacecraft in the constellation.

One such proposed constellation is the Nanosat constellation, which can consist of over 100 identical spacecraft, each weighing 10 kg, deployed into highly elliptical orbits around Earth. The apogees range from 12 to 60 Earth radii, and the perigee of each is 3 Earth radii. This results in orbital periods ranging in length from 1 day for the lowest altitude flyer to over 10 days for the spacecraft with the highest apogee. Contact with ground stations can occur only when the spacecraft is within 5 Earth radii of the station.

In the study, each of the spacecraft in the Nanosat constellation was modeled as having 864 megabits of on-board memory into which data is stored at a rate of 1 kilobit per second. This of course means that the recorder will be filled to capacity after 14400 minutes (10 days) if data cannot be downloaded at some interim point. For the outermost flyers, whose orbital periods exceeded 14400 minutes, this was indeed the case, and for those spacecraft we could not avoid losing some amount of data in each orbit.

However, through the use of NPAS playback requests the analysts were able to ensure the full requirements satisfaction of those spacecraft whose fulfillment was not a geometric impossibility. For each spacecraft in the constellation, one generic playback request was created to manage the filling and dumping of the on-board recorder. The requests were prioritized to allow spacecraft with higher altitudes to schedule before those with lower altitudes could.

NPAS playback schedule requests assisted here both in the realistic modeling of actual user equipment but also in managing optimal resource usage for planning purposes. The result was that the data loss was minimized for those spacecraft whose recorder capacity was exceeded between successive station views; the remaining 96 spacecraft in the constellation, through efficient station and spacecraft resource management, were able download 100% of their recorded data. Recommendations to the project were to increase the on-board data storage capacity of the few outer satellites if 100% of science data is required.

A second recent assessment request involved support of "clusters" of identical SN users in the 2001 timeframe. Using the multiple access return (MAR) resources of the SN Tracking Data Relay Satellite System (TDRSS), multiple sets of 5 users each desired to receive from 100 to 50 percent continuous coverage based on their type of operation. Although the TDRSS contains 5 MAR services per satellite (for each of 4 nominal first generation TDRSS satellites in operation), an additional 14 shared MAR services were proposed. Each are currently limited to 5 but can actually support 19 MAR services. This "shared augmentation" of services would be a fairly simple modification of the existing ground control equipment.

Four sets of these clusters of users would reside in a LEO nominally the same as a low inclination Space Shuttle orbit. 2 sets would be 50% users and 2 sets would be 100% MAR users. One set in each group would be "closely" located and one set in each group dispersed in orbit around the Earth. In addition, another 2 sets of 5 each would be dispersed on the ground. These users would have only real-time observation and transmit capability, thus minimizing cost.

This analysis was modeled in a straightforward manner using the NPAS scheduling option that maximizes requested support (also minimizes hand-over.) In addition, the supplemental shared 14 MAR resources were constrained using a special network concurrent shared resources option. Results were very favorable with all users receiving 98% or more of their stated support.

Conflict Explanation Utility

One recent addition to the NPAS toolbox is the Conflict Explanation Utility, referred to simply as the Explainer. It was designed primarily to assist NPAS modeling analysts determine how exactly a new set of requirements affects the performance of existing elements in a model. This can be particularly helpful when the introduction of new requirements has an indirect effect on a seemingly unrelated requirement that already exists in the model.

In its most basic form of operation, the Explainer is designed to assist NPAS analysts determine if a particular schedule request attained its maximum geometric satisfaction. If it happens that the geometric maximum is not reached by the request, the utility helps to determine what may be hindering the scheduling of the request. Many times the schedule degradation is due to higher priority requests that consume resources required by the analyst's input request. Sometimes, however, the poor performance is due to competing requirements within the same request (for example, restricting the spacecraft to schedule at only one station, but that spacecraft never has available coverage at the station). Lastly, the utility also alerts the user to simple conditions which appear to affect the scheduling process, for example, a minimum support time that is too large, or a requested service or antenna type may not be available at a requested station, and so on.

The utility is most helpful when new requirements are added to an existing model. This occurs often and it is not always clear why a new load on one resource affects the satisfaction of a seemingly unrelated requirement, for instance, a mission that is of lower priority but in a different orbit and that requests a different resource. This utility is designed to unravel the chain reaction that occurs when the scheduling of one request bumps another to a different resource, which in turn bumps another request to a still different resource, and so on. For example, it could easily happen that the introduction of a new load at high priority on the TDRSS station at location 174 degrees West longitude affects a medium priority request seeking to schedule a TDRSS station at location 41 degrees West longitude. It can be a time-consuming effort for an analyst to trace back through the scheduling process to determine why this seemingly unrelated scheduling task is affected in this situation. The Explainer simplifies this process by automating the conflict backtracking so that the NPAS analyst can report to the customer with confidence the cause-and-effect relationship seen.

Furthermore, the Explainer can optionally suggest how a request might be modeled differently to better utilize existing resources. This targeted request is not necessarily the request initially input by the analyst (the request that originally could not meet its satisfaction) nor the request for the new set of requirements: it could easily be one of the intermediate requests that played a part in the chain.

An NPAS analyst invokes the Explainer once a schedule run is made and a poorly performing schedule request is identified. Since the scheduling of this input request may have been affected by higher priority requests, the Explainer reads the modeling information pertaining to both the input schedule request and all higher priority requests to build lists of applicable constraints. An example of such a constraint would be if this request was intended to schedule only when the spacecraft was in-sunlight. Then the geometric visibility data for all these requests is processed by the Explainer and schedules with full tolerance (that is, all valid scheduling choices are included in the schedule, even if this results in far more events than desired) are created. Note that "optimal" events are identified during this process, and these are used to determine what the geometric maximum satisfaction for each schedule request should be. The process to this point is what is called the "Geometric State" construction; it describes what could be scheduled under ideal circumstances and also includes a list of alternate selections.

Following this, the "Schedule State" is determined. This begins when all requests of higher priority than the input request, in addition to the input request itself, are fed into the scheduling engine. As events are scheduled, the reason why a particular event has been chosen over another is noted and stored. For a higher-priority event, the reason would most likely be due to normal intra-requirement constraints such as scheduling only when the spacecraft is in-sunlight, selecting the longest pass, etc. If a scheduled event happens not to be one of the optimal events identified during the Geometric State construction, then the system tries to determine whether this is due to constraints imposed by the request upon itself (for example, a minimum separation from an event it previously scheduled) or if it is due to a higher priority request consuming the resources required by the optimal request. In the latter case, a "lookup" is performed to determine what higher priority request events are directly conflicting with the optimal event for the current request. Following this, one link of a dependency chain is built linking the event scheduled by the current request to that scheduled by the higher-priority request, for each interfering higher priority event. Any other reasons why the current event was chosen (for example, spacecraft is in-sunlight, the station was the most desirable) are still noted for the event, but the dependency chain is another part of the explanation why this particular event was chosen.

The Schedule State therefore is the result of the process in which the cause-and-effect relationships are determined for events of all requests that can affect the NPAS analyst's input request. Once this has been completed, The Explainer can begin to determine why the analyst's input request did not attain its maximum satisfaction. There are two cases here: one in which a non-optimal event was scheduled, and one in which no event could be scheduled in a given schedule period. In the former case, the descriptions associated with each event node are simply reported to the user. In the latter, the Explainer follows all links back through the dependency chain for each possible event in the period. (A schedule period here might be one orbit or one day, for example).

Once all non-optimal events and non-scheduled periods are examined for the request, the software determines which intra-request constraints or higher-priority requests had the most significant impact on the analyst's input request. Through the dependency chain backtracking, these impacts could of course be either direct or indirect.

The system can also suggest possible remedies, within given parameters, which can result in the input request attaining a more favorable satisfaction. The parameters allow the analyst to instruct the system to ignore certain facets of the model that should not be modified, like the network architecture or the scheduling of unmovable events. The suggestions take into account knowledge of how the scheduling software operates internally. For instance, the utility might recommend that a station priority list be reordered for some request, or that visibility events should be sorted in AOS/LOS order instead of by station priority.

To summarize, the Explainer is helpful to the NPAS analyst because it provides a full accounting of the schedule process, indicating why a particular event did schedule or why none could be scheduled for a given time span. The analyst can use this information simply as something to report to the customer, or it can be a starting point for further model optimization. The Explainer assists in this regard as well as it can suggest different ways that requests in the model can be modified to maximize the satisfaction of the input request.

Neural Network Application

In response to SN customer representative/management requests for rapid determination of the ability of baseline network configurations to support additional users (such as during a meeting in which the requirements are originally proposed), the NPAS team devised a neural network (NN) application. The underlying neural network application is a NASA COTS product, NETS (Baffes 1989). In constructing a neural network, the NPAS team identifies several types of customer support requirements that may be requested. The executing end user would choose the requirements closest to those desired and the resultant expected mission satisfaction and total network loading (both as percentages) would be given.

To date, we have determined that the best accuracy across the multiple SN baselines is to generate a separate NN for each baselined network model in each model timeframe. The design of the system is such that the core of the system is of minimal size, and the menus and all associated options are dynamically constructed at run time from flat database files. The benefit here is that new scenarios can be created by the end user and easily imported into the utility without having to recompile the software. This also allows for the rapid construction of a new NN if any of the component baseline models are changed.

Due to the inherent CV problem with the GN (covered before), we have not yet perfected the application of NN in this area. There is currently too much flux on the composition of the GN (for example, the recent addition of the commercial DataLynx service for NASA Code Y support) to justify fielding such an application.

Conclusion

Given the trend of higher numbers of satellites on minimal cost networks, proven and adaptable tools are required to provide proper feasibility assessments for advanced mission planning of telecommunications tracking. Providing for the inherent differences in ground-based and space-based antenna systems is necessary, as is the ability to combine such networks when required. A knowledgeable team to analyze input requirements, ask additional questions to bound true problems, appropriately model, and analyze the phases of support is requisite for meaningful customer commitments.

References

Levine, A., and Joesting, D. 1997. Load Forecasting of a Space-based Telecommunications Network: NASA's TDRSS. In workshop notes of the International Workshop on Planning and Scheduling for Space Exploration and Science, Oxnard, CA.

Stern, D., Levine, A., and Pitt, K. 1994. Accuracy Analysis of TDRSS Demand Forecasts. In proceeds of the Third Symposium on Space Mission Operations and Ground Data Systems, 331-318, Greenbelt, MD: NASA Conference Publication 3281.

Simons, M., and Larsson, G. 1994. Analysis of Space Network Loading. In proceedings of the Third Symposium on Space Mission Operations and Ground Data Systems, 1071-1077, Greenbelt, MD: NASA Conference Publication 3281.

Baffes, Paul. NETS User's Guide, Version 2.0. Software Technology Branch, Lyndon Johnson Space Center, September 1989.