Appears in Proceedings of the 1997 IEEE Aerospace Conference, Aspen, CO, February 1997.

This paper is organized in the following manner. We begin by characterizing the current mode of operations of the DSN. Next we describe a high level architecture for automating DSN operations which includes three levels: resource allocation, generation of individual track plans, and execution monitoring of the track plans. Then we begin the main body of the paper which describes the DANS scheduling system. We begin by describing the problem representation within the DANS system. We then describe the priority-based pre-emption scheduling algorithm. Next we provide an example of the algorithm performing rescheduling. Finally, we describe the various reasons why rescheduling occurs and describe related work[1].

2. AUTOMATED SCHEDULING OF DSN RESOURCES - OVERVIEW

Each week, a complex matching process between spacecraft project communication service requests and NASA Deep Space Network (DSN) resources occurs. In this process, project requests and priorities are matched up with available resources in order to meet communications needs for earth-orbiting and deep space spacecraft. This scheduling process involves considerations of thousands of possible tracks, tens of projects, tens of antenna resources and considerations of hundreds of subsystem configurations. Once this initial schedule is produced (8 or more weeks before implementations), it undergoes continual modification due to changing project needs, equipment availability, and weather considerations. Responding to changing context and minimizing disruption while rescheduling is a key issue.

The high level resource allocation problem for the smaller DSN antennas (26M and smaller) is currently handled by the OMP scheduler. OMP accepts generalized service requests from spacecraft projects of the form "we need three 4-hour tracks per week" and resolves conflicts using a priority request scheme to attempt to maximize satisfaction of high priority projects. OMP deals with schedules for NASA's 26-meter subnet involving thousands of possible tracks and a final schedule involving hundreds of tracks.

More recently, there has been a desire to automate the scheduling of the larger (34 and 70 meter) antennas. This paper describes the Demand-based Automated Network Scheduling System (DANS), an automated scheduling system being developed at the Jet Propulsion Laboratory (JPL) to schedule DSN 34 and 70 meter antenna resources. DANS is an evolution of the highly successful OMP system. DANS uses priority-driven, best-first, constraint -based search and iterative optimization techniques to perform priority-based rescheduling in response to changing network demand. In this techniques, DANS first considers the antenna allocation process, as antennas are the central focus of resource contention. After establishing a range of antenna options, DANS then considers allocation of the 5-13 subsystems per track out of the tens of shared subsystems at each antenna complex (signal processing Center (SPC)) used by each track. DANS uses constraint-driven, branch and bound, best first search to efficiently consider the large set of possible subsystems schedules.

The DSN domain contains many resources. In the existing configuration (as of July 1996), it consists of 11 signal processing centers (SPC), 45 antennas, and 161 subsystems. They are located at different sites around the world. The majority of the antennas can be classified as 26, 34, and 70 meter antennas. The 26 meter antennas on average handles 600 activities per week. The 34 and 70 meter antennas performs over 200 activities per week. Additionally, this workload is expected to increased dramatically in the next several years.

3. DOMAIN CHARACTERISTICS

In addition to the basic antenna resource allocation problem, the DSN scheduling problem is further complicated by three factors: (1) context-dependent priority; (2) subsystem allocation; and (3) the possibility of reducing

---

[1] For a further description of DSN antenna operations as an applications are for planning and plan execution see (Chien et al. 1996).

the length of the tracks. DSN track priorities are context dependent in that they are often contingent on the amount of tracking the project has received so far in the week. For example, a project might have priority 3 to get 5 tracks, priority 4 to get 7 tracks and priority 6 to get 9 tracks (where lower priorities represent more important tracks). This reflects that 5 tracks are necessary to maintain spacecraft health and get critical science data to ground stations; 7 tracks will allow a nominal amount of science data to be downlinked; and 9 tracks will allow for downlinking of all science data (e.g., beyond this level additional tracks have little utility). An important point is that specific tracks are not labeled with these priorities (e.g., the project is allowed to submit 5 tracks at priority 3, 2 at priority 4 and so on). Rather when considering adding, deleting, or moving tracks the scheduler must consider the overall priority of the project in the current allocation context.

In addition to allocating antennas, DSN scheduling involves allocating antenna subsystems which are shared by each Signal Processing Center (such as telemetry processors, transmitters and exciters). Allocating these complicates the scheduling problem because it adds to the number of resources being scheduled and certain subsystems may only be required for parts of the track.

Finally, the DSN scheduling problem is complicated by the fact that the track duration can be relaxed. For example, a project may request a 3 hour track but specify a minimum track time of 2 hours. When evaluating potential resource conflicts the scheduler must consider the option of shortening tracks to remove resource conflicts. Currently OMP and DANS use a linear weighting scheme in conjunction with a modified SIMPLEX algorithm to trim tracks in accordance with prioritizations.

DANS accepts two types of inputs: 1, an 8-week prior-to-operation schedule from the Resource Allocation and Planning (RAP) team, and 2, activity requests from each individual flight project. The 8-week schedule is the baseline for creating a conflict-free schedule. Many of the scheduled activities at that time are tentative at best, and therefore subject to revision due to changing project status. Also, the schedule is for the antenna resources only; DSN subsystem scheduling is not considered at all in the 8-week schedule.

The activity requests are used by the flight projects to add and delete activities on an existing schedule due to changing project requirements and/or resource availability. The DANS objective is to satisfy as many activity requests as possible while maintaining a conflict-free status with minimum disruption to the existing schedule. DANS is intended for use by the operation personnel to maintain and update the DSN schedule throughout each schedule.

Another issue is the placement of activities onto the schedule. The possible times for a spacecraft track are limited by spacecraft orbit views, which are the periods in which the spacecraft is visible from a ground station. Also, the range from the antenna to spacecraft dictates the quantity and types of antenna(s) required for each activity. Sometimes, an array of multiple antennas instead of a single one is required to communicate with the spacecraft. In addition, the uplink and downlink activities can occur on different antennas, and can be several hours apart imposing additional dependencies between activities.

There are two types of activities in the DSN domain: spacecraft activities and ground activities. Spacecraft activities are submitted by projects and used to interact with spacecraft. They are required to satisfy the domain constraints above. Ground activities represent hardware maintenance. Antenna time which is not occupied by spacecraft activities is used for ground activities such as non-regular maintenance requirements and testing, with maintenance having higher priority.

Each DSN spacecraft activity is divided into 3 steps: pre-calibration (precal), tracking, and post-calibration (postcal). The time periods for each step are specified in ranges of values. The time periods are unique for each activity type, and depend on the antenna type and subsystem usage. DANS models this dependency by either shrinking or shifting activities to maximize resource utilization as dictated by activity type, antenna type, and subsystem usage.

DANS is required to schedule two different kinds of DSN resources: antennas and subsystems. Antennas and subsystems are unit resources and as such can not be shared by more than one activity. Subsystem resources are hardware such as transmitters which are required to work with an antenna during communication. Normally, there are many pieces of hardware that support each antenna and DANS is required to generate a schedule which allocates all necessary subsystems.

4. RESOURCE REPRESENTATION

The following data structures are used in DANS to represent the DSN domain information, and for supporting the inference process: capacity timeline, project, and pass classes. These data structures are described below.

Capacity Timeline Class

The DSN consists of many Signal Processing Centers (SPC) situated around the globe. Each SPC may contain one or more antennas and many subsystems associated with the SPC. These subsystems can be used in conjunction with the antennas to perform tracks. DANS represents these resources in a hierarchical manner as shown in Figure 1. At the top of the hierarchical tree is a SPC resource. Each SPC contains one or more antennas, which are the children of the SPC. There are also many SPC subsystem resources residing also as children of the SPC. For each antenna in the hierarchy tree, there are many DSS subsystems associated with it. To further illustrate this hierarchy, Figure 2 shows the hierarchy for the SPC-10 DSN complex.

All the DSN resources are represented by the capacity timeline class (CptyTimeline). The CptyTimeline models a resource's usage at any instant of time for the duration of a schedule. It is composed of one or more instances of the capacity unit class (CptyUnit). A CptyUnit is used to represent a constant resource usage within a time period. Shown in Figure 3 is the timeline representation of the DSS-26 antenna resource.

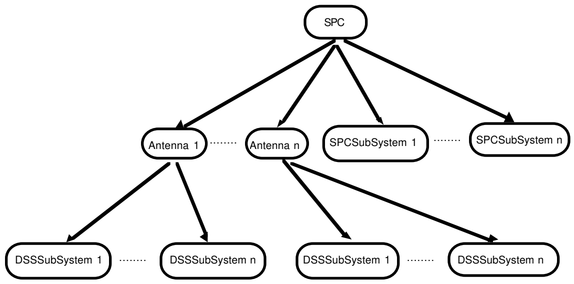

Figure 1: DSN Resources Hierarchy

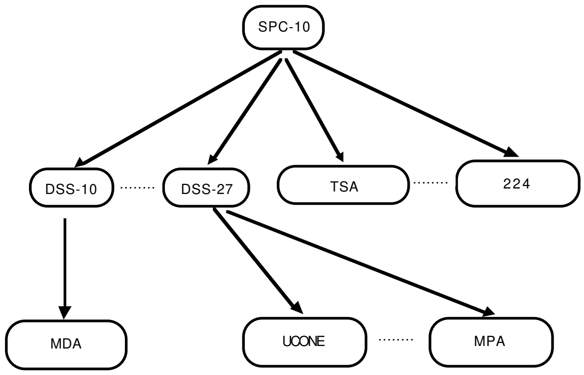

Figure 2: SPC-10 Resource Hierarchy Example

Since the CptyTimeline represents the state of a resource for the schedule's duration, it is the most often used data structure in the system. The CptyTimeline is constantly being modified and updated during the inference process to reflect the state of the resource at that instant of time. In order to increase the system performance, a time slice caching scheme is used to expedite the query process. The scheme equally divides the timeline into a number of buckets. During a query, the system will have to find the bucket that contains the moment first, then it will search sequentially within that bucket to match the query. For example, when the system looks for a CptyUnit which contains the 1:30 moment as shown in Figure 3, it first identifies Bucket #2 as the container which includes the 1:30 moment. Then it traverses down the timeline starting from the beginning of Bucket #2 until it locates the CptyUnit that contains the 1:30 moment.