Automated Planning for a Deep Space Communications Station

Tara Estlin, Forest Fisher, Darren Mutz and Steve Chien
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109-8099
{firstname.lastname}@jpl.nasa.gov

Abstract

This paper describes the application of Artificial Intelligence planning techniques to the problem of antenna track plan generation for a NASA Deep Space Communications Station. The described system enables an antenna communications station to automatically respond to a set of tracking goals by correctly configuring the appropriate hardware and software to provide the requested communication services. To perform this task, the Automated Scheduling and Planning Environment (ASPEN) has been applied to automatically produce antenna tracking plans that are tailored to support a set of input goals. In this paper, we describe the antenna automation problem, the ASPEN planning and scheduling system, how ASPEN is used to generate antenna track plans, the results of several technology demonstrations, and future work utilizing dynamic planning technology.

1 Introduction

The Deep Space Network (DSN) (JPL 1994) was established in 1958 and since then it has evolved into the largest and most sensitive scientific telecommunications and radio navigation network in the world. The purpose of the DSN is to support unpiloted interplanetary spacecraft missions and support radio and radar astronomy observations in the exploration of the solar system and the universe. There are three deep space communications complexes, located in Canberra, Australia, Madrid, Spain, and Goldstone, California. Each DSN complex operates four deep space stations – one 70-meter antenna, two 34-meter antennas, and one 26-meter antenna. The functions of the DSN are to receive telemetry signals from spacecraft, transmit commands that control the spacecraft operating modes, generate the radio navigation data used to locate and guide the spacecraft to its destination, and acquire flight radio science, radio and radar astronomy, very long baseline interferometry, and geodynamics measurements.

From its inception the DSN has been driven by the need to create increasingly more sensitive telecommunications devices and better techniques for navigation. The operation of the DSN communications complexes requires a high level of manual interaction with the devices in the communications link with the spacecraft. In more recent times NASA has added some new drivers to the development of the DSN: (1) reduce the cost of operating the DSN, (2) improve the operability, reliability, and maintainability of the DSN, and (3) prepare for a new era of space exploration with the New Millennium program: support small, intelligent spacecraft requiring very few mission operations personnel (R. Hill *et al.* 1995).

This paper addresses the problem of automated track plan generation for the DSN, i.e. automatically determining the necessary actions to set up a communications link between a deep space antenna and a spacecraft. Similar to many planning problems, track plan generation involves elements such as subgoaling to achieve preconditions and decomposing high-level (abstract) actions into more detailed subactions. However, unlike most classical planning problems, the problem of track generation is complicated by the need to reason about issues such as metric time, DSN resources and equipment states. To address this problem, we have applied the Automated Scheduling and Planning Engine (ASPEN) to generate antenna track plans on demand.

ASPEN (Fukanaga *et al.* 1997) is a generic planning and scheduling system being developed at JPL that has been successfully applied to problems in both spacecraft commanding and maintenance scheduling and is now being adapted to generate antenna track plans. ASPEN utilizes techniques from Artificial Intelligence planning and scheduling to automatically generate the necessary antenna command sequence based on input goals. This sequence is produced by utilizing an "iterative repair" algorithm (Minton & Johnston 1988; Zweben *et al.* 1994), which classifies conflicts and resolves them each individually by performing one or more plan modifications. This system has been adapted to input antenna tracking goals and automatically produce the required command sequence to set up the requested communications link.

This work is one element of a far-reaching effort to upgrade and automate DSN operations. The ASPEN Track Plan Generator has been demonstrated in support of the Deep Space Terminal (DS-T), which is a prototype 34-meter deep space communications station intended to be capable of fully autonomous operations (Fisher *et al.* 1998; 1999).

This rest of this paper is organized in the following manner. We begin by characterizing the current mode of operations for the DSN, and then describe the track plan generation problem. Next, we introduce the ASPEN planning and scheduling system and describe it's modeling language and search algorithm(s). We then present an operations example of using this system for track plan generation and discuss several successful demonstrations that were performed with Mars Global Surveyor using a 34-meter antenna station in Goldstone, CA. Finally, we discuss some related work and describe current efforts to expand this system to incorporate a dynamic planning approach which will allow for closed-loop control and automatic error recovery when executing a DSN antenna track.

2 How the DSN Operates

The DSN track process occurs daily for dozens of different NASA spacecraft and projects which use the DSN to capture spacecraft data. Though the process of sending signals from a spacecraft to Earth is conceptually simple, in reality there are many earthside challenges that must be addressed before a spacecraft's signal is acquired and successfully transformed into useful information. In the remainder of this section, we outline some of the steps involved in providing tracking services and in particular discuss the problem of track plan generation.

The first step in performing a DSN track is called network preparation. Here, a project sends a request for the DSN to track a spacecraft involving specific tracking services (e.g. downlink, uplink). The DSN responds to the request by attempting to schedule the necessary resources (i.e. an antenna and other shared equipment) needed for the track. Once an equipment schedule and other necessary information has been determined, the next step is the data capture process, which is performed by operations personnel at the deep space station. During this process, operators determine the correct steps to perform the following tasks: configure the equipment for the track, perform the actual establishment of the communications link, and then perform the actual track by issuing control commands to the various subsystems comprising the link. This process can be viewed as a problem of software module reconfiguration (Chien *et al.* 1998) where a set of activities is reconfigured (e.g. activities are swapped in and out, sequenced, parameterized) in order to provide control scripts for antenna commanding and control.

After a track has been generated it is then executed at the station. Throughout the track operators continually monitor the status of the link and handle exceptions (e.g. the receiver break lock with the spacecraft) as they occur. All of these actions are currently performed by human operators, who manually issues tens or hundreds of commands via a computer keyboard to the link subsystems. This paper discusses the application of the ASPEN planning system to automatically generate DSN track plans (i.e. the steps necessary to setup and perform the requested track) and dramatically reduce the need for many manual steps.

3 Track Plan Generation: The Problem

Generating an antenna track plan involves taking a general service request (such as telemetry - the downlink of data from a spacecraft), an antenna knowledgebase (which provides the information on the requirements of antenna operation actions), and other project specific information (such as the spacecraft sequence of events), and then generating a partially-ordered sequence of commands. This command sequence will properly configure a communications link that enables the appropriate interaction with the spacecraft. To automate this task, the ASPEN planning and scheduling system has been applied to generate antenna operation procedures on demand.

ASPEN has been adapted to use high-level antenna track information to determine the appropriate steps, parameters on these steps and ordering constraints on these steps that will achieve the input track goals. In generating the antenna track plan, the planner uses information from several sources (see Figure 1):

Project Service Request - The service request specifies the DSN services (e.g. downlink, uplink) requested by the project and corresponds to the goals or purpose of the track.

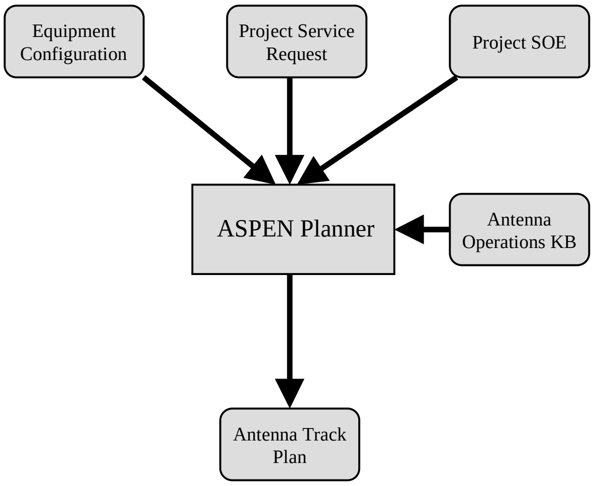

Figure 1: ASPEN Inputs and Outputs

Project SOE -. The project sequence of events (SOE) details spacecraft events occurring during the track - including the timing of the beginning and ending of the track and spacecraft data transmission bit rate changes, modulation index changes, and carrier and subcarrier frequency changes.

Antenna Operations KB - The Antenna Operations Knowledge Base stores information on available antenna operations actions/commands. This KB dictates how actions can be combined to provide essential communication services. Specifically, this includes information such as action preconditions, postconditions, and command directives and also includes any other relevant information such as resource and state descriptions

Equipment Configuration - This configuration details the types of equipment available and includes items such as the antenna, antenna controller, the receiver, etc.

4 The ASPEN Modeling Language and Search Algorithm

ASPEN is a reusable, configurable, generic planning/ scheduling application framework that can be tailored to specific domains to create conflict-free plans or schedules. It's components include:

- An expressive modeling language to allow the user to naturally define the application domain
- A constraint management system for representing and maintaining antenna and/or spacecraft operability and resource constraints, as well as activity requirements
- A set of search strategies
- A temporal reasoning system for expressing and maintaining temporal constraints
- A graphical interface for visualizing plans/schedules

A brief introduction into the ASPEN modeling language is given below. For more details on ASPEN, see (Fukanaga *et al.* 1997).

4.1 Modeling Language

The ASPEN modeling language allows the user to define activities, resources, and states which describe a particular application domain. A domain model is input at start-up time, so modifications can be made to the model without requiring ASPEN to be recompiled. The modeling language has a simple syntax, which can easily be used by operations personnel. Each application model is comprised of several files which define and instantiate activities, resources, and states.

The central data structure in ASPEN is an activity. An activity corresponds to the act of performing a certain function (e.g. configuring the antenna receiver) and represents an action or step in a plan/schedule. Once instantiated it has a start time, an end time, and duration. Activities can also use one or more resources and reason about domain states. Figure 2 shows several activity definitions utilized for antenna-track plan generation. Shown is a "Pre_track" activity that introduces into the plan the steps required to set up the antenna and subsystems for the actual track, and an "Acquire_signal" activity that uses the antenna receiver to acquire the spacecraft signal.

Activity parameters are used to store values in activities or reservations. Lines 8 and 9 contain parameters that specify the number of communication channels (or ways) utilized in the track and the time the track began. Parameter values can be set in an activity definition, passed in from other activities, or as in this case, determined by checking the value of a particular state (as shown on lines 10 and 11). These parameter values are then later referenced when generating the actual command that will execute this step in the final antenna track plan.

Activities can also contain decompositions, as shown in the first activity definition in Figure 2. This

```
1    Activity Pre_track {
2       Decompositions =
3           (Begin_pre_track, Configure_subsystems, Point_antenna, On_point_check,
4           Start_APC where ordered)
5    };
6
7    Activity Acquire_signal {
8       int way;
9       time_param bot_time;
10      Timeline_dependencies =
11          bot_time ← bot_time_sv, way ← way_sv;
12      reservations =
13          BVR,
14          Antenna_sv must_be "on_point",
15          Signal_sv change_to "acquired";
16   };
```

Figure 2: ASPEN Activity Examples

activity contains a decomposition into several subactivities (e.g. Configure_subsystems, Point_antenna). These subactivities are activities that can be scheduled any time within the parent activity time interval subject to any constraints within the subactivity definitions. Thus as soon as a "Pre_track" activity is instantiated in a plan, it's subactivities are also instantiated. Decompositions may also be "ordered", such as the one shown here, where all sub-activities must occur in the order specified.

Reservations are used to reserve a portion of a resource or state for the duration of an activity. The second activity in Figure 2 contains a reservation on the Block-V Receiver (BVR). There are two main types of resource reservations in ASPEN: atomic and aggregate. Line 13 of Figure 2 shows an example of an atomic reservation that reserves the BVR for the duration of the activity. No other activities can use the BVR during this time. An example of an aggregate reservation would be to use N units of power or fuel or some other depletable resource.

State reservations can be used to require a certain state be true or change the value of a state variable. Line 14 of Figure 2 requires that the antenna be "on_point" (indicating that the antenna is pointing at the correct set of coordinates) before attempting to acquire the spacecraft signal. Line 15 changes the state of the signal state variable to "acquired" indicating that the spacecraft signal has been successfully acquired by the receiver.

One other utilized feature that is not shown is temporal constraints between activities. Examples of these constraints are: starts_before, starts_after, contains, etc. These constraints can be used to specify partial orderings over certain activities. For example, in the antenna model, it's specified that the activity for generating receiver predicts (where predicts dictate settings for the receiver) must be ordered before the activity which delivers the predicts to the receiver (e.g. Generate_bvr_predicts ends_before start of Deliver_bvr_predicts).

Besides activities, other defined model elements include resources and states. Resource definitions contain a profile of a physical resource over time. There are three main types of resources: atomic, depletable, and non-depletable. Atomic resources are physical devices that can only be used (reserved) by one activity at a time, such as a receiver or antenna controller. Depletable resources are resources that can be used by more than one activity at a time, but their capability is diminished after use, such as a battery or other power source. Non-depletable resources are similar to depletable resources except that their capacity does not diminish and thus they do not need to be replenished, such as memory bus. Most of the resources utilized for antenna track plan generation are atomic resources that represent different pieces of equipment.

A device or subsystem may also be represented by a state variable that gives information about its state over time. A state variable contains a state profile, which is defined as an enumerated type. Some examples of possible states are that an antenna can be "on_point", "off_point" or "stowed", a receiver can be "locked" or "unlocked" and the Conscan subsystem can be "on" or "off." States can be reserved or changed by activities and a state variable must equal some state at every time. Also, if there are several different states possible for a particular state variable, allowable state transitions can be defined where only certain transitions between those states are possible.

4.2 Conflict Detection

Conflicts arise within a plan when a constraint has been violated. This constraint could be temporal or involve a parameter, resource or state. In order to reason about temporal constraints, ASPEN utilizes a Temporal Constraint Network (TCN) that describes temporal relationships between activities. The TCN can be queried as to whether the temporal constraints currently imposed between activities are consistent. Also used is a Parameter Dependency Network (PDN) that reflects any defined dependencies between activity parameters. A dependency between two parameters is defined as a function from one parameter to another. These dependencies are maintained by the PDN which checks that at any given time all dependency relations are satisfied.

Resource timelines are used to reason about the usage of physical resources by activities. Conflicts are detected if two or more activities are utilizing an atomic resource at the same time or if the aggregate usage of a resource exceeds its capacity at any given time. State timelines represent attributes, or states, that can change over time where each state can have several possible values. As activities are placed/moved in time, the state timeline updates the values of the state, and detects possible inconsistencies or conflicts that can be introduced as a result.

4.3 Planning/Scheduling Algorithms

The search algorithm(s) utilized in a planning/ scheduling system typically search for a valid, possibly near-optimal plan/schedule. The ASPEN framework has the flexibility to support a wide-range of scheduling algorithms. For this application, we mainly utilized a repair-based algorithm (Fukanaga *et al.* 1997; Minton & Johnston 1988; Zweben *et al.* 1994) to classify conflicts and resolve them. For track plan generation, ASPEN begins by generating a complete schedule that's possibly invalid using a greedy, constructive algorithm. Then at every iteration, the schedule is analyzed, and repair heuristics that attempt to eliminate conflicts in the schedule are iteratively applied until a valid schedule is found. Conflicts are resolved by performing one or more plan modifications such as adding, moving or deleting an activity. Domain-dependent heuristics can be also added to direct the search towards more optimal solutions.

5 Track Plan Generation: An Example

Given a set of tracking requests, ASPEN can generate a conflict-free track plan within the order of seconds that will correctly set up the requested communications link. In order to begin the planning process, the tracking service request, the equipment configuration, and the project SOE are parsed and relevant information is placed in a initial setup file which lists the requested track goals and any relevant initial state information. For example, Figure 3 shows three activity instantiations that request that a "Pre_track", "Track" and "Post_Track" activity be placed in the final plan at specific times.

```
Pre_track pre_track1 {
  Start_time = 1998-213/13:32:26;
  End_time = 1998-213/13:47:26;
};

Track Track1 {
  Start_time = 1998-213/13:47:26;
  End_time = 1998-213/16:40:00;
};

Post_track post_track1 {
  Start_time = 1998-213/16:40:00;
  End_time = 1998-213/16:50:00;
};
```

Figure 3: Activity Instantiations

ASPEN then decomposes these activities into the necessary steps that set up the antenna and subsystems (i.e. "Pre_track"), that perform the track (i.e. "Track"), and that perform the necessary shutdown procedures once the track had ended (i.e. "Post_track"). Other initial state information is provided in a "Set_state_values" activity which sets up the appropriate state variables. The information includes the spacecraft ID, antenna ID, the tracking goals, the carrier and sub-carrier frequency, the symbol rate, etc. ASPEN is also provided with the model files that hold the relevant activity, parameter, resource and state definitions, which were explained in the previous section.

Once the initial goals and state information are loaded, ASPEN utilizes its iterative repair algorithm to create a conflict-free track plan that provides the requested services. This final plan contains a large amount of information, including a list of grounded activities (where each activity has been assigned a start time and end time), and a list of constraints over those activities, including temporal, parameter,

```
Config_equip:

Start jsc_asn.prc(dss,sc,pass,&ret_status)
If (!ret_status) then
   Write("fatal error:  cannot start pass")
   Goto fatal_err
Endif

Start ugc_hi.prc(&ret_status)
If (!ret_status) then
   Write("fatal error:  cannot control UGC")
   Goto fatal_err
Endif

Start apc_hi.prc(&ret_status)
If (!ret_status) then
   Write("fatal error:  cannot control APC")
   Goto fatal_err
Endif
...
Point_antenna:

Ret_status = exec("APC DCOS")
Start apc_track.prc(&ret_status)
If (!ret_status) then
   Write("fatal error:  cannot point ant")
   Goto fatal_err
Endif
```

Figure 4: Antenna Control Script

resource and state constraints. ASPEN also displays the final resource and state timelines which show the states of those entities over the course of the plan.

The actual antenna control script that will be used to execute the track is output in a separate file which contains the command sequence necessary to control and breakdown the link. In the model definition, a command (or set of commands) can be specified for each defined activity. These commands are then output in the correct sequence based on the final plan constraints. An example of this file format is shown in Figure 4. This control script is then sent to an antenna operator or execution agent where it will be used to perform the requested track.

Thus, through the use of AI planning and scheduling techniques we are able to reason about the constraints and interactions of track plan activities in order to sequence and configure the final track. This task can be viewed as a complex software module reconfiguration problem (Chien *et al.* 1998) where software modules (i.e. track plan activities) are selected and reconfigured in order to generate the necessary antenna command sequence.

6 DS-T Demonstrations

In order to validate ASPEN's ability to create a valid track plan, this system has been demonstrated in support of the Deep Space Terminal (DS-T) (Fisher *et al.* 1998; 1999) being developed at the NASA Jet Propulsion Laboratory. DS-T is a prototype 34-meter deep space communications station intended to be capable of fully autonomous operations. When requested to perform a track, the DS-T station automatically performs a number of tasks (at appropriate times) required to execute the track. First, a Schedule Executive sets up the track schedule for execution and provides the means for automated rescheduling and/or manual schedule editing in the event of changes. The Configuration Engine is then responsible for retrieving all the necessary data needed for station operations. Next, the Script Generator (ASPEN) generates the necessary command sequence to perform the track. Finally, a Station Monitor and Control process executes the generated script and records relevant monitor data generated during the track.

Demonstrations of the DS-T architecture were performed in April, May and September 1998 where ASPEN was used to automatically generate the necessary command sequences for a series of Mars Global Surveyor (MGS) downlink tracks using the equipment configuration at Deep Space Station 26 (DSS26), a 34-meter antenna located in Goldstone, CA. These command sequences were produced and executed in a fully autonomous fashion with no human intervention. In addition, the September demonstration was for a 6 day period where DS-T was used to perform all Mars Global Surveyor coverage scheduled for the Goldstone antenna complex. This corresponded to roughly 13 hours of continuous track coverage per day.

Figure 5 presents statistical data from a representative day of this 6-day demonstration. The graph represents when MGS was in view of the stations at each of the three DSN complexes (Madrid, Goldstone, and Canberra). For this day, DST collected above 90% of the transmitted data frames which is on par with the operator-controlled stations. DS-T, which is located at Goldstone, tracked MGS through the five track segments indicated in the figure. Each track segment is labeled with a mode. This mode indicates whether information is just being downlinked from a spacecraft to a station (1 way), information is being both uplinked and downlinked to a station (2 way), or information is being uplinked to one station and downlinked to different station (3 way).

For each track segment, Figure 5 shows the percentage of frames successfully collected by DS-T. During

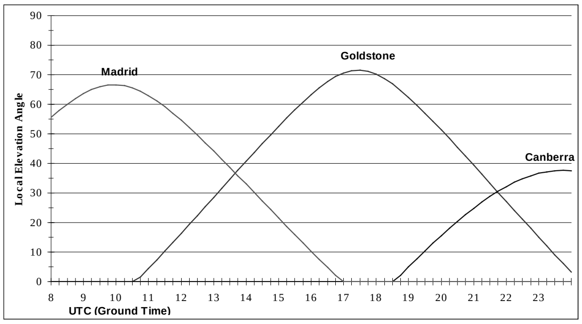

Track Segments:				Frames Collected:
#1	11:15:00 - 14:00:23	3way/65	(Madrid)	75%
#2	14:00:23 - 15:42:26	LOS		0%
#3	15:42:26 - 15:52:23	1way		91%
#4	15:42:26 - 15:52:23	3way/25	(Goldstone)	96%
#5	15:52:23 - 20:03:34	1way		90%
#6	20:03:34 - 23:40:00	3way/34	(Canberra)	23%

Figure 5: September 16, 1998 MGS track

segment 1 and 2 the elevation of the dish is low in the sky, as shown by the graph. Under these circumstances there is considerably more atmospheric interference which explains the lower percent of frame collection for these segments. Conversely, for segment 4, which is a long segment with the spacecraft high in the sky, the data collection percentage is quite high. In segments 3 and 5 data collection is slightly lower due to data being lost during a change in mode (e.g. 3way to 1way). The LOS label on track segment 2, indicates that there was a scheduled "loss of signal" (LOS) during that time and thus no frames were collected.

As a component of the DS-T demonstrations, the Script Generator (ASPEN) performed flawlessly, dynamically producing instantiated control scripts based on the desired service goals for the requested communications pass. Overall, the use of this technology resulted in three primary benefits:

- Enabling autonomous operations by eliminating the need for hundreds of manual inputs. Currently, the task of creating a communications link is a very manual and time-consuming process which requires input of over a hundred control directives and the constant monitoring of several dozen displays.
- Reducing the level of expertise required of an operator to perform a communications track. Currently, this process requires a very high level of expertise from the operator, however through the development of a KB by a domain expert much of this expertise is captured within the system.
- Providing a declarative representation of operations procedures. The KB developed for this task documents the procedural steps required for performing antenna communication services.

7 Related Work

There are a number of existing systems built to solve real-world planning or scheduling problems (Tate, Drabble, & Kirby 1994; Wilkins 1994; Zweben *et al.* 1994). The problem of track plan generation combines elements from both these fields and thus traditional planners and schedulers cannot be directly applied. First, many classical planning elements must be addressed in this application such as subgoaling to achieve activity preconditions (e.g. the antenna must be "on_point" to lock up the receiver) and decomposing higher-level (abstract) activities into more detailed sub-activities. In addition, many scheduling elements are presents such as handling metric time and temporal constraints, and representing and reasoning about resources (e.g. receiver, antenna controller) and states (e.g. antenna position, subcarrier frequency, etc.) over time.

One other system has been designed to generate antenna track plans, the Deep Space Network Antenna Operations Planner (DPLAN) (Chien *et al.* 1997). DPLAN utilizes a combination of AI hierarchical-task network (HTN) and operator-based planning techniques. Unlike DPLAN, ASPEN has a temporal reasoning system for expressing and maintaining temporal constraints and also has the capability for representing and reasoning about different types of resources and states. ASPEN can utilize different search algorithms such as constructive and repair-based algorithms, where DPLAN uses a standard best-first based search. And, as described in the next section, ASPEN is currently being extended to perform dynamic planning for closed-loop error recovery, where DPLAN has only limited replanning capabilities.

8 Future Work: Closed-Loop Control through Dynamic Planning

Currently, we are working on extending the current ASPEN Track Plan Generator to provide a Closed Loop Error Recovery system (CLEaR) for DSN track automation. CLEaR is a real-time planning system where the approach taken is to dynamically feed monitor data (sensor updates) back into the planning system as state updates. As these dynamic updates come in, the planning system verifies the validity of the current plan. If a violation is found in the plan, the system will perform local modification to construct a new valid plan. Through this continual planning approach, the plan is disrupted as little as possible and the system is much more responsive and reactive to changes in the real (dynamic) world.

This CLEaR effort is also being integrated with a Fault Detection, Isolation and Recovery (FDIR) system. FDIR is an expert system providing monitor data analysis. As is often the case with system, monitor (sensor) data is often related in different ways that becomes difficult for a human to detect. The advantage of combining these two systems is that FDIR can first interpret the vast amount of data and summarize it into a set of meaningful values for a planning system to react to. We think of this union as intelligent analysis and intelligent response, much like a careful design and implementation; one without the other is of little use.

9 Conclusions

This paper has described an application of the ASPEN automated planning system for antenna track plan generation. ASPEN utilizes a knowledge base of information on tracking activity requirements and a combination of Artificial Intelligence planning and scheduling techniques to generate antenna track plans that will correctly setup a communications link with spacecraft. We also described several demonstrations that have been performed as part of the DS-T architecture where ASPEN was used to generate plans for downlink tracks with Mars Global Surveyor. Finally, we described a planned extension of this system which will allow for closed-loop error recovery and fault detection using dynamic planning techniques.

Acknowledgments

This work was performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. We thank members of the ASPEN scheduling team and members of the DS-T automation team for contributing to this work.

References

Chien, S.; Jr., R. H.; Govindjee, A.; Wang, X.; Estlin, T.; Griesel, A.; Lam, R.; and Fayyad, K. 1997. A hierarchical architecture for resource allocation, plan execution, and revision for operation of a network of communication antennas. In *Proceedings of the 1997 IEEE Conference on Robotics and Automation.*

Chien, S.; Fisher, F.; Mortensen, H.; Lo, E.; Greeley, R.; Govindjee, A.; Estlin, T.; and Wang, X. 1998. Using artificial intelligence planning techniques to automatically reconfigure software modules. In *Proceedings of the Eleventh International Conference of Industrial and Engineering Application of Artificial Intelligence and Expert Systems.*

Fisher, F.; Chien, S.; Paal, L.; Law, E.; Golshan, N.; and Stockett, M. 1998. An automated deep space communications station. In *Proceedings of the 1998 IEEE Aerospace Conference.*

Fisher, F.; Mutz, D.; Estlin, T.; Paal, L.; and Chien, S. 1999. The past, present and future of ground station automation within the dsn. In *Proceedings of the 1999 IEEE Aerospace Conference.*

Fukanaga, A.; Rabideau, G.; Chien, S.; and Yan, D. 1997. Towards an application framework for automated planning and scheduling. In *Proceedings of the 1997 International Symposium on Artficial Intelligence, Robotics and Automation for Space.*

JPL. 1994. Deep Space Network. Technical Report 400-517, Jet Propulsion Laboratory.

Minton, S., and Johnston, M. 1988. Minimizing conflicts: A heuristic repair method for constraint satisfaction and scheduling problems. *Artificial Intelligence* 58:161–205.

R. Hill, J.; Chien, S.; Fayyad, K.; Smyth, C.; Santos, T.; and Bevan, R. 1995. Sequence of events drive automation of the deep space network. *Telecommunications and Data Acquisition* 42-124.

Tate, A.; Drabble, B.; and Kirby, R. 1994. O-Plan2: An open architecture for command planning and control. In Zweben, M., and Fox, M., eds., *Intelligent Scheduling.* San Francisco, CA: Morgan Kaufmann. 213–240.

Wilkins, D. 1994. *Practial Planning: Extending the Classical AI Paradigm.* San Francisco, CA: Morgan Kaufmann.

Zweben, M.; Daun, B.; Davis, E.; and Deale, M. 1994. Scheduling and rescheduling with iterative repair. In Zweben, M., and Fox, M., eds., *Intelligent Scheduling.* San Francisco, CA: Morgan Kaufmann. 241–256.