Automating Operations for a Network of Communications Antennas

Randall W. Hill, Jr., Steve A. Chien, and Kristina V. Fayyad
Jet Propulsion Laboratory, California Institute of Technology

Keywords: Artificial Intelligence, Planning and Scheduling, Architecture

Introduction

The Deep Space Network (DSN) [1] 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 require 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.

The purpose of this paper is to describe an architecture for automating the process of capturing spacecraft data. In particular, we will describe how the components of the architecture transform a flight project service request into an executable set of DSN operations that fulfill the request through automated resource allocation, goal-driven plan generation, and plan execution and monitoring. The architecture was designed to assist the DSN in meeting the three NASA goals mentioned in the last paragraph. We successfully demonstrated a prototype of this architecture in February 1995 at NASA's experimental DSN station, DSS-13, on a series of Voyager tracks.

This paper is organized in the following manner. We begin by characterizing the operation of the DSN at the time that this research was performed. Next we describe the architecture of our automated prototype -- we give a functional description of each of the components, which includes the OMP scheduling system for automated resource allocation, DPLAN, an automated procedure generation system, and LMCOA, a plan execution and monitoring system. In addition we provide examples of the inputs and outputs to each of the components to illustrate what occurs at each step in the process of capturing spacecraft data. Next, we will describe the knowledge representation used for DPLAN's decomposition rules and for the resultant temporal dependency networks (TDN's) used by LMCOA. Finally, we describe the results of the technology demonstration at DSS-13.

How the DSN Operates

Voyager-1 is cruising at 17.5 kilometers/second toward the outer edge of the solar system. Though its onboard systems are mostly asleep during this phase of its mission, Voyager's health metrics are continually sent to Earth via a telemetry signal radiated by its 40-watt transmitter. It will take eight hours at the speed of light for the signal to reach its destination, Earth, a billion miles away. Upon arrival, the telemetry signal is received by an extremely sensitive ground communications system, NASA's Deep Space Network (DSN), where it is recorded, processed, and sent to the Mission Operations and Voyager project engineers, who assess the health of the spacecraft based on the contents of the signal.

The type of activity just described occurs daily for dozens of different NASA spacecraft and projects that 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 transformed into useful information.

Network Preparation

Figure 1 shows a partial map of some of the operational processes of the DSN (see [2] for a more complete description of the DSN processes). The first stage is called

Network Preparation and it encompasses the process boxes in Figure 1 labeled Schedule Resource, Resource Management, and Generate Ground Network SOE. The project initiates Network Preparation by sending a request for the DSN to track a spacecraft. The DSN responds to the request by attempting to schedule the resources (i.e., an antenna and other shared equipment)

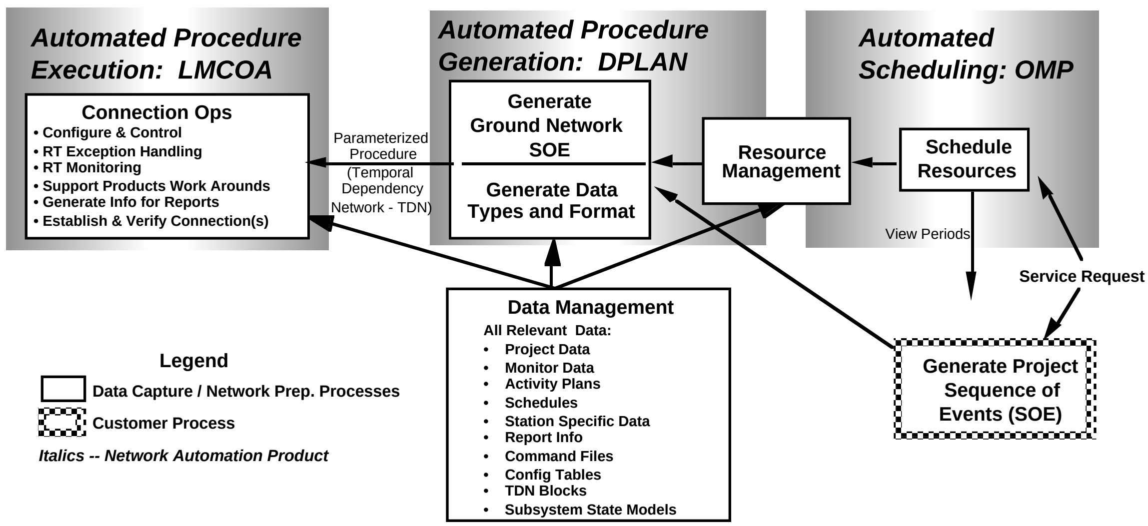

Figure 1: Automation technology applied to DSN operations processes

needed for the track. Next, the project prepares a Sequence of Events (SOE) describing the time-ordered activities that should occur during the track. The SOE includes actions that the DSN should take, (e.g., begin tracking the project's spacecraft at 1200 hours), and it also includes events that will occur on the spacecraft being tracked (e.g., the spacecraft will change frequency or mode at a designated time and the DSN should anticipate the event). The project SOE is sent to the DSN, which then generates its own version, called a Ground Network SOE. The Ground Network SOE is a more elaborate version of the project SOE in that it expands the activities from high level descriptions (e.g., begin tracking the spacecraft) into a finer level of detail for use by the operations personnel at the deep space station. The Ground Network SOE is sent to the deep space station, along with the required support data (see the Data Management process box in Figure 1).

Data Capture

The data capture process is performed by operations personnel at the deep space station--they configure the equipment for the track, establish the communications link, which we hereafter refer to as a 'link', and then perform the track by issuing control commands to the various subsystems comprising the link. Throughout the track the operators continually monitor the status of the link and handle exceptions (e.g., the receiver breaks lock with the spacecraft) as they occur. All of these actions are currently performed by human operators, who manually issue tens or hundreds commands via a computer keyboard to the link subsystems. The monitoring activities require the operator to track the state of each of the subsystems in the link (usually three to five subsystems), where each subsystem has many different state variables that change over time.

Automation of the DSN

In the last section we described the process for transforming a flight project service request into an executable set of DSN operations. As we have already pointed out, many of the steps of the described processes are intensely manual. We propose to implement portions of the process map shown in Figure 1 with a set of tools that we have built, and integrate these tools into an automation architecture. The specific tools and their functions are shown in Figure 1 -- OMP is applied to the resource scheduling process, DPLAN is used for automatically generating DSN operations procedures, and LMCOA automatically executes the operations procedures.

OMP: Automated Scheduling

The high level resource allocation problem for the DSN is handled by the Operations Mission Planner (OMP) scheduling system, which takes 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. OMP performs a vital function in the automation architecture, but we will not say anymore about it here and will focus instead on the other two components, DPLAN and LMCOA.

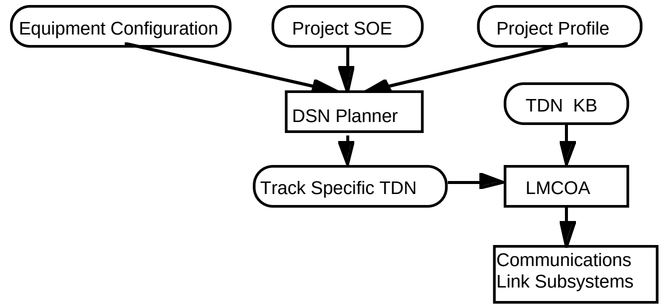

Figure 2: DPLAN and LMCOA Inputs and Outputs

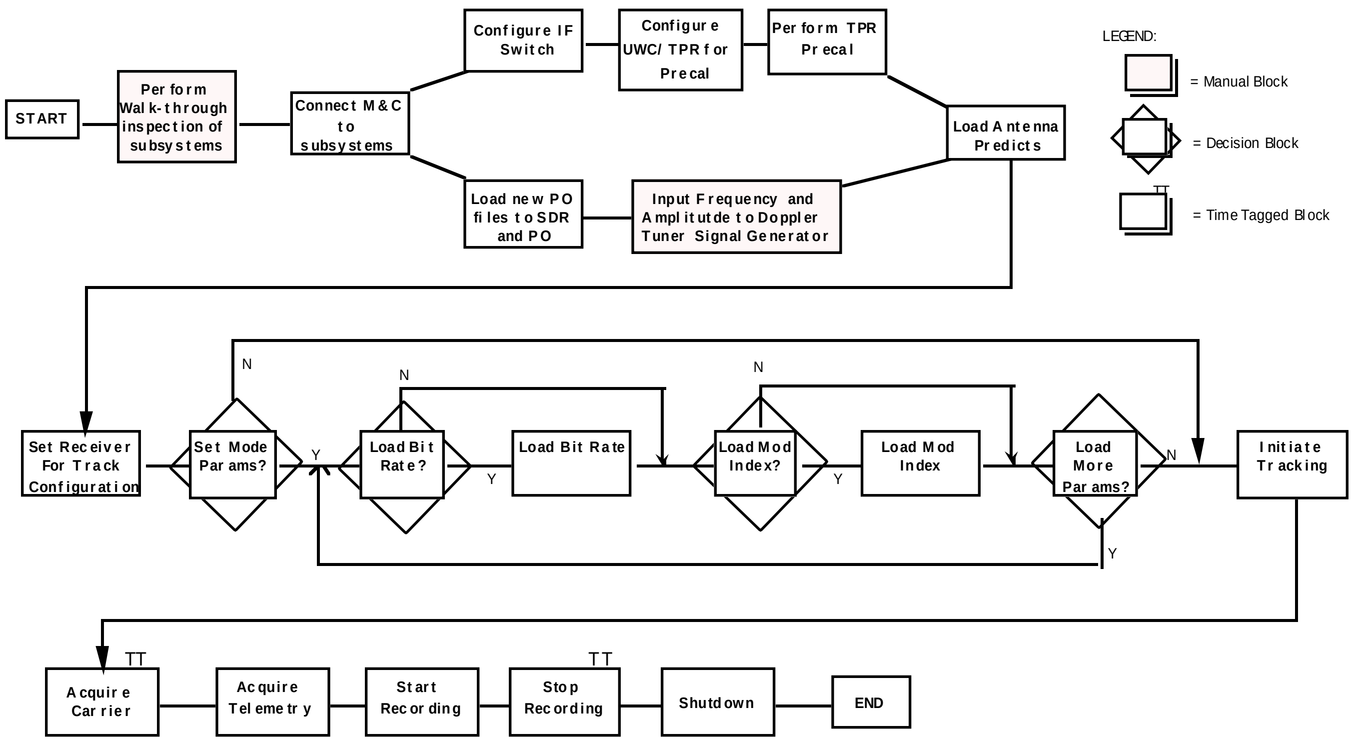

Figure 3: Temporal Dependency Network for Voyager Track

DPLAN: Automated Procedure Generation

The automated track procedure generation problem involves taking a general service request (such as telemetry - downlink of data from a spacecraft) and an actual equipment assignment (describing the type of antenna, receiver, telemetry processor, and so on), and generating the appropriate partially ordered sequence of commands (called a Temporal Dependency Network or TDN; see Figure 3) for creating a communications link to enable the appropriate interaction with the spacecraft. The DSN Antenna Operations Planner (DPLAN) uses an integration of AI Hierarchical Task Network (HTN) and partial order operator-based planning techniques to represent DSN antenna operations knowledge and to antenna operations procedures on demand from the service request and equipment assignment.

The DPLAN planner uses high level track information to determine appropriate steps, ordering constraints on these steps, parameters of these steps to achieve the high level track goals given the equipment allocation. In generating the TDN, the planner uses information from several sources (see Figure 2):

Project SOE: The project sequence of events specifies events from the mission/project perspective. As a result, the project SOE contains a great deal of information regarding the spacecraft state which is relevant to the DSN track, as well as a large amount of spacecraft information unrelated to DSN operations. Relevant information specified in the project SOE includes such items as the one-way light time (OWLT) to the spacecraft, notifications of the beginning and ending times of tracks, spacecraft data transmission bit rate changes, modulation index changes, and carrier and subcarrier frequency changes.

Project profile: This file specifies project specific information regarding frequencies and pass types. For example, the Project SOE might specify frequency = HIGH, and the project profile would specify the exact frequency used. The project profile might also other signal parameters and default track types.

TDN KB: The Temporal Dependency Network (TDN) knowledge base [3][4] stores information on the TDN blocks available for the DSN Planner and LMCOA to use. This knowledge base includes information regarding preconditions, postconditions, directives, and other aspects of the TDN blocks. It also includes information on how to expand the block parameters and name into the actual flatfile entry in a TDN.

Equipment Configuration: This details the types of equipment available and unique identifiers to be used to specify the exact pieces of equipment to be used in the track. These include the antenna, antenna controller, the receiver, and so on.

DPLAN uses the SOE and project profile to determine the overall goals of the track. The DSN planner reduces the high level track goals into executable steps by applying knowledge about how to achieve specific combinations of track goals in the context of specific equipment combinations. This information is represented in the form of task reduction rules, which detail how a set of high level goals can be reduced into a set of lower level goals in a particular problem-solving context. Each task reduction rule rigorously details the scope of its expertise in terms of track and equipment combinations. The information of scope of applicability of the rule can be considered in terms of a track goal hierarchy and equipment goal hierarchy, where the rule applies to all contexts below the rule in the relevant hierarchy (all specializations of its scope).

Using this problem specification, the DSN planner then uses task reduction planning techniques (also called hierarchical task network or HTN) [5] and operator-based planning techniques [6] to produce a parameterized track-specific TDN to be used to conduct the track. The actual planner used to generate the TDN is a modified version of the task reduction planning component of the Multimission VICAR Planner system [7]. This track-specific TDN and the SOE can then be used by LMCOA to operate the actual antenna to conduct the requested antenna track.

LMCOA: Automated Procedure Execution

The automated execution component, called the Link Monitor and Control Operator Assistant (LMCOA) uses the TDN (Figure 3) generated by DPLAN to perform the actual track and is responsible for monitoring the execution of the TDN. This involves ensuring that the expected conditions and subsystem states are achieved, certain types of closed-loop control and error recovery are performed in a timely fashion, and the correct dispatching of commands to the subsystems controlling the link occurs. LMCOA uses an operator-based representation of the TDN to represent necessary and desired conditions for execution of procedures and tracks relevant subsystem state.

The LMCOA performs the operations procedures for a tracking activity by executing a Temporal Dependency Network (TDN), which is a procedure that is automatically generated by DPLAN, as described in the last section. DPLAN composes the TDN so that it contains the procedures (TDN blocks) needed for a specific tracking activity, and it orders them according to its knowledge of the dependencies that are defined among the blocks as well as by what it knows about the pre- and postconditions of the blocks. The knowledge about interblock dependencies and about block pre- and postconditions is passed to the LMCOA, whose task it is to execute the end-to-end procedure. The LMCOA receives the TDN in the form of a directed graph, where the precedence relations are specified by the nodes and arcs of the network. The blocks in the graph are partially ordered, meaning that some blocks may be executed in parallel. Temporal knowledge is also encoded in the TDN, which includes both absolute (e.g. Acquire the spacecraft at time 02:30:45) and relative (e.g. Perform step Y 5 minutes after step X) temporal constraints. Conditional branches in the network are performed only under certain conditions. Optional paths are those which are not essential to the operation, but may, for example, provide a higher level of confidence in the data resulting from a plan if performed. More details about TDNs are provided in [3].

To execute a TDN, LMCOA performs the following functions: (1) it loads the parameterized TDN into the Execution Manager; (2) it determines which TDN blocks are eligible for execution and spawns a process for each TDN block; (3) it checks whether the preconditions of each TDN block have been satisfied; (4) once the preconditions are satisfied, it issues the TDN block commands; and (5) it verifies whether the commands had their intended effects on the equipment. The Operator interacts with the LMCOA by watching the execution of the TDN; the operator can pause or skip portions of the TDN, check the status of commands within individual blocks, and provide inputs where they are required. When part of a TDN fails, the LMCOA supports manual recovery by the operator by highlighting the point of failure and providing information about the preconditions or postconditions that failed to be satisfied.

Results

In February 1995 a comprehensive demonstration was conducted to validate the concept of integrating and using the AI software described in the preceding paragraphs to track a spacecraft with the DSN. In the demonstration, DPLAN generated the TDN shown in Figure 3 for a Voyager Telemetry Downlink track using the equipment configuration at Deep Space Station 13 in Goldstone, California, which included a 34-meter beam-wave guide (BWG) antenna and a telemetry processor. The TDN generated by DPLAN was successfully executed by LMCOA--a communications link was established with Voyager and the 34-meter BWG antenna tracked the spacecraft, with minimal human control. As a result of this demonstration, DPLAN and the concepts implemented in LMCOA are currently being transferred and implemented in the Network Control Project (NCP), which will replace two major DSN subsystems--the Monitor and Control (M&C) subsystem and the Network Planning and Preparation (NPP) subsystem.

References

[1] Deep Space Network, Jet Propulsion Laboratory Publication 400-517, April 1994.

[2] Final Report of the Services Fulfillment Reengineering Team, JPL Interoffice Memorandum, March 14, 1995.

[3] K. Fayyad and L. P. Cooper. "Representing Operations Procedures Using Temporal Dependency Networks," *SpaceOps '92*, Pasadena, CA, November 1992.

[4] K. Fayyad, R.W. Hill, Jr., and E.J. Wyatt. "Knowledge Engineering for Temporal Dependency Networks as Operations Procedures." *Proceedings of AIAA Computing in Aerospace 9 Conference*, 1993, San Diego, CA.

[5] K. Erol, J. Hendler, and D. Nau, "UMCP: A Sound and Complete Procedure for Hierarchical Task Network Planning," *Proceedings of the Second International Conference on AI Planning Systems*, Chicago, IL, June 1994, pp. 249-254.

[6] J. S. Pemberthy and D. S. Weld, "UCPOP: A Sound Complete, Partial Order Planner for ADL," *Proceedings of the Third International Conference on Knowledge Representation and Reasoning*, October 1992.

[7] S. Chien, "Automated Synthesis of Complex Image Processing Procedures for a Large-scale Image Database," *Proceedings of the First IEEE International Conference on Image Processing*, Austin, TX, November 1994, Vol 3, pp. 796-800.