The Overall Architecture of the HST Science Planning and Scheduling System

Roberto J. Samson

Space Telescope Science Institute
3700 San Martin Drive
Baltimore, MD 21218
samson@stsci.edu

Abstract

This paper presents an overview of the architecture and function of the Hubble Space Telescope (HST) Science Planning and Scheduling System (SPSS). SPSS is the system responsible for scheduling HST activities that must be in compliance with scientific as well as spacecraft health and safety constraints. The paper will focus on describing the constraint data and processing involved in producing flight calendars.

Introduction

The Science Planning and Scheduling System (SPSS) performs the detailed and accurate scheduling of Hubble Space Telescope (HST) activities that are required to capture the science data specified in groups of proposals. The activities must be scheduled in compliance with constraints imposed for spacecraft health and safety and science integrity, as well as maximize the science data returned by the telescope. Generally, activities are scheduled to the nearest second with pointing uncertainties of well under one arcsecond. These activities are normally scheduled a few weeks in advance of actual flight.

HST activities are nominally scheduled over a seven-day period and are defined in an SPSS "**calendar**" file. The flight calendar builder creates a calendar from a list of "**candidate**" observations selected from a pool maintained in a relational database. Prior to the calendar building stage, proposal instructions are transformed to internal data structures and undergo rigorous validation and planning phases.

This paper provides an overview of the process involved in producing HST flight calendars. It will focus on the SPSS data and functions required to schedule activities within proposer and spacecraft constraints. The first section describes the SPSS inputs from upstream systems such as TRANS and SPIKE. This is followed by a detailed listing of the constraints that are evaluated by SPSS. We then describe the manner in which processing is broken down into distinct "planning" and "scheduling" stages in SPSS. Finally, the automated process of building an efficient calendar from a pool of observations is presented.

SPSS Inputs

The SPSS software is built around a Standard Query Language (SQL) relational database called the Project Management Data Base (PMDB). Data from several upstream systems generate the scheduling constraints that are consolidated in the PMDB. The following systems and the data they provide are described in this section:

- TRANS - transformed proposal instructions from RPS2
- POMS - crafted parallel observation data
- SPIKE - long range plan and relative constraints
- MOSS - moving target ephemeris
- PDB - spacecraft health and safety constraints
- GSSS - Guide Star data
- FDF - HST and TDRS ephemeris

With the exception of the systems at the Flight Dynamics Facility (FDF), which reside at the Goddard Space Flight Center, all of these systems operate at the Space Telescope Science Institute (STScI).

Figure 1 illustrates how these systems interface with SPSS via the PMDB and auxiliary files. This figure also shows the SPSS processing tasks described later in this paper.

Proposer Constraints

The science and engineering proposal specifications and constraints ingested by the upstream proposal submission system RPS2 are transformed into a hierarchical data structure for SPSS processing. This data structure is created by the front-end system called TRANS. The structure levels, from highest to lowest are

- Proposals
- Scheduling Units (SUs)
- Observation Sets (OBSETs)
- Alignments

- Exposures

Note that the concept of a "**visit**" exists in the HST front-end systems. SPSS has no visit concept. An exposure can be regarded as a shutter-open activity. An alignment is a set of one or more exposures for a specific pointing of HST. An observation set is a set of one or more alignments that normally use the same set of guide stars. A scheduling unit is a group of one or more related observation sets. A proposal contains one or more related scheduling units.

The Parallel Observation Matching System (POMS) uses pointing information from scheduled "**primary**" observations to craft "**parallel**" SUs based on parallel observing proposals. Parallel observations are scheduled at the same pointing but using different science instruments or detectors than primary observations. POMS receives a representation of a weekly schedule of primary observation activities from SPSS and generates parallel SU data that are loaded into the PMDB and scheduled onto the final flight calendar (see Figure 1).

Scheduling units can also be grouped into link sets. Link sets describe, often complex, relative timing or orientation constraints between sets of SUs that can span over long time periods (years). Proposer specified relative constraints are loaded into the PMDB via the SPIKE software.

Long Range Plan

The SPIKE system processes the proposal database and produces a Long Range Plan (LRP). The LRP spreads observations out evenly over an observing cycle and provides each scientist with an approximate time when their observation will be scheduled. The LRP consists of lists of windows for each Scheduling Unit. The nominal case would be one eight week long "**plan window**" for each SU, however windows could be shorter, or could cover the entire cycle. The SPIKE plan windows are stored in a front-end relational database called ASSIST, which are accessed directly and used by the SPSS software.

Long Range Planning Versus Short Term Scheduling

Note the distinction between long range planning in SPIKE and short term scheduling in SPSS. The set of plan windows provided by SPIKE for every scheduling unit is used by SPSS to identify those scheduling units that can be executed within a particular week. The number of scheduling units that meet this criteria typically exceed the maximum available observing time for the seven-day calendar. Calendars with an excess number of these candidate SUs are referred to as **oversubscribed**. Automated SPSS scheduling tools are used to maximize the efficiency of the final calendar in SPSS. When building a flight calendar, candidate SUs with plan windows that do not extend after the end of the seven-day calendar are given higher scheduling priority. Candidates that "fall-off" a calendar are planned for subsequent weeks.

Three primary factors distinguish long range planning performed by SPIKE and short term scheduling performed by SPSS. The first factor is the orbit uncertainty of the HST spacecraft due primarily to atmospheric drag. The orbital true anomaly of the spacecraft can only be determined a few weeks in advance with acceptable accuracy needed for detailed scheduling.

The second factor is computational workload. SPIKE effectively reduces the time span for many time consuming calculations to its generated plan windows. SPIKE predicts, to a reasonable accuracy, when during a cycle a particular visit can schedule, without having to schedule activities on a calendar. SPSS, on the other hand, can determine with absolute certainty when an SU may schedule – by physically attempting to schedule all of the SU's activities on a calendar. To compensate for SPIKE's plan window uncertainty, a verification process is in place to detect mismatches between the SPIKE and SPSS schedulability windows.

The third factor is the possible dependencies that exist between any two scheduling units on a flight calendar. With the exception of relative timing and orientation constraints, SPIKE models each observation independently. In SPSS, however, this property does not hold. The HST science instruments, for example, impose well-defined warm-up/cooldown state transition restrictions. These restrictions cause overall observation durations to vary depending on the placement relative to other observations in the timeline.

Moving Target Data

Moving target observations require target positional data that are not specified in proposals. Positional data for moving targets such as planets, moons, and asteroids are modeled using piecewise Chebychev polynomials and visibility windows calculated by the Moving Object Support System (MOSS).

A Target data structure maintained in the PMDB contains detailed fixed and moving target position and characteristics information.

Spacecraft Health and Safety Constraints

Spacecraft health and safety constraints such as sun, moon, and antisun avoidance regions, power and solar array constraints are maintained in the Project Data Base (PDB). Some of this data are made available to SPSS in the form

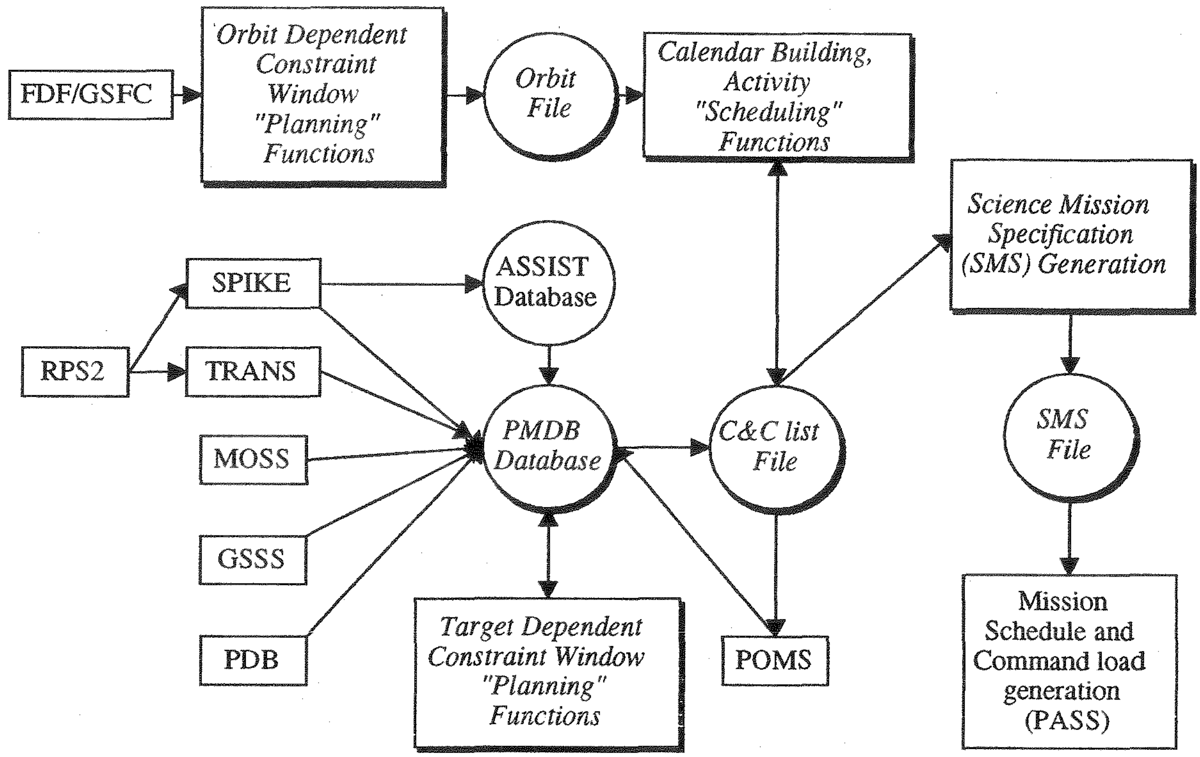

Figure 1: Diagram showing the high level interfaces between SPSS and the front-end proposal processing systems. Shadowed symbols with italicized text denote SPSS data and functions.

of namelists while others are loaded into relations in the PMDB. Science instrument operational constraints are directly maintained in the PMDB.

Guide Star Data

Guide Stars form another data structure. Single stars and pairs of stars for specific observation sets are determined by the Guide Star Selection System (GSSS) and stored in the PMDB. SPSS determines acquisition datasets for the guide star data. These datasets form spacecraft roll ranges for the observation set guide stars.

HST and TDRS Ephemerides

Predictive HST and TDRS (Tracking and Data Relay Satellite) position data are obtained from NASA's Flight Dynamics Facility (FDF) at the Goddard Space Flight Center (GSFC). In SPSS, the raw ephemeris is basically reduced to a set of non-linear equations that are used to estimate the HST and TDRS orbital elements as a function of time. This reduced data is stored in the PMDB.

SPSS Processing

SPSS Constraint Window Computation

As its name implies, SPSS processing is primarily composed of the "**planning**" phase, when scheduling windows are precomputed to reduce processing time, and the "**scheduling**" phase, during which spacecraft activities are scheduled on the nominal seven-day calendar.

Scheduling windows are defined as periods when some constraint is satisfied. In the planning phase, the constraints for which these windows can be computed are those that are either HST orbit dependent (e.g. SAA avoidance) or target dependent (e.g. sun avoidance), but not both. Scheduling windows for HST orbit dependent constraints are stored in an "Orbit file", and windows based on constraints that are solely target dependent are stored in the PMDB for every observation set. The orbit file windows are typically computed over a period of 10 weeks, which is the duration of HST and TDRS predictive ephemeris provided by the Flight Dynamics Facility (FDF) at the Goddard Space Flight Center. Target dependent constraint windows are stored in the PMDB as computed over a time period specified by the user. During the scheduling phase, windows for constraints that are both HST and target position dependent (e.g. moving target earth avoidance) are computed on the fly.

Table 1 is a summary of the constraints that must be met when scheduling spacecraft activities. Note that constraint types generally fall under two categories or types: those that are ensure science integrity and those that maintain the health and safety of the HST.

Constraint	**Constraint Type: H = Health and Safety S = Science Integrity**	**Subsystem responsible for computing windows**	**Computed during SPSS planning or scheduling stage?**	**SPSS Constraint Window Storage Area**
Sun Avoidance	H	SPSS	Planning	PMDB
Anti-Sun Avoidance	H	SPSS	Planning	PMDB
HST Roll Orientation	H	SPSS	Planning	PMDB
Power : • Nominal Roll • Off-Normal • Solar Array Shadowing • Thermal	H	SPSS	Planning	PMDB
Moon Avoidance	S	SPSS	Planning	PMDB
Earth Avoidance	S	SPSS	Scheduling	On demand
South Atlantic Anomaly Avoidance	S	SPSS	Planning	Orbit File
HST in Earth Shadow	S	SPSS	Planning	Orbit File
Moving Target Visibility	S	MOSS	N/A	PMDB
Phase Windows	S	PMDB	Planning	PMDB
Science Instrument Parallel Usage	S	SPSS	scheduling	On demand
Minimum/Maximum Separation	S	SPSS	Scheduling	On demand
Maximum Interruption	S	SPSS	Scheduling	On demand
Fine Guidance Sensor Constraints: • Earth Avoidance • SAA Avoidance • Guide Star Constraints (Roll Angle)	S S S	SPSS SPSS SPSS	Scheduling Planning Planning	On demand PMDB PMDB
Fixed Head Star Tracker Constraints: • Moon Avoidance • Earth Avoidance	S S	SPSS SPSS	Scheduling Scheduling	On demand On demand
User Specified (US) Windows	S	SPSS	Scheduling	On demand
Timing Link Set	S	SPSS	Scheduling	On demand
Orientation Link Set	S	SPSS	Scheduling	On demand
Plan Windows	S	SPIKE	N/A	PMDB

Table 1: Summary of SPSS constraint types. Note that pre-computed windows that are stored in the PMDB are eventually copied to the C&C list for use in scheduling. Windows that are calculated "On demand" are stored in internal memory as needed. For more detailed information on these constraints, see the on-line references listed in this paper.

Calendar Building and Activity Scheduling

For a given weekly schedule, SU data are copied from the PMDB to a file called a Candidate and Calendar List or "C&C List" file. This file contains a list of candidate SUs and a calendar of scheduled activities that support the observations specified in each SU. The final calendar contains the detailed schedule of all spacecraft activities for the time span. A file is used instead of direct SQL manipulations to improve performance. This file is essentially a snapshot of what is in the PMDB at the start of a calendar building week. If necessary, the C&C list can be updated with new data in the PMDB (e.g. if proposal specifications change) on demand.

Scheduling constraints apply to virtually every activity scheduled. The scheduling process determines what constraints apply to a given activity, either by rules built into SPSS, or by database specified flags. After determining which constraints apply for a given activity, windows are computed (or retrieved, if pre-computed) for times where the constraints are satisfied, and the activity is scheduled at the earliest opportunity within the window. If a constraint cannot be satisfied, then the scheduling fails, otherwise the activity is delayed until the constraint has cleared.

SPSS schedules all of the current HST science instruments, Fine Guidance Sensors (FGSs), Gyros, Fixed Head Star Trackers (FHSTs), and other spacecraft hardware. The spacecraft activities that SPSS schedules include:

- Vehicle Maneuvers : Slews, Small Angle Maneuvers, and Track 51 Linear Scans (moving target tracking)
- Pointing Control System : Guide Star Acquisitions, Guide Star Reacquisitions, Guide Star Handoffs
- Target Acquisitions
- Fixed Head Star Tracker Updates: Delay Mode Roll Updates, Maneuver Mode Full Updates
- Science Instrument Reconfigurations
- Science Activities: Data Collection, Calibration, Interleaver Science, Parallel Science
- Tracking & Data Relay Satellite Contacts
- Tape Recorder Usage

As mentioned earlier, the SPSS software must schedule observations under a large number of constraints. For example, some instruments are sensitive to the South Atlantic Anomaly (SAA). SPSS determines time spans of HST passage through various SAA models and schedules observations to avoid these periods. Observations must avoid the bright and dark Earth by different avoidance angles. SPSS must model the Earth's terminator and determine avoidance windows depending on the instantaneous Sun, HST, and fixed/moving target positions. Various automatic, manual, and graphical are used to schedule the observations on the calendar. The observations are scheduled to maximize the scientific return from HST.

The scheduling unit (SU) is the basic unit of scheduling within SPSS. That is, a scheduler will specify that a scheduling unit is to be added to a calendar. To be considered successful, every supporting activity must be scheduled constraint free. The calendar builder is initially provided with an oversubscribed list of scheduling units with plan windows (from SPIKE's long range plan) that overlap the seven-day calendar. Oversubscription poses the following problem: which of the scheduling units are to be placed on the calendar, and in what order so as to maximize "efficiency"?

An automated scheduling tool that implements a heuristic based greedy search algorithm is used to arrive at a best ordering of scheduling units. The problem can be modeled as the need to search for the best possible path down an n-ary tree, where each node represents a scheduling unit that is scheduled in the best possible place (not necessarily in chronological order) on an updated calendar.

The calendar is built by starting at the root node where a heuristic is applied to determine which path to take to the next level in the tree. Once a path is determined, the scheduling unit is permanently scheduled to create an updated calendar, which serves as input to the selection process down subsequent levels in the tree. A complete solution is obtained when a leaf is reached or when the heuristic is unable to make a valid selection.

The heuristic consists of two parts. The first part involves the ranking of each scheduling unit based on several factors, which includes a user specified priority factor and "difficulty to schedule" adjustments. In the second part, the top N-ranked scheduling units are added alone on the input calendar, where N is a user-specified parameter. A new score is computed for each scheduling unit/place combination. The scheduling unit with the highest place score is then added permanently on the calendar.

The greedy search has its shortcomings in that the resulting calendars have required manual repair to accommodate tight constraints. This algorithm can likely be improved by including knowledge of SU schedulability windows – i.e. windows where an SU will schedule on an empty calendar.

SPSS Outputs and Post-Calendar Processing

When the calendar is completed, other SPSS tools update the PMDB with basic information about the scheduled

SUs. Finally, a Science Mission Schedule (SMS) file is generated from the calendar and data in the PMDB. An SMS is sequence of detailed spacecraft commands. Another subsystem, PASS, verifies the SMS, inserts health and safety activities, and generates spacecraft ready command loads.

SPSS Software and Hardware Characteristics

The SPSS software consists of nearly one million physical lines of F77 and C/C++ code in about one hundred executable programs. Most of these programs are supported on both AXP/VMS and Sparc/Solaris platforms.

Summary

SPSS has been used operationally in the planning and scheduling of HST activities for the past seven years. Several modifications have been made to SPSS during this period to accommodate new instruments as well as new spacecraft capabilities and constraints. The system has been proven to adapt well to these changes.

References

More detailed documentation can be found on the World Wide Web at the following URLs:

http://www.pst.stsci.edu/~samson/overview_doc.cgi
http://www.pst.stsci.edu/~samson/windows_doc.cgi