

Commentary on Validating Resource Usage in Least Commitment Planning

Nazma Ferdous and Mark Giuliano

By Rob Sherwood
Jet Propulsion Laboratory
4800 Oak Grove Dr.
Pasadena, CA 91109
Rob.Sherwood@jpl.nasa.gov

This paper describes some improvements to the planning process for the Hubble Space Telescope (HST). The HST observations are planned using a two step process:

1. A 56-day long term plan is created using the least commitment SPIKE planning system with a resolution of 1 day
2. A 1-week short term plan is created using a least commitment planner with a resolution of one-second

SPIKE uses a very simple model to calculate the resource usage that results in schedules that may appear to be legal but actually use more resources than are available. Similarly some schedules appear to have resource conflicts when they are actually legal schedules. The main thrust of this paper was to solve the resource modeling problems in SPIKE.

The SPIKE scheduler assigns resources evenly over plan windows based on user constraints such as sun and moon avoidance, guide star restrictions, and earth occultation. The plan windows represent legal times for observations to be scheduled. The long term schedule tracks 2 resources:

1. Orbits used for observing the target
2. Five resources related to how orbits are affected by the South Atlantic Anomaly (SAA)

The long term scheduler keeps track of orbits that are impacted by the SAA. If the spacecraft flies over the SAA in an orbit, and not at a time when the Earth is occulting the target, the orbit is affected by the SAA (SAA impacted and SAA hiding orbits). Observations cannot be scheduled during SAA crossings. Because only a third of the orbits do not fly over the SAA anomaly, these orbits are very valuable (SAA free orbits). The scheduler has a priority to try and use the SAA hiding orbits first, followed by the SAA free orbits.

The short term scheduler adds additional spacecraft constraints such as slewing between observations, instrument configuration, data upload and download times, and thermal and power constraints. The two-phased scheduling approach is used for 3 reasons:

1. Exact in-track position is not known until a few weeks before the observation
2. Scientists have last minute requests
3. Observations plans change as the spacecraft capability changes

The long term planning resource modeling problems in SPIKE were addressed by using new software that verifies resource usage – the plan validation scheduler. The plan validation scheduler uses an iterative deepening repair algorithm. Without going into the details of the planning steps, the end result is that the plan validation tool flattens the resource profile and points out regions with actual resource contention. Users can then manually fill in areas that are underutilized. This tool was added to the SPIKE planner.

The author makes several comments in the paper about how the new tool has streamlined the long term planning process. These improvements have led to a savings in both labor and time to build long range schedules. In addition, the plans being generated have included more observations that are more resource balanced. This has been discussed qualitatively, but it would be helpful in future publications to show these gains quantitatively. The developers on their team have discussed creating a plan quality matrix, but this has not been implemented yet.

Another measure of plan quality is the amount of manual replanning. From discussions with the author, this is currently 5-10%. Could this number be pushed closer to 0% with more optimization? Perhaps this is not possible because some manual planning must occur for targets of opportunity, engineering activities, and late knowledge of orbital ephemerides. This might be another area to examine more closely.

I had some discussions with the author about imaging campaigns that last several weeks. With a one-week planning window, some campaigns will cross over the week planning boundaries. There could be additional gains in optimizing global plans over longer periods. There are so many interruptions in science observation time from Earth occultations and the SAA, that using one week planning windows has not resulted in much penalty

to optimizing longer campaigns. They are typically broken down into smaller few-orbit long observations. Still this is an area that could result in additional efficiencies. With all the restrictions on the planning window, it could be beneficial to plan for longer periods and perhaps even overlap multiple image campaigns. That would allow some additional flexibility for the planner to place observations on the schedule. There could be additional science gains from longer observation windows as well. It might be worthwhile to discuss this option with the Hubble scientists. There may be opportunities to coordinating multiple synergistic imaging campaigns with either different instruments or related targets.

The authors have presented a tool for validating resource usage on long range plans developed for HST. This tool has resulted in better plans that are generated more quickly. Additional areas of optimization have been discussed, as well as quantitative measures of plan quality.