A Commentary On:
Generating a Long Range Plan for a New Class of Astronomical Observatories

David E. Smith

NASA Ames Research Center
Mail stop 269–2
Moffett Field, CA 94035
de2smith@ptolemy.arc.nasa.gov

Introduction

The paper *Generating a Long Range Plan for a New Class of Astronomical Observatories* by Laurence Kramer, describes how scheduling for the Hubble Space Telescope (HST) is broken up into two phases:

1. *Long Range Planning* (LRP) – where the observations for an extended period (a year or more) are each assigned to plan windows of typically 4 to 8 weeks duration.
2. *Short Term Scheduling* – where a detailed schedule of observations is constructed for a single impending week.

The paper goes on to describe a proposal for adapting this same architecture to the Space Infrared Telescope Facility (SIRTF). For SIRTF, observations are typically much less constrained than for HST. However, resource management is much more critical, because there is a finite supply of cryogen for cooling the instruments, and switching instruments uses both coolant and time. As a result, the author proposes to construct explicit instrument windows during the LRP process and to use those windows to help assign plan windows for observations.

To start, I want to comment briefly on the general nature of observation scheduling problems. Second, I wish to comment on the issue of maintaining both flexibility and stability in a schedule, and offer a different way of thinking about stability. Finally, I wish to comment on the proposed use of instrument windows for SIRTF and suggest an alternative.

Observation Scheduling

There is a great deal of similarity between the scheduling problem for HST and SIRTF, and observation scheduling problems for other facilities such as automated ground-based observatories [5], airborne observatories (KAO and SOFIA) [6], and earth observing satellites (EOS). For all of these observatories, investigators submit specific observing proposals, usually well in advance. Proposals are then reviewed and ranked in some fashion. The result is a large collection of requested observations with constraints on date, time, sky condition, and other operational parameters. The challenge is to maximize scientific return, subject to the constraints. The problem is that for all of these facilities there are significant sources of operational uncertainty – weather conditions for ground-based observatories and EOS, water vapor conditions and air traffic delays for KAO and SOFIA, and uncertainty due to slewing, target acquisition, and observation duration.

Are these planning problems or scheduling problems? Both and neither. On the one hand, there are usually more observations than can be performed in a given period, so the task involves *choosing* some subset of the available observations. Choosing the operations to perform sounds like planning, but the choices are simple ones; the choice of one operation does not lead to other choices (as with preconditions). In addition, these problems involve rich temporal constraints and reasoning about continuous resources. While planning problems can include these characteristics, the solution techniques are not robust or mature (see [11] for discussion of this issue).

In contrast, scheduling problems often involve temporal constraints between tasks, continuous resources, and optimization. As a result, the constraint-based representation commonly used for scheduling problems is particularly relevant. However, there is little relationship between the common job shop scheduling problem and observation scheduling problems. In particular, there is no parallelism in observation scheduling – only one target can be observed at a time. The optimization criteria is also much different.

In sum, observation scheduling is a rather different animal, and neither the scheduling or the planning communities have a well developed set of tools for addressing this class of problems.

Flexibility and Stability

In the paper, Kramer argues that successful scheduling systems for HST and SIRTF need to generate schedules that are both flexible and stable. Flexibility is needed because of uncertainty – observation requests change, orbits are uncertain, and instrument behavior and performance are not entirely predictable. Stability is needed because astronomers have expectations about when their observations will be performed and they must make commitments based on those expectations. As a result, changes to the observing schedule must be made with care.[1]

The properties of flexibility and stability are fundamentally at odds with one another. Flexibility is obtained by delaying commitment – that is, scheduling decisions are not locked in until the last moment. In contrast, stability is guaranteed by early commitment – locking down decisions about the date and time of each observation. In Spike, this trade-off is accomplished by:

1. Early partial commitment to a 4-8 week range of dates for each observation
2. Delaying commitment on the exact dates and times for individual observations until shortly before execution

From a theoretical point of view, the early partial commitment done in LRP is somewhat different than operations usually performed in constraint-based scheduling. To see this, consider the observation scheduling problem as a constrained optimization problem in which there are variables representing the possible start times for each observation. (For each observation, the value of this variable will be a set of time intervals corresponding to the intervals in which the observation is possible.) By constraint propagation techniques such as arc-consistency and edge finding [4, 9], we can narrow down these intervals by virtue of other constraints on the observation, constraints on telescope operation, and time constraints between observations (see [7] for details). But for HST and SIRTF, this can still leave a wide range of possible start times for each observation. The purpose of the LRP process for HST and SIRTF is to narrow down these ranges so that 1) the remaining short-term scheduling problem is somewhat more constrained and hence easier to solve, and 2) astronomers can be given some idea when their observations are likely to be performed. When viewed as constraint optimization, the LRP is making commitments to subsets of the possible values for observation start time variables, instead of making commitments to individual values for those variables.

Schedule stability is enforced, because on subsequent rescheduling the LRP will not rescind these interval choices (except for observations that have been modified). In effect, the LRP refuses to backtrack on window narrowing choices that were made at an earlier date. This approach certainly guarantees the desired schedule stability, but could be problematic in other more highly constrained scheduling domains. The trouble is that there is no guarantee that a solution will exist to the resulting (more constrained) scheduling problem. In fact, scheduling systems typically need to search through many alternative choices for variable values in order to find acceptable schedules. The LRP usually gets away with this kind of irrevocable commitment, because the HST and SIRTF problems are generally very underconstrained.

For a domain in which this kind of commitment will not always work, how can we preserve a reasonable level of schedule stability? I contend that stability is an optimization criterion in rescheduling – there is some penalty associated with delaying an observation beyond its original predicted window, and the extent of the delay is presumably correlated to the magnitude of the penalty. (There is no penalty for performing an observation early, but doing so might displace another observation, which could result in a penalty.) Using this optimization criterion, a scheduler would prefer schedules in which all observations take place prior to the end of their predicted window, but alternatives could be considered and evaluated. Note that treating schedule stability as an optimization criterion, and allowing window revision does not require that we reschedule from scratch each time. It is still possible, and rational, to reschedule using local search methods, treating the previous schedule as the seed for the local search.

Of course, the LRP could use this approach when rescheduling. It would allow greater flexibility in rescheduling, and could allow more principled recovery from larger scale failures such as equipment problems aboard the telescope or unusual events that preempt existing observations.

Instrument Windows

For SIRTF, observations are typically much less constrained than for HST. This means bigger initial time windows. However, resource management is much more critical, because there is a finite supply of cryogen for cooling the instruments, and switching instruments uses both coolant and time. In the paper, the authors propose to 1) assign time windows for tasks with very small constraint windows, 2) build explicit instrument windows around those time windows, and 3) assign time windows for the remaining observations, splitting or adjusting instrument windows if necessary to accommodate observations.

I am not convinced that the construction of instrument windows is the best approach to this problem, because the instrument windows must be chosen, then continually revised to accommodate additional observations. In the final schedule, instrument windows are merely consequences of the choices of observation windows. In other words, it is the observation constraints that determine the best set of instrument windows, not the other way around. The paper argues that the observations are so under constrained that instrument windows are needed to help constrain the assignment of observation windows. If this were so, there would be no need to revise the initial set of observation windows to accommodate the remaining tasks.

Switching instruments on SIRTF is a classic example of a *setup step* common to many job shop scheduling problems. For example, consider a milling machine on a factory floor. Changing the bit or the clamping arrangement on the table is a time consuming and hence costly operation. If the previous task uses the same bit and the same clamping arrangement, no cost is incurred. Thus, there is considerable incentive to

---

[1] The need for both flexibility and stability is not unique to observation scheduling; flexibility and stability are required in any domain where there is significant uncertainty, but commitments must be made early. This is more common than not in real world scheduling problems.

schedule tasks in such a way as to minimize the number and cost of such setup steps. In general, simple search techniques like A*, Limited Discrepancy Search [8], and HBSS [3] should work quite well. To see how this might work for SIRTF, consider the partial schedule shown in Figure 1.

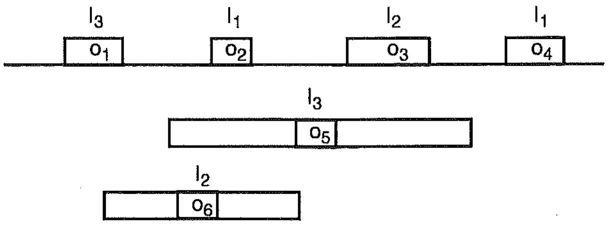

Figure 1: The search for Instrument windows.

Initially there are several observations in the schedule that can only be performed at specific times. For each observation, the required instrument is indicated above the observation. Initially, we have a lower bound on the instrument change cost for this schedule. Somewhere between o_1 and o_2, we will at least need to change from instrument I_3 to I_1. Similar transitions must occur between the other observations. If we let c_{ij} be the cost of changing from instrument i to j, the cost for this partial schedule is $c_{31}+c_{12}+c_{21}$.

Now consider an observation o_5 that has a much larger range of possible start times. Given the existing schedule, there are three places we could put this observation: before o_2, between o_2 and o_3, and after o_3. If o_5 occurs before o_2, there is no additional instrument change cost; only the transition from I_3 to I_1 is required. If o_5 occurs between o_2 and o_3, the cost increases because we must now transition from instrument I_1 to I_3, then back to instrument I_2. The results are similar if o_5 occurs after o_3. For subsequent observations a bit more reasoning is required, because we need to consider the possible ordering relative to flexible observations like o_5. However, the general idea is to search the space of relative observation orderings, in order to minimize instrument change cost. The resulting best ordering will dictate the best set of instrument windows.

In doing this search, one could also make use of probabilistic instrument usage profiles similar to those used by Sadeh [10] and Fox [1, 2]. In Figure 2, I have illustrated a simple profile for instrument I_2 from the example in Figure 1.

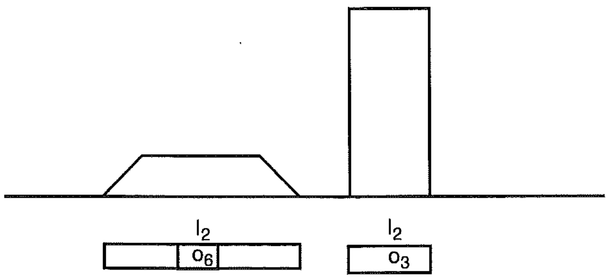

Figure 2: Probabilistic usage profile for instrument I_2.

For o_3, there is no time flexibility so I_2 is used with probability 1 throughout the interval. For o_6 there is flexibility, so the probability mass is distributed in a trapezoid. A composite usage profile for an instrument is constructed by summing up the profiles for all observations using that instrument.. These profiles provide heuristic guidance on the best windows to assign for a task – namely those for which there is high demand for the instrument and lower demand for other instruments.

In summary, I believe that for SIRTF, a more systematic search of the space of possible observation windows is practical, and that these observation windows will dictate the resulting instrument windows.

References

1. Beck, J., Davenport, A., Sitarski, E., and Fox, M. 1997. Texture-based heuristics for scheduling revisited. In *Proc. 14th Nat. Conf. on AI.* 241–248.

2. Beck, J. C. and Fox, M. 1998. A generic framework for constraint-directed search and scheduling. In *AI Magazine* 19(4).

3. Bresina, J. 1996. Heuristic-Biased Stochastic Search. In *Proc. 13th National Conf. on AI (AAAI-96).*

4. Carlier, J. and Pinson, E. 1989. An Algorithm for Solving the Job Shop Scheduling Problem. *Management Science, 35(2).* 164-176.

5. Drummond, M., Bresina, J., Edgington, W., Swanson, K., Henry, G., and Drascher, E. 1995. Flexible Scheduling of Automatic Telescopes over the Internet. In *Robotic Telescopes: Current Capabilities, Present Developments, and Future Prospects for Automated Astronomy.* Edited by G. Henry and J. Eaton. ASP Conf. Series, Vol. 79. Astronomical Society of the Pacific.

6. Frank, J. 2000. SOFIA's choice. In *International Workshop on Planning and Scheduling for Space Exploration and Science* (this volume), San Francisco, CA.

7. Kramer, L. and Giuliano, M. 1997. Reasoning About and Scheduling Linked HST Observations with Spike. In *International Workshop on Planning and Scheduling for Space Exploration and Science*, Oxnard, CA.

8. Harvey, W. and Ginsberg, M. 1993. Limited Discrepancy Search. In *Proc. 14th Int. Joint Conf. on AI (IJCAI-93).*

9. Nuijten, W. 1994. *Time and Resource Constrained Scheduling.* Ph.D. dissertation, U. Eindhoven.

10. Sadeh, N. 1994. Micro-opportunistic scheduling: the micro-boss factory scheduler. Technical report CMU-RI-TR-94-04, Carnegie Mellon Robotics Institute.

11. Smith, D., Frank, J., and Jónsson, A. 2000. Bridging the Gap Between Planning and Scheduling. To appear in *Knowledge Engineering Review* 15:1.