Scheduling Tools for Astronomical Observations

Mark D. Johnston (johnston@stsci.edu)

Space Telescope Science Institute
3700 San Martin Dr.
Baltimore MD 21218

Abstract.

A key challenge for the new generation of 8m class facilities will be ensuring that the best observing conditions are fully exploited to carry out unique science — science that can *only* be accomplished when these conditions obtain. Flexible scheduling is one approach to this problem: by scheduling a mix of observations from a pool of available observing programs, it is possible to select best matches to changing conditions. In many ways, this mode of observing is similar to that used for satellite observatories (except that hands-on intervention on the ground is much easier!) Among the issues to be dealt with in flexible scheduling are those of appropriate approaches and tools. This paper discusses some of the options for portable scheduling tools and the use of scheduling technology in observatory operations. Such use is not limited to schedule generation, but can also fruitfully encompass feasibility and tradeoff studies of different observing strategies, and simulations of operations under various conditions. To facilitate communication and interaction among observatories involved in advancing scheduling technology, it would be of wide benefit to adopt a common "scheduling data interchange format" for transferring scheduling information from one site/system to another, and that benchmark problem sets be defined and published for testing and evaluating different approaches.

1. Introduction

The use of scheduling system software offers the astronomical community the potential to make the most productive and efficient use of world-class observing facilities. By making these scheduling systems portable, it will be possible to take the lessons learned in one situation and apply them elsewhere, thus providing more capability at lower cost. This paper discusses some of the issues related to this topic:

- some caveats on astronomical scheduling terminology
- the role and application of scheduling in a ground-based observatory setting
- different approaches to integrating scheduling into observatory operations

- capabilities to look for in candidate systems
- distributed scheduling
- three recommendations for specific steps that can be taken now to facilitate the future use and exchange of scheduling technology

2. Terminology Caveats

The discussion of scheduling in a ground-based astronomy context does not have a well-defined common terminology. This complicates the discussion of new approaches, since there are major disagreements about the wisdom of changing approaches to telescope operations that have served the community well for generations. Among the terms to be found in the literature, or in use at workshops like this one, are: flexible scheduling, queue scheduling, service observing, remote observing, classical observing, integrated scheduling, and block scheduling. These terms often have quite different meanings to different people.

For the purposes of this discussion it is important to separate the issues of *policy* applied to telescope scheduling from the *tools* used to carry out that policy. Appropriate tools will facilitate the implementation of policy, not determine the policy itself. Such tools could support a wide range of observing "modes", ranging from full-night allocation to astronomers who travel to the telescope (and there make use of tools to help optimize the night's activities), to a full service modes where the observatory operations staff choose the night's exposures based on a matching a pool of candidate exposures to changing seeing conditions.

3. Why Schedule?

Automated scheduling has been common in space-based astronomy for years (Johnston and Miller 1994), but is relatively new to most of the ground-based community. A brief look at the general advantages of scheduling satellite observatories reveals the following driving factors:

- interleaving observations boosts efficiency, thus obtaining more science from the facility over a limited lifetime
- scheduling to optimize critical onboard resources allows for longer mission lifetimes
- the quality of individual observations is improved by scheduling them in consideration of the entire science program, rather than arbitrarily by assignment in blocks of time

Of course a major motivation for automated scheduling of satellite missions is the scarcity of crucial resources and the absence of any possibility for intervention should things go wrong. For the new generation of ground-based observatories under construction now, the scarcity of the best observing conditions provides an important motivation. But a crucial difference from space missions is the possibility to intervene, and even the desirability of using access to human

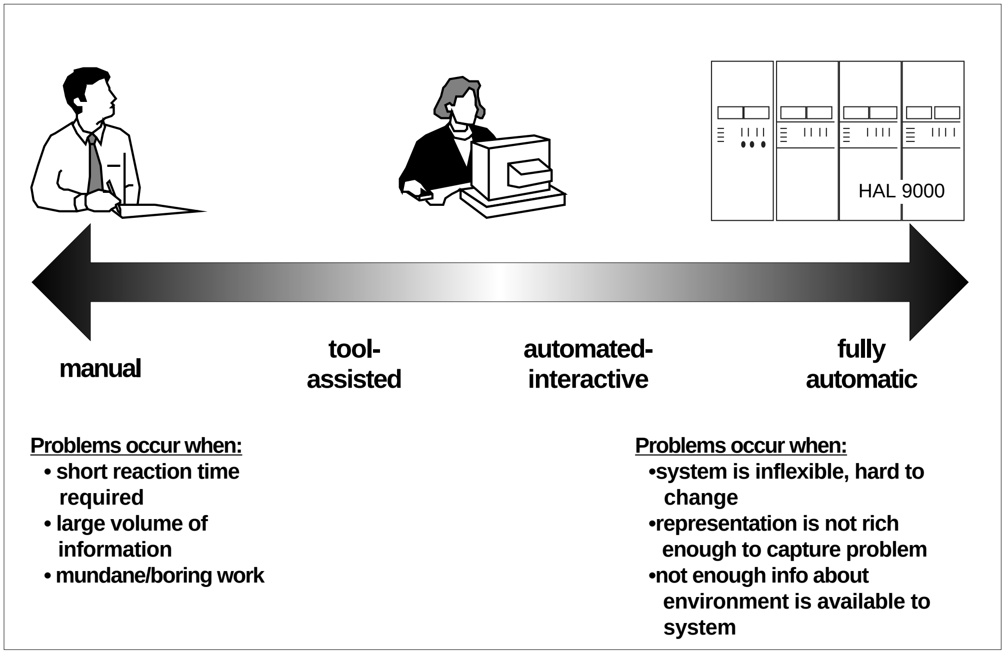

Figure 1. Scheduling automation can be introduced in a spectrum of ways, ranging from not at all to full "lights-out" operations. There are problems with both extremes.

expertise as a means to minimize costs and react to the need to change. Thus the challenge for scheduling in a ground-based observing context is to reap the rewards of higher quality of science and efficiency of operation, while preserving the advantages of a "hands-on" facility that needs to adapt quickly to changing conditions of all kinds.

4. Automated Scheduling in a Ground-Based Context

4.1. The Role of Scheduling

One often thinks of scheduling only in a strict operations context, but in fact the role of scheduling is substantially broader (Chavan et al. 1995, Miller 1995a). There is of course a major role in operations:

- to develop and maintain the observation schedule
- to reschedule when conditions warrant, either due to changes in the environment, or in reaction to discovery and changes in science objectives

In addition, there is an important place for scheduling support in planning at the observatory level:

- to decide what program mix to accept in a scheduling period

- to carry out feasibility and tradeoff studies
- to conduct resource planning and bottleneck detection
- to develop metrics that inform the staff and users how effectively the facility is being utilized

A further role is in planning at the observer level:

- to write observing proposals with a better chance of being allocated observing time
- in designing the plan for a night's observations

Finally, there is even a *multi-facility* role, in planning and implementing multi-wavelength and coordinated campaigns. In this case one is interested in maximizing the chances for success across more than one observatory, often involving both ground- and space-based facilities.

4.2. Approaches and Timescales

There are a range of approaches that can be taken to scheduling, as illustrated in Figure 1. At one end of the spectrum is a completely manual approach, which is perfectly adequate when problems are small, not too complicated, and do not change frequently.

At the next stage of automation one can introduce tools to support manual scheduling: e.g. to summarize what scheduling opportunities exist, what the current schedule contains, etc. Again, this approach is feasible so long as the problem is not too large or volatile.

Further automation leads to what can be called *automated-interactive* mode, wherein the human operator still completely controls the scheduling function but does so by running a system which creates and maintains the schedule. In this mode the system takes on a *decision support* role for the user, enabling him or her to handle large, complex, and volatile scheduling problems.

A problem emerges with increasing automation, however: it has too often been found that large automated systems lead to inflexibility, since the systems themselves become hard to change. And, given the richness of real-world scheduling problems, it is very difficult to build in from the outset a complete description of the problem that will not have to change with experience. A related problem is that of providing to the scheduling system certain information that humans gather very easily, such as "rumors" of problems with a particular piece of equipment, or "obvious" changes in observing strategy based on what's seen in the data so far. These factors suggest that there is an optimal degree of automation which effectively assists the human operator, but is not so rigid and constraining that it creates more problems than it solves.

Scheduling can be introduced on a variety of timescales. *Classical mode* observing (Figure 2a) often refers to allocating whole nights to observers who travel to the telescope. Some time is typically also scheduled for engineering and maintenance activities. In this mode, scheduling tools can help optimize the assignment of observers to nights, taking into account target visibility, phase of moon, etc. While this scheduling task has typically be done manually in the past, the advantages of automatic support become more apparent when there are many telescopes to schedule and many interacting constraints to satisfy.

A more sophisticated mode of operation is illustrated in Figure 2b, which can be termed "mixed mode" observing. In this scenario, the highest level scheduling is performed on a *long-term schedule* of some months in duration. This period is allocated to blocks for observers to travel to the telescope (as in classical mode), as well as into periods dedicated to "service observing" which will be conducted by the telescope operations staff. The service observations are themselves scheduled in a *medium-term schedule*, which tracks the degree of completeness, moon phase requirements, target visibility, and so forth, and is used by the service observers to determine most likely candidate observations for each night. Then, during the night, the *short-term scheduler* is used to help match these candidate observations to changing observing conditions, using the medium-term schedule as a guide. At each level, feedback from execution is essential for proper rescheduling. Furthermore, it is desirable that the scheduler at each level be the same, to simplify the interactions between the different levels.

4.3. Desirable Capabilities

Based on the experience of running Hubble Space Telescope (HST) as a remote- and service-operated observatory, and of continuing to re-engineer the HST ground system (Johnston and Miller 1994, Johnston and Minton 1994, Miller 1995b), it is clear that there are some important characteristics to look for in a scheduling system for astronomical applications:

- adaptability — the system should be easy to change to react to new requirements and desires; the schedules should be easy to change to react to new conditions
- interactivity — users should have control over and visibility into the scheduling process at all times; responsiveness is a key requirement
- connectivity — the scheduler needs to have straightforward interfaces with with the other components of the observatory operations environment, including databases, observation specifications, and the results of execution

The Spike system, developed at Space Telescope Science Institute and used there for long-range scheduling, provides a good example of a system that has been developed explicitly with the intention of being applied to other kinds of scheduling problems. There are other systems as well, some still in research phases, others are various stages of being fielded. The interested reader is referred to the URL http://www.stsci.edu/spike for more information and pointers to a wide range of scheduling resources, both applied and theoretical.

It is worth noting that, while scheduling has typically been regarded as the province solely of observatory operations, it can be useful to provide individual observers access to the same tools. For Hubble Space Telescope, the feedback from users of bundling some of the Spike scheduling software with proposal preparation tools has been very positive. A major advantage of this strategy that

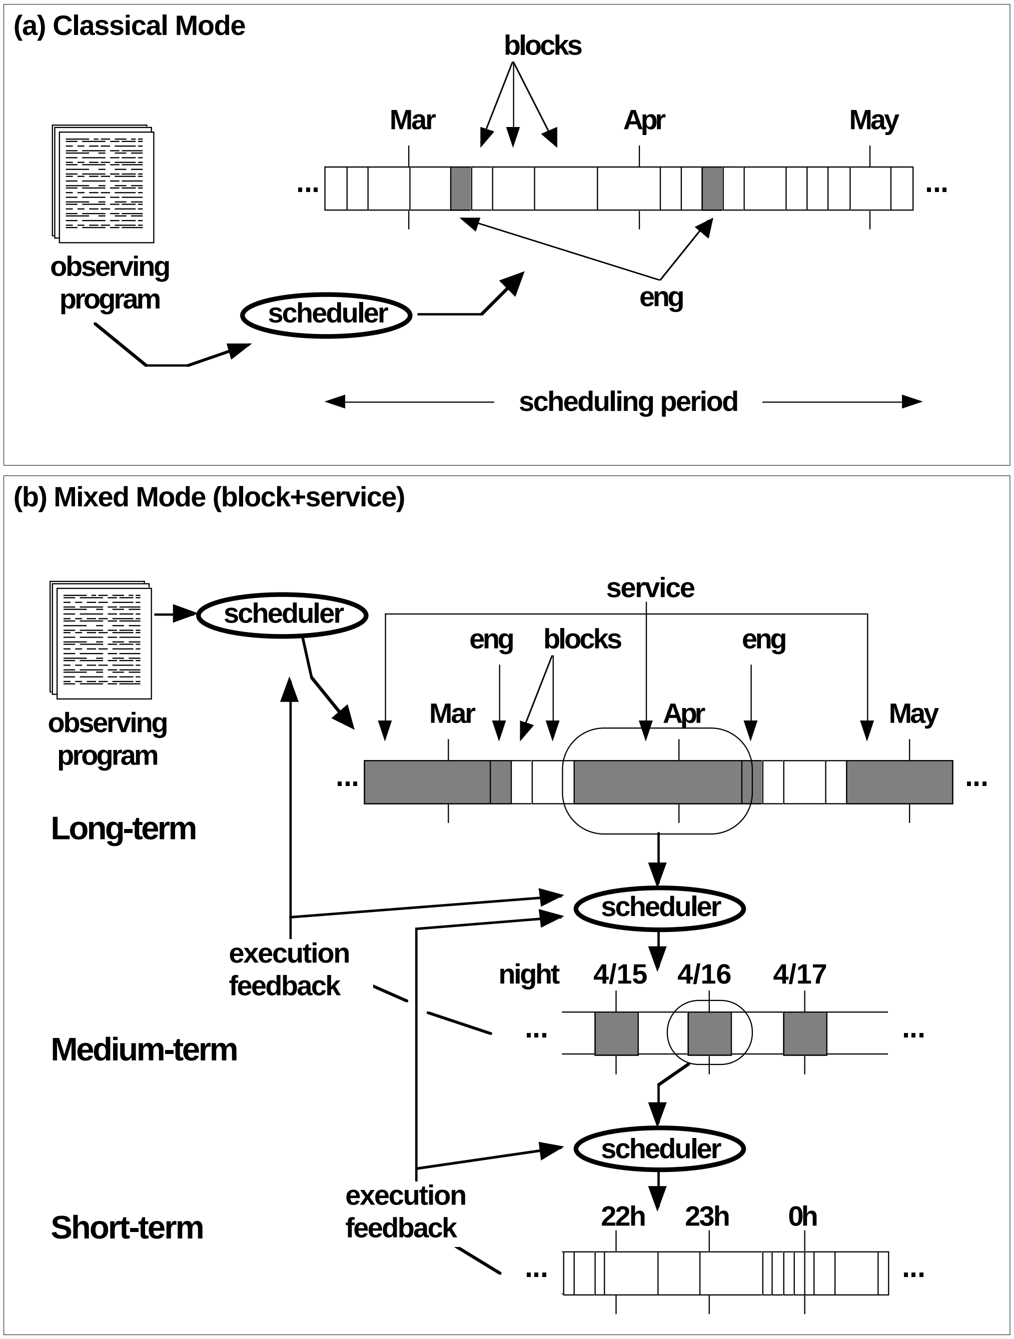

Figure 2. Two possible scenarios for introducing automated scheduling into ground-based observatory operations: (a) time allocation in a "classical" observing mode, and (b) scheduling at three increasingly fine levels of time granularity.

it allowed users to make some key scientific decisions about their observations without having to iterate over and over with Space Telescope Science Institute staff.

4.4. Distributed Scheduling

One area which offers great future potential is that of distributed scheduling. This has advantages not only in terms of performance, but also for new capabilities such as parallel contingency scheduling.

Performance gains can be very dramatic. Figure 3a shows the architecture of a master-slave configuration using an extended version of Spike. The controller (client) partitions the scheduling problem, distributes it to the servers, receives the results, and develops and implements a strategy to resolve remaining conflicts. This system has been prototyped and evaluated (Figure 3b) on a 12-workstation local network of Sun SparcStations (one controller and 11 servers). When compared with the results of scheduling the same problem on a single workstation, there was a ×4 speedup, with only a slight reduction in solution quality. While but a prototype, this illustrates the potential of distributed processing for this kind of problem. It is a particularly good way to make effective use of times when idle workstations are available.

The prospect for parallel contingency scheduling is illustrated schematically in Figure 4. The user would interact with the "master" system (controller client), which would farm out portions of the work to a set of servers. In this architecture the servers would make different assumptions about future conditions, and then construct and iterate on optimized schedules to correspond to those conditions. For example, one server could make the assumption that conditions will remain as they are. Others could assess the implications of changes to cloud conditions and seeing, hypothesizing that they are either improving or getting worse. Still others could evaluate scheduling options should the current activity finish early or late. All of these options could be brought together to the user via the master to present in a comprehensible way. The user could then, on very short notice, make a decision to purse a contingency plan when appropriate conditions develop. Likewise, exploring the "what-if" ramifications of a particular decision would be fast and straightforward.

5. Recommendations for the Future

In considering the future of astronomical scheduling applications, there are three recommendations that can be put forward:

1. **Develop a multi-Observatory interchange format for scheduling data**

   The advantages of such a format would include: facilitating planning and scheduling of coordinated and multi-wavelength campaigns, and comparison of scheduling approaches on the same problem data sets. There are several starting points for developing such a format, including ATIS (used for controlling automatic photometric telescopes), the Model Interchange Format proposed for Space Station, and some internal formats developed at Space Telescope Science Institute for exchanging scheduling data among

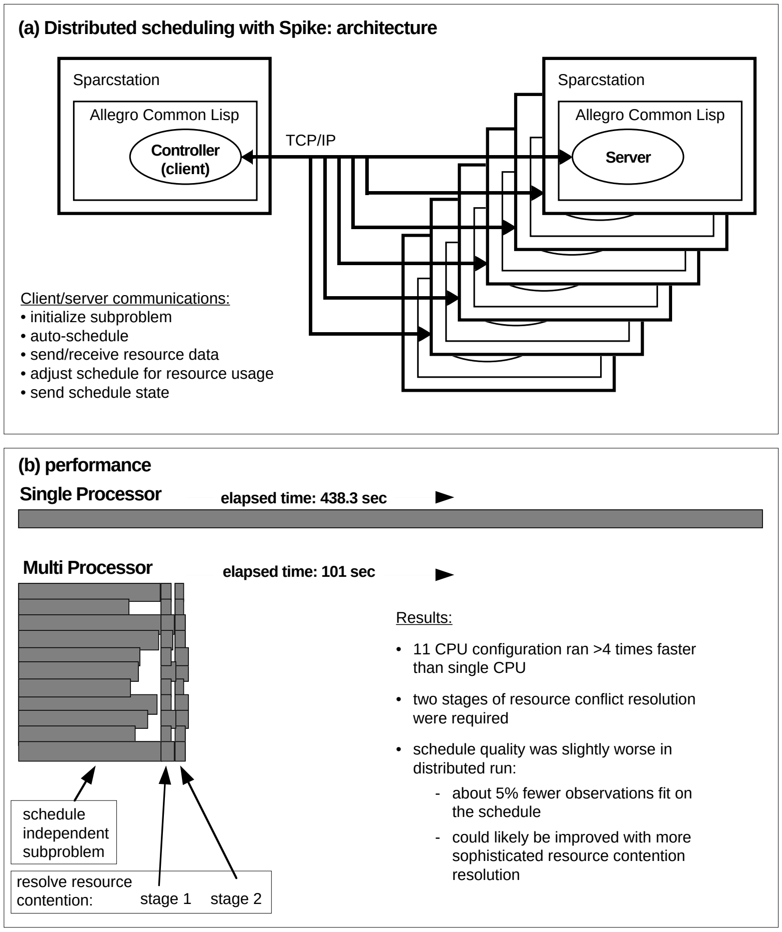

Client/server communications:

- initialize subproblem
- auto-schedule
- send/receive resource data
- adjust schedule for resource usage
- send schedule state

Results:

- 11 CPU configuration ran >4 times faster than single CPU
- two stages of resource conflict resolution were required
- schedule quality was slightly worse in distributed run:
  - about 5% fewer observations fit on the schedule
  - could likely be improved with more sophisticated resource contention resolution

Figure 3. Distributed scheduling with Spike for augmenting performance: (a) client-server architecture (b) results of a 12-workstation experiment compared to the same problem on a single workstation.

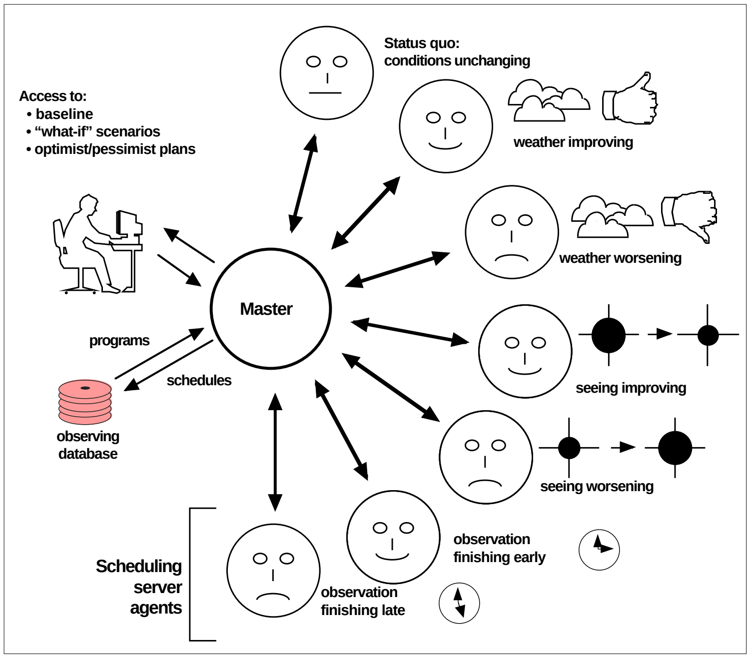

Figure 4. Distributed scheduling for contingency: a single client can control multiple servers, each considering variant future schedules.

the HST ground system elements. It is important that the format accommodate the needs of space missions, since multi-wavelength campaigns will certainly involve them in a crucial way.

Like FITS, this should be a transport format, not necessarily the internal format of any one scheduling system. With careful design, it should be possible to allow format extensions that are upwardly compatible, and for systems to ignore those parts of the language that they do not handle.

2. **Publish benchmark problem sets and solutions for realistic observing scenarios**

   A set of commonly-used benchmarks would allow schedulers to assess different approaches on realistic problems, test systems for reasonable behavior, and evaluate system improvements. It would also encourage groups in computer science or operations research to apply new techniques to astronomical scheduling problems, thus providing a fertile source of new ideas and techniques. Use of a set of common benchmarks would obviously be made much easier with agreement on a common format as suggested above.

3. **Provide more information to other sites about methodology, approaches, tools, results, ...**

   Workshops such as this one make clear that there is a great deal of work going on at many sites that would be of interest to others. While there is no way to substitute for face-to-face introductions and interactions, it would be very valuable for each group to spend some effort to inform the rest of the community of their problems, progress, and plans. The WWW provides an ideal way to publish this information once it has been collected. At STScI we are publishing a new page http://www.stsci.edu/spike, and we encourage others to do likewise.

Acknowledgments. The author is grateful for the hospitality of the European Southern Observatory where some of this work was done. Space Telescope Science Institute is operated by the Association of Universities for Research in Astronomy for the National Aeronautics and Space Administration.

References

Chavan, A. M., Johnston, M, and Albrecht, M. A. 1995, "Building Scheduling Tools for Ground Based Telescopes" (this conference).

Johnston, M. D. and Minton, S. 1994. "Analyzing a Heuristic Strategy for Constraint-Satisfaction and Scheduling" in Intelligent Scheduling, ed. M. Fox and M. Zweben (Morgan Kaufmann: San Mateo), 257-289.

Johnston, M. and Miller, G. 1994. "Spike: Intelligent Scheduling of Hubble Space Telescope Observations" in Intelligent Scheduling, ed. M. Fox and M. Zweben (Morgan Kaufmann: San Mateo), 391-422.

Miller, G. 1995a, "A System for Long-Term Scheduling of Ground-Based Observatories" (this conference).

Miller, G. 1995b, "Hubble Space Telescope Planning and Scheduling — Experience, Lessons Learned and Future Directions" (this conference).