



SPACE
TELESCOPE
SCIENCE
INSTITUTE

PREPRINT SERIES

No.306

AUTOMATED TELESCOPE SCHEDULING

Mark D. Johnston

August 1988

SPACE TELESCOPE SCIENCE INSTITUTE
3700 San Martin Drive Baltimore, MD 21218

AUTOMATED TELESCOPE SCHEDULING

Mark D. Johnston

Space Telescope Science Institute
3700 San Martin Drive
Baltimore, MD 21218

Published in the proceedings of the
Symposium on Coordination of Observational Projects

Strasbourg, November 1987

Cambridge University Press

AUTOMATED TELESCOPE SCHEDULING

Mark D. Johnston

Space Telescope Science Institute, Baltimore MD 21218, USA

Abstract: With the ever increasing level of automation of astronomical telescopes the benefits and feasibility of automated planning and scheduling are becoming more apparent. Improved efficiency and increased overall telescope utilization are the most obvious goals. Automated scheduling at some level has been done for several satellite observatories, but the requirements on these systems were much less stringent than on modern ground or satellite observatories. The scheduling problem is particularly acute for Hubble Space Telescope: virtually all observations must be planned in excruciating detail weeks to months in advance. Space Telescope Science Institute has recently made significant progress on the scheduling problem by exploiting state-of-the-art "artificial intelligence" software technology. What is especially interesting is that this effort has already yielded software that is well-suited to scheduling groundbased telescopes, including the problem of optimizing the coordinated scheduling of more than one telescope.

INTRODUCTION

Telescope observing time is a scarce resource: oversubscription by factors of several are typical. This basic fact will not change with the construction of newer and more advanced facilities. It is thus important to consider how the utilization of existing and planned telescopes can be increased to the maximum extent possible. By giving early consideration to these issues it will be possible to incorporate significant improvements in the next generation of astronomical observatories.

This paper focuses on one area which has generally received little attention for ground-based telescopes: automated telescope scheduling. For those observing programs which do not require the physical or remote presence of the observer ("absentee" or "service" observing), significant efficiency improvements can be expected by optimizing the observation sequence in response to changing environmental conditions, instrument availability, and other factors. This mode of observing offers the possibility of interleaving programs to optimize telescope utilization while at the same time executing observations under their most advantageous conditions. Exploiting this possibility will require the availability of automated scheduling tools for use by the telescope operations staff.

Automated scheduling at some level has been done for several satellite observatories, but the requirements on these systems were much less stringent than on modern ground or satellite observatories. Hubble Space Telescope (HST) is a case in point: essentially all observations must be pre-planned in detail weeks to months in advance. Planning over timescales of a year or more is also required in order to ensure that the shorter term schedules are mutually consistent. Because of certain (but hopefully not too frequent) disruptions to a schedule in progress it will also be necessary to re-plan on all timescales with very little notice. These factors, plus the complexity of HST observing

constraints implied by its low earth orbit, and the sheer number of observations to plan, make the HST scheduling problem a difficult one indeed.

Recent work at Space Telescope Science Institute (STScI) has addressed these issues by initiating a project to devise automated scheduling tools for use by the HST planning staff. This project exploits "artificial intelligence" (AI) software development and implementation techniques to provide a powerful framework for representing the scheduling problem and the strategies for its solution. The results to date of this effort have been very encouraging, as described further below.

From one perspective, the HST scheduling problem can be viewed as a setting for many of the same issues that must be addressed in scheduling groundbased telescopes. The HST constraints are more numerous and complex and the timescales are shorter, but the general nature of the problem is very similar. It has become clear that the scheduling tools being developed at STScI can be applied to other telescope scheduling problems, in particular to other satellite and groundbased observatories, and to the problem of planning joint or cooperative observing programs among observatories. It also appears likely that planning ambitious but complex multi-observatory projects would be facilitated by making planning tools of this type widely available to astronomers who wish to use them.

SPACE TELESCOPE PLANNING AND SCHEDULING

Hubble Space Telescope [1] is currently scheduled for launch by the Space Shuttle in mid-1989. Following an initial checkout period it will be made available for use by astronomers around the world. HST has a full complement of scientific instruments: two imaging cameras, two spectrographs, and a high-speed photometer; the guidance system can also be used for astrometry. Each instrument has many possible observing configurations and modes that are selectable by command.

Astronomers request the use of HST by submitting observing proposals in response to a yearly Call for Proposals [2,3]. After proposals are peer-reviewed and selected, the detailed specifications of the exposures are entered into a database. These specifications include target and exposure descriptions along with a potentially large number of configuration parameters that specify how the data is to be taken. For a typical one-year observing program several tens of thousands of exposures are expected to be defined.

There are a variety of properties of and relationships among these exposures that may be specified by the proposer. Their relative order and time separation may be important. Some exposures are designated as calibrations or target acquisitions for others. Some must be executed at specific times, or at specific phases in the case of periodic phenomena. Some are especially sensitive to stray or scattered light. Exposure durations may vary depending on background light intensity. Some exposures must be executed without interruption while others can be broken up as needed. In some cases a specific orientation of an instrument aperture is required. Some exposures are conditional on the results of other exposures.

In addition to proposer-specified constraints, there are a large number of other constraints that must be considered when scheduling HST operations [4]. These range from "strict" constraints that cannot be violated under any circumstances, to "good operating practices" that represent scheduling goals. HST is not allowed to point closer

than 50° to the sun and 15° to the bright moon. Slewing the telescope is relatively slow (90° in $\sim 15^m$) so it is important to minimize the time spent in maneuvers. Many constraints are a direct result of HST's low orbital altitude (500 km) and consequent 95^m orbital period. A typical target is occulted by the earth for $\sim 40^m$ of each orbit. Up to half of the orbits in a day are contaminated for up to $\sim 20^m$ by HST's passage through the South Atlantic Anomaly, a high particle density region during which data cannot be collected. Scattered earthlight changes dramatically over the course of an orbit. HST communications with the ground is via the two geosynchronous Tracking and Data Relay Satellites (TDRS), which are visible for only part of each orbit and are shared with other users. Use of TDRS must be requested and scheduled ahead of time, and the TDRS schedule can be disrupted at the request of a high-priority user such as Space Shuttle.

As a result of these and other constraints the operation of HST is almost entirely pre-planned. The most extended plans must ensure that the overall HST scientific program is feasible and balanced. Shorter term schedules are used to lay out detailed observation sequences and make requests for TDRS communications contacts. The most detailed schedules specify spacecraft and instrument commands and the times they are to be executed. Command loads for the onboard computers are prepared and transmitted several times a day.

There are limited provisions for real-time interaction by the observer: small configuration changes (e.g. filters) may be requested as a result of the analysis of acquisition exposures. Pre-planned branching sequences are also allowed, wherein several alternative command sequences are transmitted to HST's onboard computers. At the specified decision point the observer chooses one of the alternatives.

Many constraints are accurately predictable, but the scheduling of HST is complicated by some factors that are intrinsically uncertain. The location of HST in orbit is extrapolated using an orbit model which is fit to past position and velocity measurements: this extrapolation is not valid for more than a few months into the future. The existence of suitable guide stars for pointing control can only be estimated: failure to acquire guide stars can be the result of stellar variability or of binaries that cannot be resolved from the ground.

As a consequence of all of these factors, the optimal scheduling of HST is an extremely complex problem. With tens of thousands of observations each year, constraints that range from policy statements to orbital events, and the potential for having carefully prepared schedules completely disrupted by unexpected events, it is clear that even the best human schedulers will need extensive software support to accomplish their task.

HST SCHEDULING SOFTWARE

The scheduling software currently available at STScI is known as the Science Planning and Scheduling System (SPSS) and is part of the Science Operations Ground System (SOGS) developed by TRW. SPSS is a large "classical" software system: it is written in FORTRAN for DEC VAX computers, and is intimately coupled to a central database of scheduling information derived from the observing proposals. While SPSS has been successfully used to generate detailed schedules of a few days duration, there are several factors that prevent its use on the long-range planning problem:

- SPSS scheduling algorithms only examine a few possible times to schedule exposures, and can therefore easily miss good scheduling opportunities
- SPSS always considers detailed orbital events and conditions, even when they are uncertain or unpredictable. This makes it computationally very expensive to construct and evaluate long-range plans
- a significant number of scheduling constraints are not considered by SPSS, and, because of the design and implementation of the system, they are difficult to add to the software
- the throughput of the overall system (people plus software) is and remains a serious concern

As a result of these problems, STScI initiated in early 1987 a project (SPIKE) to provide software support tools for long range planning and to help HST planners use SPSS more effectively [5]. It was an early decision in the SPIKE project to exploit artificial intelligence software development technology. This refers to a collection of software development techniques that have evolved in the course of computer science research as effective ways to represent and solve certain kinds of problems. These techniques have moved from the laboratory to commercial use as their effectiveness has been demonstrated. For the purposes of the SPIKE project, the most important of these techniques are: a language (Lisp) that is particularly appropriate for manipulating complex data structures and symbolic data; object oriented programming with inheritance and message passing; rule-based programming facilities; integrated graphics and window tools; and a rich development environment including an integrated editor, debugger, and data inspector, an online database of function cross-references and documentation, and the ability to develop in either interpreted or compiled modes. For SPIKE the advantages of using these techniques are primarily a rapid software development cycle, a concise but expressive representation of scheduling data, flexibility in the definition and modification of scheduling constraints, and the ability to incorporate a graphics-oriented user interface to help the planner understand and modify the schedule. In addition, the development and operation of SPIKE tools on single-user workstations has the advantage of adding no additional burden to the SPSS computers while giving the individual user an assured response.

The initial focus of the SPIKE planning tools is to make available to HST planners the ability to handle the long-range planning problem. In practice, this means the evaluation of exposures from the proposal database with the goal of finding the best times to schedule them. In this, the earliest phase of HST planning, a typical plan may cover a year or more, and it is required to place exposures on the plan within "windows" that are between a few days and a month in duration. "Best" in this context means that no predictable strict constraints are violated, that the probability of violating unpredictable constraints is minimized, that the time windows on exposures reflect the factors that make the window most favorable, and that resource usage over the plan is balanced.

SPIKE PLANNING TOOLS

Computer techniques for optimal scheduling have been investigated for a number of applications. Unfortunately the complexity of "real" scheduling problems has limited the applicability of much of this work to idealized cases or to model problems. In recent years several AI research efforts have considered scheduling as a domain where AI techniques can be fruitfully applied. Of particular interest in this area is the factory scheduling work of Smith et al. [6]. While the factory scheduling problem

shares a number of common features with telescope scheduling (most notably a similar set of precedence and efficiency constraints), there are some unique features of telescope scheduling that can be exploited to advantage. The most important of these is that target visibility and related constraints can be stated and predicted far in advance, and can therefore be used to limit the search for good schedules that would otherwise be overwhelming.

This is the approach adopted in the SPIKE planning tools project. Associated with each activity to be scheduled (typically an exposure or group of exposures) is a "suitability function", a function of time whose value represents how desirable it is to start an activity at a specified time. Suitability functions are derived from constraints, of which there are several types:

- 1) *Absolute constraints* represent factors that are independent of what is planned. They are typically used for celestial and orbital constraints which are predictable well in advance. Because of the uncertainty in orbit model extrapolation they are often statistical in character. For example, a constraint on an exposure that it must be taken when HST is pointing more than 5° from the earth limb, be in earth shadow, and have an uninterrupted view of the target for at least 30^m can be represented as the average amount of time these conditions are satisfied over some interval which encompasses many orbits. This would reflect the desirability of scheduling such an exposure in, say, a week in which a greater number of such viewing opportunities occur.
- 2) *Relative constraints* represent factors that depend on when one or more other activities are planned. An example of a constraint of this type is a precedence constraint: if activity A must precede activity B, and if there are any time constraints on A such that it has a known earliest end time, then the suitability function of B can quantify the fact that B cannot be started until after the earliest end of A. Two other important types of relative constraints are minimum and maximum time separations.
- 3) *Segment constraints* represent factors that apply to intervals of time (time segments) as opposed to activities. For example, they can represent the fact that the total exposure time or data volume that can be scheduled during a given time interval is limited to some maximum value.

The suitability function of an activity is the product of the suitability functions derived from all of its constraints: this mirrors an intuitive notion of how to combine different sources of support for planning an activity at a given time. Suitability functions are represented internally as piecewise constant functions. With this choice, combining suitability functions with arithmetic operators (e.g. multiply) again yields a function of the same form. It also has the advantage of computational efficiency.

Suitability functions represent scheduling possibilities and their degree of desirability. There must further be some way to represent choices of actual times to schedule activities. In SPIKE this is accomplished by iteratively breaking up the total planning interval (years or more) into "segments" of duration days to weeks. A schedule is therefore a consistent and complete (so far as possible) set of assignments of activities to time segments. Since each such assignment can change the suitabilities of other related activities via relative constraints, and each activity can typically be assigned to many possible segments, there still remains an enormous search problem in finding a consis-

tent, let alone optimal, set of assignments. It is nevertheless a much smaller problem than trying to schedule at a microscopic level from the beginning.

The search for a consistent set of assignments can be pursued with a variety of search strategies. Simple examples of strategies would be to schedule the highest priority activities in their best segments first, or to schedule the most constrained activities at their earliest possible times. More complex strategies could exploit simulated annealing or neural net optimization techniques [7]. The development of realistic strategies for HST will require the analysis of a large sample of real HST observing proposals, and the determination of overall measures of schedule "goodness" as a basis for comparing alternative partial schedules. This analysis is currently in progress.

The SPIKE tools are implemented in Common Lisp in conjunction with the Flavors object system and CommonWindows for graphics and user interaction. These choices allow the tools to be portable to many workstations. The development of the tools is taking place on Texas Instruments Explorer workstations, but it is possible that delivery will be on general purpose workstations such as those from Sun Microsystems or other vendors. The use of "expert system shells" such as KEE (Intellicorp, Inc.) or ART (Inference, Inc.) is currently under investigation for control and strategic decision support. Scheduling under the control of a rulebase has been demonstrated: for example, it is possible to write almost verbatim a rule of the form "if there is an unscheduled high priority activity which is highly constrained and related to activities already scheduled, then schedule it next". This will allow for easy incorporation of new scheduling strategies in the system as experience is gained with real scheduling problems.

An important aspect of the SPIKE system design is a graphics-oriented user interface. While much of the SPIKE processing is automatic, there will always be a need for diagnostic tools when scheduling problems arise, and for displays of schedules as they evolve. Since it is expected that new scheduling strategies will be developed as planners use the system, it is important that planners have visibility into the results of their work in order to effectively formulate these strategies

SCHEDULING OF GROUND BASED TELESCOPES

Little attention has been paid to automated scheduling of groundbased telescopes up to now, but this situation is changing. One of the participants at the ESO Very Large Telescope (VLT) Workshop [8] made the following remark:

"... it will become inevitable that optical astronomy allocate time completely differently than it has done hitherto. The fatalism of the weather has to be eliminated completely ... Good seeing has to be used and not just talked about and one has to manage one's telescope in such a way that it can be exploited to optimum advantage. That is one reason, why I think that absentee observing will become inevitable, and that, in the future, these very large expensive facilities will have at least two kinds of time allocations: priority 1 and priority 2. Priority 1 allocations will be guaranteed to be executed, barring disasters, under all circumstances. Priority 2 allocations will be executed if possible and feasible, and if not they'll be removed from the menu and have to be reapplied for. Then you can make quarterly schedules for you instruments, for your telescopes, which are weather dependent. You can then optimize the schedule by writing some fairly complex computer pro-

grams with possibilities of on-site human intervention to utilize your telescope to best advantage."

And from the VLT Proposal [9]:

"All too frequently at the moment someone travels to La Silla for a programme which requires excellent seeing to have only some mediocre nights, while another astronomer a week later experiences superior conditions which his programme does not need. With the VLT such wasteful procedures cannot be accepted."

While it is certainly true that some observing programs will require the physical presence of the observer, others can be executed by the telescope operations staff based on detailed specifications from the observer. This type of "absentee" observing has been successfully tried on several telescopes (see [10] for a discussion of the UK experiments). To achieve the maximum possible efficiency from this type of observing it is important to have automated scheduling tools available to help the telescope staff make and adapt their plans to changing telescope conditions.

The design of the SPIKE planning tools makes them well suited for this application. Because of the advantages obtained through the use of AI development techniques it has been possible to maintain a great deal of modularity and flexibility in the SPIKE system design. As a result it is straightforward to include new constraints for new problems. Some HST constraints are sufficiently general that they apply equally well to groundbased observing, such as precedence and time separation constraints among exposures. Others do not apply at all, and there are of course a number of new constraints that are relevant only for ground observations. For example, in groundbased scheduling it will be important to include constraints related to airmass, seeing, moon phase and visibility, etc. Given that these factors can be quantified, the SPIKE tools provide a framework in which they can be utilized to help optimize telescope scheduling.

In a similar fashion, constraints from more than one observatory can easily be incorporated in the system at the same time. For example, suppose a simultaneous observing program on an object were contemplated in which several satellite and groundbased observatories were to be involved. It would be straightforward to use the SPIKE tools to find optimal observing times based on the specific constraints of each observatory. This could greatly simplify the planning and proposal preparation phases of ambitious multi-wavelength projects.

There are, however, some issues that must be addressed before automated scheduling can be widely utilized. Optimal scheduling means that there must be a pool of observations upon which may not be executed (referred to above as Priority 1 and Priority 2). It must be accepted that, in order to optimize the overall observing program, some individual programs may not be executed. STScI has allowed for this by accepting proposals at a priority level of "supplemental". Another issue is the level of detail required of proposers in order for automated scheduling to be effective. For scheduling software to exploit preferences and constraints it must know about them, which means that they must be specified by the proposer at some level, or be derivable from the proposal specifications. STScI has simplified this process by providing a remote proposal submission system which accepts machine-readable observing proposals. The proposal submission software makes a large number of checks for completeness and consistency of the proposal specifications. Facilities of this type are likely to be necessary for

groundbased telescopes which plan to make extensive use of automated scheduling capabilities.

CONCLUSIONS

The experience of HST in the domain of automated telescope scheduling makes it clear that it is a technology that can be fruitfully applied to ground-based and multi-observatory scheduling. The use of artificial intelligence software technology makes it straightforward to adapt the HST planning tools to other telescope scheduling problems. For the next generation of astronomical observatories, now in the design stage, automated scheduling offers a significant potential for increases in observing efficiency and telescope utilization.

REFERENCES

1. *The Space Telescope Observatory*, 1982, NASA CP-2244, ed. D. Hall
2. *Call for Proposals*, October 1985, Space Telescope Science Institute
3. *Proposal Instructions*, October 1985, Space Telescope Science Institute
4. *Design Reference Mission*, 1982, Sherrill, T.J., Lockheed Missiles and Space Company, Inc.
5. "Expert System Tools for Hubble Space Telescope Observation Scheduling," 1987, Miller, G., Rosenthal, D., Cohen, W., and Johnston, M., in *Proceedings of the 1987 Goddard Conference on Space Applications of Artificial Intelligence*
6. "Constructing and Maintaining Detailed Construction Plans," 1986, Smith, S., Fox, M., and Ow, P., *AI Magazine*, Fall 1986, p. 45
7. "Optimization Algorithms: Simulated Annealing and Neural Network Processing," 1987, Jeffrey, W., and Rossner, R., *Ap. J.* 310, 473
8. H. Van der Laan in *Very Large Telescopes, their Instrumentation and Programs*, IAU Colloq. No. 79, April 1984, ESO, Garching
9. *Proposal for the Construction of the 16m Very Large Telescope*, 1987, ESO, Garching
10. "The UK Remote and Service Observing Programmes," 1986, Longair, M.S., Stewart, J.M., and Williams, P.M., *Quarterly Journal of the Royal Astronomical Society* 27, 153

Acknowledgements: The author is grateful to G. Miller, J. Sponsler, and S. Vink (STScI) and D. Rosenthal (NASA Ames Research Center) for many useful discussions, and for the hospitality and support of the Space Telescope European Coordinating Facility, ESO, Garching, where some of this work was conducted.

Space Telescope Science Institute is operated by the Association of Universities for Research in Astronomy for the National Aeronautics and Space Administration.