Telescope Network Scheduling
Rationale and Formalisms

Sotiria Lampoudi[1] and Eric Saunders[2]
[1]*Computer Science Department, University of California, Santa Barbara, CA 93106, U.S.A.*
[2]*Las Cumbres Observatory Global Telescope Network, 6740 Cortona Dr., Suite 102, Goleta, CA 93117, U.S.A.*

Keywords: Scheduling, Astronomy.

Abstract: Scheduling of university and institutional telescopes is typically performed manually, and astronomers are used to interacting with a human to explain their requirements for resources and time. Las Cumbres Observatory Global Telescope (LCOGT) is deploying a worldwide network of robotic telescopes. At LCOGT manual scheduling is infeasible due to: 1) the number of resources and observations that must be scheduled, 2) the scheduling-time dependencies that arise when concurrent or consecutive access to telescopes is required, and 3) the need to rapidly re-calculate the schedule to accommodate near-real-time requests from high priority observing programmes (e.g. transient followup programmes), and changing resource availability due to weather and other reasons. In this paper we develop a formalism capable of expressing the complex requirements and preferences of astronomers concerning resource and time allocation on a telescope network, and formulate the offline *telescope network scheduling problem* as the problem of choosing and scheduling (i.e. assigning concrete start and end times to) a maximum priority, non-overlapping subset of an input list of requests.

1 INTRODUCTION

Robotic telescope networks are gaining traction (Hessman, 2006), due to the novel astronomical observation modes they enable, and their lower operational and administrative costs. These resource pools can be homogeneous (Vestrand et al., 2002; Akerlof et al., 2003) or heterogeneous (Nather et al., 1990), and designed from the ground up or federated, much like computational clusters and grids. However, because astronomical observations must be performed at specific times, rather than at the first opportunity, scheduling telescope networks is a form of advance reservation scheduling that is more challenging than the batch scheduling usually found on computational clusters and grids.

Las Cumbres Observatory Global Telescope (LCOGT) has acquired two 2m robotic telescopes (Faulkes North and South), and is additionally deploying a worldwide network of 0.4m and 1m optical telescopes and attached imagers and spectrographs, dedicated to time-domain astronomy (Martinez et al., 2010).

Scheduling of individual, non-robotic university and institutional telescopes is typically performed manually. But manual scheduling is infeasible on the LCOGT network for a number of reasons. Primary among them is the sheer number of resources and observations that must be scheduled. Secondary is the desire to encourage requests that take advantage of features unique to a spatially distributed telescope network. Such requests are for concurrent or consecutive use of resources, and they introduce scheduling-time dependencies, making manual scheduling more difficult. Finally, it is desirable that schedule (re)calculation be fast, to accommodate near-real-time observing requests (e.g. for transient followup (Brown et al., 2007)) with minimal disruption to long-standing requests, and rapid changes in resource availability due to weather or technical issues.

In this paper we survey the literature on telescope scheduling, define a formalism that allows astronomers to express their complex requirements and preferences regarding resource and time allocation on a telescope network, formulate the *multiple telescope scheduling problem*, and briefly discuss its attributes.

2 RELATED WORK AND MOTIVATION

The field of telescope scheduling has not evolved sys-

Lampoudi S. and Saunders E..
Telescope Network Scheduling - Rationale and Formalisms.
DOI: 10.5220/0004274203130317
In *Proceedings of the 2nd International Conference on Operations Research and Enterprise Systems* (ICORES-2013), pages 313-317
ISBN: 978-989-8565-40-2

tematically; the majority of publications report on attempts to engineer real-world scheduling systems for production use, and therefore contain little theoretical justification of their design choices, almost no theoretical classification of their problems and solution, and few attempts to generalize them into forms amenable to theoretical exploration. These are omissions we seek to remedy in this paper.

It is helpful to categorize the literature on telescope scheduling according to whether it concerns scheduling requests from single or multiple "users". In this context the term "user" really refers to a scientific objective, rather than an individual. A single user can be an astronomer, a group of astronomers sharing a single science goal, or a computational agent with a single science goal.

In the case of a single user, although the astronomical literature refers to a "scheduling" problem, what is being studied bears no similarity to, for example, computational resource scheduling. In single-user "scheduling" the objective is the optimal choice of targets to observe (Colome et al., 2010; Bakos et al., 2011, for example) and/or times during which to observe them (Saunders et al., 2006, for example), in order to maximize scientific return, under the assumption of perfect resource availability and with no concern for contention. This optimisation goal can only be achieved with very detailed domain-specific, astrophysical information. For the purposes of this paper we will refer to this as an *observation selection* problem, not a scheduling problem.

In multi-user telescope scheduling we have a problem that bears a closer resemblance to computational resource scheduling. The resource pool is time- and/or space- shared, and the contention between user requests must be resolved in a way that optimises objectives and/or satisfies constraints. Each individual user's input to the multi-user scheduler may be the output of a single-user observation selection step. Multi-user scheduling can concern a single resource (Kinzel, 2010, for example), or many, as in the LCOGT network described in this paper.

In telescope scheduling practice the observation selection problem and the multi-user scheduling problem have historically been entangled. This situation is likely due to the fact that traditionally astronomers will sit down with a human scheduler and manually produce schedules that take into account both scientific value and contention at once. The Robonet I (Fraser, 2006) and II (Tsapras et al., 2009) schedulers are cases in point: they attempt to optimise target selection (for a single user, whose science goal is transient followup) using complex scoring functions (Snodgrass et al., 2008), but are tightly integrated with the RCS scheduler, which is responsible for all scheduling of the Faulkes telescopes, and which is also used by all the other users of those telescopes. This system was acquired by LCOGT along with the Faulkes telescopes, and its management requires human intervention to resolve conflicts between automated (Robonet), professional astronomer and educational user groups.

A step in the direction of disentangling single- (observation selection) and multi- user scheduling was the eStar meta-scheduler (Allan et al., 2006), an agent-based scheduler acting on behalf of a user (after the observation selection step) to submit requests for resources to any telescopes and telescope networks that were compatible with it. However, as eStar depended on the resources to simply accept or reject its requests according to their internal criteria, and prioritised among those accepting its requests using its own unrelated criteria, little theoretical argument could be made regarding its tendency to achieve objectives, satisfy constraints, or otherwise optimise the scheduling process. No empirical studies addressing these questions were published.

Our approach conceptually disentangles the single-user (observation selection) and multi-user parts of the telescope scheduling problem by factoring them respectively into a formalism that hides the astronomy-specific aspects of the problem, and an optimisation problem definition based on that formalism.

3 CONCEPTS AND FORMALISMS

3.1 Manual vs Automatic Scheduling

The process of scheduling university and institutional telescopes typically goes as follows. A Time Allocation Committee (TAC) is responsible for allocating to each scientific project a total amount of time to be used over a semester on their resources, and for assigning a numerical priority to each project. Before the semester begins, individual investigators contact a human scheduler and, in English, request exclusive use of particular resources for particular intervals. They usually assign their requests their own, usually three tier, level of priority ("High, Medium, Low"). The human scheduler then resolves scheduling conflicts by taking into account a subjectively weighted mix of the conflicting projects' TAC-assigned priorities, the investigators' self-assigned priority for the particular observation, and the urgency (expressed in terms of expectation of future opportunities to perform an observation) and importance (expressed as number of past observations that become useless if this observation is not performed) of the requests, if those are available. Observing programmes requiring near-real-time access to resources are allocated Target of Opportunity (ToO) time on the resources, allowing them to manually pre-empt ongoing observations, and assume control. The human scheduler patches up the schedule in response to disruption from ToO requests and weather, to whatever extent possible.

In the following section we develop an "assembly language" for specifying users' resource and time requirements. Higher level specifications of users' needs can be "compiled" into this language by interfaces, and automated systems can encode their needs directly in it.

3.2 The Reservation Formalism

At the core of advance scheduling of all types is the concept of a *reservation*, a representation of a request for exclusive use of a resource at a particular time. In this paper reservation with a lowercase 'r' will refer to the concept, whereas Reservation with an uppercase 'R' will refer to a formalism, or data structure.

Formally, a Reservation is a 4-tuple (d, p, t, W), where:

- d is the duration of the exclusive use of the resource,
- p is the priority,
- t is the resource, and
- W is a list of "windows of opportunity", $\{w_1, \ldots, w_i\}$, where $w_j \geq d, \forall j \in [1, i]$.

Windows of opportunity give the times during which the observation can occur; that is, the concrete start and end times of the Reservation must both fit within a single window of opportunity.

In the vocabulary of astronomy, a Reservation is a request by a project with priority p to use telescope t for duration d during one of the windows listed in W.

To explain the design considerations that led us to adopt this particular formalism, let us consider two typical examples of requests made via email to the human scheduler of the Faulkes North telescope in the first semester of 2011.

In both examples the first column gives the date, the second the start time and the third the end time. A different window is specified on each line. Note that an informal English sentence further specifying the intended meaning of each request was necessary. Note also – and we will not dwell on this aspect of transitioning from human language to automated interfaces, although it is interesting and possibly important to human computer interface design – that in the first example, "or two" is redundant and can be omitted without loss of information, since, if one night is sufficient, two are clearly also sufficient.

Example 1. "[For this observation] one or two of the nights should be sufficient."

```
11-05-01 05:30 15:30
11-04-27 05:30 15:30
11-05-03 05:30 15:30
```

The next request invents a notation for a date range. This range implicitly contains five windows of opportunity (on consecutive days between May 1st and 5th). The investigator also wishes to include so-called "slack", i.e. some flexibility in the start and end times, but that has to be expressed in English.

Example 2. "[This observation] can also start (and end) up to 2 hours later if that fits better."

```
2011-05-01<->2011-05-05  06:30   11:00
```

The Reservation formalism allows us to compactly capture these two typical request scenarios, which previously had to be expressed in a mix of machine-readable notation and English. Both examples illustrate the need for multiple windows of opportunity. Example 2 illustrates the additional need to accommodate slack in the windows, which we meet by including a duration field that is separate from the start and end of the windows.

3.3 The Compound Reservation Formalism

Intuition and experience suggest that the astronomical observations being requested by a particular project are usually not independent of each other, since they are components of a larger scientific programme. Aside from having a common target, requests from the same project are often linked to each other in ways that are more complicated. In what follows we describe two types of "scheduling-time dependencies", so called because they can be resolved at the time the schedule is being assembled.

The first kind of scheduling-time dependency is one in which the *scheduling status* of a request ("did it get scheduled or not?") depends on the scheduling status of the other requests in the dependency. For example, it can be the case that two observations are useful only if both of them are successfully completed; if only one can be scheduled, then it will be of no value. Or it can be the case that the same data can be obtained using either of two different telescopes, a logical disjunction between observations that cannot be captured within a single Reservation, as defined above.

The second type of scheduling-time dependency is one in which the *times* of observations participating in the dependency are related by some formula. This is a use case commonly known in astronomy as a "time series with cadence", and it is used when the timing of the first observation is flexible, but the gaps between subsequent observations must obey some mathematical rule. The Reservation formalism as defined above, given its fixed windows of opportunity, cannot help encode such a dependency. Our practical way of dealing with this limitation is by allowing astronomers to request such a dependency at a higher level, and then "compiling" it into the fixed windows of Reservations after making some (arbitrary or intelligent) constraining decisions about the position of one of the observations in the dependency.

Scheduling-time dependencies are distinct from another class of dependencies that we do not address in this paper, and which cannot be captured by our formalisms, namely those for which the decision of whether to schedule an observation participating in a dependency hinges on some *post-completion* function (not the mere scheduling status) of its antecedents.

To take full advantage of the fact that a scientific programme can include the first type of scheduling-time dependencies between observations, we define a second formalism, the Compound Reservation. A Compound Reservation is a logical sentence containing Reservations connected by a logical operator. The two logical operators that we have found relevant to the telescope network scheduling problem are AND and ONE-OF. Luckily, most astronomy use-cases are covered by Compound Reservations that include only one level of logical operators, although, with minor changes to the model, multiple levels of nesting are possible.

We call a Compound Reservation containing a single Reservation and no operators a *singleton*.

The AND operator in our formalism is the traditional conjunction operator. The Compound Reservation (r_1 AND r_2) means simply that either both r_1 and r_2 should be scheduled, or neither should be scheduled. (r_1 AND r_2 AND … AND r_i) is also defined as one would expect.

The ONE-OF operator in our formalism is equivalent to a "one-hot circuit" in digital circuit design. (r_1 ONE-OF r_2) means that either r_1 or r_2 should be scheduled, but not both. For two arguments ONE-OF is equivalent to XOR. The reason we use the notation ONE-OF rather than XOR is that most implementations of an XOR for more than two arguments yield a parity checker (by chaining many two-argument XOR gates). A parity checker is a circuit that evaluates to True when an odd number of its arguments are True, whereas what we want is a circuit that evaluates to True when *exactly one* of its arguments is True. Since the term "one-hot" is not commonly used outside digital design, we use the more intuitive label ONE-OF for this operator.

Note that independent Compound Reservations, which we join in a list by separating them with commas, i.e. R_1, R_2, R_3, are implicitly linked by a *regular* logical OR, since any non-overlapping subset of them may be scheduled.

Some examples of the kinds of astronomical observation requests that the Compound Reservation formalism enables are:

- observing a target from one of multiple alternative telescopes (using ONE-OF)
- generic time-series of observations (using AND)
- concurrent observation of the same target from multiple locations or using multiple instruments (using AND)
- tracking a stationary target in spite of the earth's rotation, or a moving target, using a succession of telescopes (using AND)

4 THE TELESCOPE NETWORK SCHEDULING PROBLEM

We define the *telescope network scheduling problem* as the following: given a list of Compound Reservations $\{R_1,\ldots,R_i\}$, produce a consistent schedule, i.e. one that satisfies the logical constraints and contains no overlaps, that maximises the sum of the priorities of the scheduled Reservations.

Similar problems have been formulated in operations research and computer science (e.g. truck scheduling (Lee et al., 2012) and satellite scheduling (Barbulescu et al., 2004; Frank et al., 2001)). The telescope network scheduling problem combines four attributes of previously formulated scheduling problems. First, it concerns *interval scheduling*, i.e. scheduling events that must occur at particular points in time, in a non-overlapping fashion. Second, it allows *slack*, that is, flexibility in the start and end times. This is captured in the Reservation formalism by allowing windows of opportunity to be of greater length than the separately defined duration of a Reservation. Third, it is a *multi-resource problem*. Multiple resources may be available concurrently, but they are not interchangeable due to their distinct locations.

Thus, the resource being requested in a Reservation must always be specified. Fourth, through the use of ANDs and ONE-OFs, the multiple telescope scheduling problem contains one type of logical *scheduling-time dependencies* between reservations.

For the purposes of the LCOGT scheduler, we are interested in the offline version of the telescope network scheduling problem, that is, the version in which the full list of Compound Reservations is provided at the outset. Thus, we seek algorithms that take as input a list of Compound Reservations and return one or more consistent schedules that maximise the sum priority of the scheduled Reservations. In practice, such an algorithm will be wrapped in a control algorithm that triggers the re-calculation of the schedule in response to observation failure, resource availability changes, requests for near-real-time observations and other user actions, which can be thought of as changes to the input list.

5 SUMMARY

In this paper we have described a formalism that allows astronomers to express their complex scientific needs and preferences regarding resource and time allocation on a telescope network. This formalism serves as the assembly language into which even higher-level descriptions of astronomers' needs can be translated.

REFERENCES

Akerlof, C. W., Kehoe, R. L., McKay, T. A., Rykoff, E. S., Smith, D. A., Casperson, D. E., McGowan, K. E., Vestrand, W. T., Wozniak, P. R., Wren, J. A., Ashley, M. C. B., Phillips, M. A., Marshall, S. L., Epps, H. W., and Schier, J. A. (2003). The ROTSE-III Robotic Telescope System. *Publications of The Astronomical Society of the Pacific*, 115:132–140.

Allan, A., Naylor, T., and Saunders, E. S. (2006). The eSTAR network agent architectures for astronomy. *Astronomische Nachrichten*, 327(8):767–770.

Bakos, G. Á., Hartman, J. D., Torres, G., Kovács, G., Noyes, R. W., Latham, D. W., Sasselov, D. D., and Béky, B. (2011). Planets from the HATNet project. In *European Physical Journal Web of Conferences*, volume 11 of *European Physical Journal Web of Conferences*, page 1002.

Barbulescu, L., paul Watson, J., Whitley, L. D., and Howe, A. E. (2004). Scheduling space-ground communications for the air force satellite control network. *Journal of Scheduling*, 7:7–34.

Brown, T., Rosing, W. E., Baliber, N., Hidas, M., and Street, R. (2007). Surveys, Temporal Variability, and the Las Cumbres Observatory Global Telescope. In *American Astronomical Society Meeting Abstracts #210*, volume 39 of *Bulletin of the American Astronomical Society*, page 173.

Colome, J., Casteels, K., Ribas, I., and Francisco, X. (2010). The TJO-OAdM Robotic Observatory: the scheduler. volume 7740, page 77403K. SPIE.

Frank, J., Jónsson, A., Morris, R., and Smith, D. E. (2001). Planning and scheduling for fleets of earth observing satellites. In *in: Proceedings of Sixth Int. Symp. on Artificial Intelligence, Robotics, Automation and Space*.

Fraser, S. N. (2006). Scheduling for robonet-1 homogenous telescope network. *Astronomische Nachrichten*, 327(8):779–782.

Hessman, F. V. (2006). Prospects for a global Heterogeneous Telescope Network (HTN). *Astronomische Nachrichten*, 327:763.

Kinzel, W. M. (2010). Jwst planning and scheduling operations and concepts. volume 7737, page 77370Y. SPIE.

Lee, S., Turner, J., Daskin, M. S., de Mello, T. H., and Smilowitz, K. (2012). Improving fleet utilization for carriers by interval scheduling. *European Journal of Operational Research*, 218(1):261 – 269.

Martinez, J. J., Brown, T. M., Conway, P., Elphick, M., Falarski, M., Hawkins, E., Rosing, W., and Shobbrook, J. (2010). Lcogt sites and site operations. volume 7737, page 77370O. SPIE.

Nather, R. E., Winget, D. E., Clemens, J. C., Hansen, C. J., and Hine, B. P. (1990). The whole earth telescope - A new astronomical instrument. *Astrophysical Journal*, 361:309–317.

Saunders, E. S., Naylor, T., and Allan, A. (2006). Optimal placement of a limited number of observations for period searches. *Astronomy and Astrophysics*, 455:757–763.

Snodgrass, C., Tsapras, Y., Street, R., Bramich, D., Horne, K., Dominik, M., and Allan, A. (2008). The WEBplop observation prioritisation system. In *Manchester Microlensing Conference*.

Tsapras, Y., Street, R., Horne, K., Snodgrass, C., Dominik, M., Allan, A., Steele, I., Bramich, D., Saunders, E., Rattenbury, N., Mottram, C., Fraser, S., Clay, N., Burgdorf, M., Bode, M., Lister, T., Hawkins, E., Beaulieu, J., Fouqu, P., Albrow, M., Menzies, J., Cassan, A., and Dominis-Prester, D. (2009). RoboNet-II: Follow-up observations of microlensing events with a robotic network of telescopes. *Astronomische Nachrichten*, 330(1):4–11.

Vestrand, W. T., Borozdin, K. N., Brumby, S. P., Casperson, D. E., Fenimore, E. E., Galassi, M. C., McGowan, K., Perkins, S. J., Priedhorsky, W. C., Starr, D., White, R., Wozniak, P., and Wren, J. A. (2002). The RAPTOR experiment: a system for monitoring the optical sky in real time. In Kibrick, R. I., editor, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 4845 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 126–136.