Multi-Telescope Observing: the LCOGT Network Scheduler

Eric S. Saunders
Las Cumbres Observatory Global Telescope

Sotiria Lampoudi
Liquid Robotics

1 Introduction

Las Cumbres Observatory Global Telescope (LCOGT) is developing a worldwide network of fully robotic optical telescopes dedicated to time-domain astronomy. Observatory automation, longitudinal spacing of the sites, and a centralised network scheduler enable a range of observing modes impossible with traditional manual observing from a single location. These include continuous coverage of targets across sites, simultaneous observing with multiple resources, and cadenced time-series monitoring without diurnal gaps. The network also provides resource redundancy, with the potential for observations to be rescheduled in response to changing weather conditions. The scheduling model supports a wide variety of observing programs, which typically have very different constraints, goals, contingencies and timescales.

Heterogeneous requests in a networked environment present specific, unusual challenges for telescope scheduling that do not arise with single-resource schedulers. Here, we present a short overview of the LCOGT network scheduler. We describe our design goals, outline the scheduler's modular architecture, and highlight its current and planned capabilities.

2 Design goals

Las Cumbres Observatory Global Telescope (LCOGT) is a non-profit foundation dedicated to ground-based optical time-domain astronomy. Two 2m class and nine 1m class robotic telescopes, deployed across five sites and controlled and coordinated from a central headquarters in Santa Barbara, California, provide an automated observing network for a wide range of time-domain science (Fig. 1) [1].

Each telescope is fully robotic. All hardware is automated, from the weather sensors and dome to the telescope mount, instrumentation, and controlling computers.

Figure 1: The 2x2m and 9x1m telescopes making up the deployed LCOGT network, as of January 2014.

This level of automation is the enabler for the key feature of the network - its global scheduling. The control metaphor for the LCOGT network is that of a single global instrument, capable of observing an unusually large area of sky, and, for southern targets, almost always in darkness, so unfettered by the diurnal cycle. The network is homogeneous; within each telescope class, the telescopes are exact clones, with identical build, instrumentation and filters. Both spectra and CCD imaging may be performed from any site.

Most sites have more than one telescope present. This is desirable from both a cost and resource perspective, making maintenance easier and allowing the telescopes joint access to a shared site spectrograph. Multiple telescopes also provide important redundancy, increasing throughput (several observations can be concurrently made from the site), while mitigating against technical failures of individual resources.

The network is controlled and operated from a central headquarters, through a distributed "hub and spoke" software architecture. The software system is made up of a number of layers of increasing abstraction, from the low-level code that drives the hardware, through the sequencer for each telescope, and up to the agents that mediate access to each site [2]. Within the hub, the central databases that track the status of observing requests are interrogated and updated by the network scheduler in response to changing conditions and new observing requests, which are created by most users through the LCOGT website, and in some cases are submitted programmatically.

The modular architecture allows components to be tested or overridden in isolation, and also creates tolerance to network outages through the use of schedule caching at each site.

The purpose of the network is to support multiple distinct types of time domain science, across a large dynamic range of characteristic timescales. Many of these programs have competing requirements. Long-term continuous periodic monitoring, for example, is at odds with short timescale but highly urgent and potentially disruptive requests. Such requests may show up with very little warning. For example, transient survey programs generate previously unknown targets which require prompt follow-up. Thus an important design goal was that the network be highly responsive to new input, prioritise effectively between the many types of science, and allow requests to be made throughout an operational semester.

The other major goal was to fully utilise the potential arising from having a robotic network. Observations should be dynamically reschedulable in response to changing conditions at the sites. Schedule optimisation should take place across all telescopes together. Unique observing modes such as synchronous observing, hand-off between telescopes (for long unbroken time series), and cadences without diurnal aliasing should be possible. Finally, the abstraction of a single instrument should be preserved, so that observations can be described in a high-level way, independent of any particular observing site.

3 System overview

Requests are the principal unit of scheduling. They are abstract, generalised observing descriptions, which provide a way to ask for data from the network as a whole. Each request has one or more *observing windows*, which are temporal constraints indicating an outer bound within which an observation can occur. Note that an observing window may be much larger than the total duration of the observation. A window of one week, for example, means that the request may be performed at any time during that week, subject to target visibility. In general, larger windows provide more opportunities for a request to be successfully scheduled.

Requests also hold parameters such as target, autoguider, exposure and filter requirements, which are necessary to actually perform the observation. Optionally, constraints such as maximum acceptable airmass or lunar distance may also be provided.

If a request is selected for the final schedule, a *block* is created. A block is a description of an observation that has been bound to a specific telescope and time.

Figure 2 shows the flow of requests and blocks through the network. Requests enter the system through the web portal, or programmatically through the web service API. They are validated and stored in the request database. The network scheduler

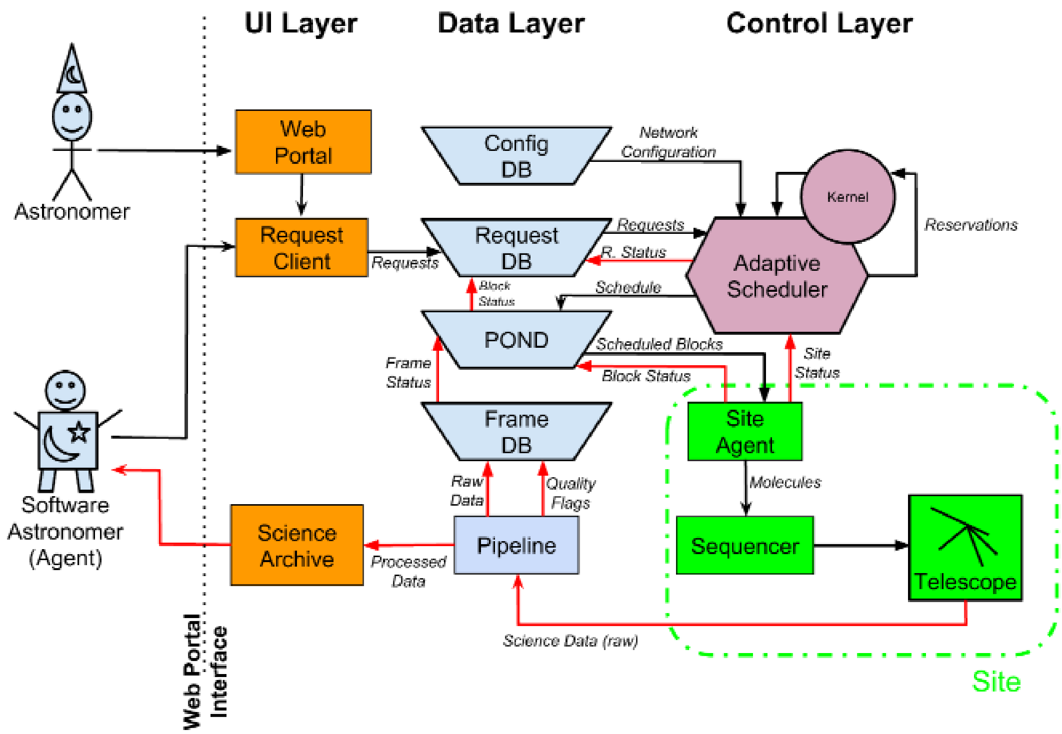

Figure 2: High-level architecture of the LCOGT observing software stack.

monitors the request database continuously. A scheduling run is triggered in response to changes in the pool of requests. A triggering change could be a new request arriving, an existing request being cancelled, or a scheduled block failing to be observed. Re-scheduling is also triggered by any change in the set of available telescope resources. If a site becomes temporarily unusable as a result of cloud cover, for example, then the schedule is recalculated, and observations are reassigned among the remaining available sites.

Re-scheduling never pre-empts a block that is already executing at a telescope at scheduling time; interrupting such blocks is inefficient and causes significant "thrashing". The only exception to this rule are requests flagged as time critical, called Target of Opportunity (ToO). These requests require immediate placement at the first available telescope. As long as ToO requests comprise only a small fraction of the total request pool, such interruptions are tolerable.

For each site, the windows of each request are reduced by the hard constraints of day and night, target visibility, and any other *a priori* constraints (e.g. airmass) specified. Figure 3 depicts how this step reduces the range of possible times available to a request.

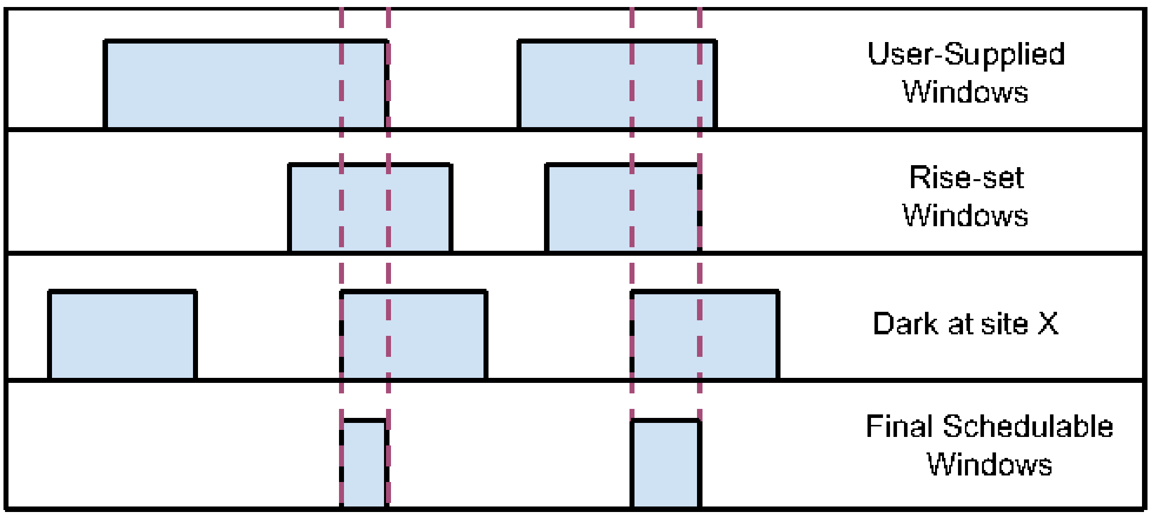

Figure 3: Intersecting request windows by hard constraints to obtain final schedulable windows.

With the windows now fully determined, the requests are mapped to an abstract formalism [3] that can be passed to the scheduling kernel. The scheduling kernel solves the well-defined mathematical problem laid out in [3] using a discretised integer linear program. The implementation of the kernel is beyond the scope of this paper; see [4] for a detailed discussion.

The computed schedule (optimised with respect to the implemented priority function) selects the requests which should be observed, and indicates where and when they should take place. Corresponding blocks are constructed, and stored in the Proposal and Observation Network Database (POND). This database is polled by the site agent at each site to determine what to observe.

Each site agent pulls several days of scheduling from the POND and stores a copy locally. This schedule caching allows the site to continue to observe for some time even in the event of a network outage between the site and the scheduler.

4 Version 1.0 capabilities

The LCOGT network begins official operations in April 2014. The central goal of the initial system is reliable operation and monitoring of the network with a small core feature set. From a scheduling perspective, it will include the following key capabilities:

- Globally optimised observation placement
- Simultaneous and cross-site observing
- Automatic weather rescheduling
- Automatic rescheduling of unsuccessful observations
- Target of opportunity observations, placed within 15 minutes of submission
- Both human and programmatic request submission interfaces
- Cadence-driven, multi-site sequences
- Support for solar-system objects
- Low-dispersion spectrographs at both 2m sites [1]

5 Future work

Many features are slated for implementation beyond version 1.0. Adding seeing and transparency constraints will allow just-in-time site selection based on current site conditions. Intra-project priorities will allow users to provide information on the relative importance of requests within their proposals. Evaluating the performance of different priority functions on simulated and historical data sets will be an ongoing process to tune the scheduling output. Automatic decomposition of large contiguous windows among sites will conveniently allow continuous observing for longer than a night. "Smart" cadences that autonomously adjust behaviour in response to missing observations are a largely unexplored area of research which promise to be powerful tools for specifying dynamic observing strategies [5]. Additionally, parts of the network may be used to monitor space junk and track satellites. Such objects present their own unique scheduling challenges.

References

[1] Brown, T. M., Baliber, N., Bianco F. B., et al.: 2013, PASP 125, 931
[2] Hawkins, E., Baliber, N., Bowman, M., et al.:2010, SPIE 7737, 77370P
[3] Lampoudi, S., Saunders, E. S.: 2013, ICORES 313
[4] Lampoudi, S., Saunders, E. S.: 2014, in prep.
[5] Saunders, E. S.: 2007, University of Exeter, "Optimal Observing of Astronomical Time Series Using Autonomous Agents", PhD thesis