Control Software and Scheduling of the Liverpool Robotic Telescope (A Revised Version Of SPIE paper 3112-28, in press)

I.A.Steele and D.Carter

Astrophysics Group, Liverpool John Moores University, Liverpool L3 3AF, U.K.

ABSTRACT

The Liverpool Telescope will be a fully robotic 2m telescope situated in La Palma and managed from Liverpool John Moores University. The data flows between the robotic components of the telescope are discussed and the functionality of those components described. We pay particular attention to the design of the telescope scheduler. Using simulations we show that a relatively simple dispatch scheduler may be tuned to produce a schedule that has acceptable trade-offs between conflicting scheduling criteria.

Keywords: telescopes, robotic telescopes, telescope control systems, telescope scheduling

1. INTRODUCTION

The Liverpool Telescope will be the first telescope produced by the New Generation Astronomical Telescopes project,[1] which is a joint initiative of Liverpool John Moores University, the Royal Greenwich Observatory and National Museums and Galleries on Merseyside. As such it has three functions:

- To provide a unique research facility for the JMU Astrophysics Group and the wider astronomical community.
- To provide images and other data for public understanding of science activities.
- To act as a technology prover and demonstrator for the NGAT company, Telescope Technologies Ltd, who are producing both robotic and traditionally operated telescopes in the 1 to 5m range.

The telescope will be situated at the ORM international observatory site on island of La Palma. The optical and tracking specification of the telescope is summarized in Table 1. The basic telescope design is ‘alt-az’, with hydrostatic bearings for each axis. Control of the telescope motion will use tape encoders and DC servo motors, the velocity loop being closed via. a PMAC based mount control system. The telescope will have an A&G box with 4 side ports and one straight through port for mounting of instruments. Instrument selection will be by a controlled, retractable, rotating mirror. A fixed auto-guider will also be mounted in the A&G box at a distance of 15 arc-minutes from the field centre. The first instrument to be mounted on the telescope will be a CCD Camera, probably using a 2048x2048 chip. This will have a 12 position filter wheel and electronic shutter and will probably be mounted at the straight through position. Subsequent instruments are anticipated to be an Infra-red Camera, a fibre bundle fed spectrograph and a tip-tilt corrector.

The normal mode of operation of the Liverpool Telescope will be fully robotic. This means that the telescope will not be supervised either locally or remotely during its routine operation. Instead it will carry out a programme of observations according to a schedule and return the data to the observer in the form they require. Such a mode of telescope operation allows science programmes to be done, such as fast response to target of opportunity observations and long term monitoring, that are virtually impossible with a conventionally scheduled telescope.[2] In addition proper matching of programmes to observing conditions can be carried out so that, for example, programmes requiring the ‘best’ seeing can be scheduled with a genuine chance of being done. Finally there are considerable cost and convenience benefits in going to a fully robotic mode of operation.[3]

In order for a telescope to operate robotically it is necessary to implement a software system that attempts to act as a human observer would in any situation (only hopefully more efficiently!). This includes such procedures as

email: ias@staru1.livjm.ac.uk, dxc@staru1.livjm.ac.uk

telescope run-up and down and focusing as well as actual observations. It also must include the ability to respond to faults and errors that occur and handle them in a manner that ensures that the telescope and its associated systems are not damaged. All This is the task of what we have called the robotic system (alternatively the Observatory Control System). This will provide functionality in the following areas:

- Scheduling observations
- Start-up and close-down procedures
- Weather monitoring (independent weather monitoring will also be provided by the telescope interlock system)
- Sequencing within a given observation
- Top level error handling and reporting (a low level interlock system will act as the 'final line of defense' however)

In addition a data reduction pipeline will be implemented. The degree of reduction carried out at the telescope will be specified by the requirements of each individual programme. The highest level of data reduction will provide fully reduced data product (i.e. magnitudes, spectra) without any user intervention. It is envisaged that the data reduction system will run as far as possible independently from the rest of the telescope systems so as not to impose any timing constraints on telescope operations. This implies some form of data reduction queue system along the lines of that developed for CGS4 on UKIRT.[4]

It should be noted that direct control of the telescope will still be possible, either from the site of the telescope itself, or via. Internet connection for authorized users. In addition in certain circumstances (for example public demonstrations in the Liverpool Planetarium and error recovery) it will be necessary that data from the telescope is returned to a remote user as near to 'real-time' as possible.

The project is currently poised on the transition between the design and implementation phases. The primary mirror has been ordered, and we are now in the stages of detailed mechanical and control design. This paper presents a brief overview of the robotic components that are being developed, based around a dataflow diagram. It then considers in more detail how we will implement scheduling of the telescope.

Table 1. Optical and Tracking Specification of the Liverpool Telescope

Item	Specification
Primary Mirror Diameter	2.0m
Primary Mirror f-ratio	f/3
Final f-ratio	f/10
Image Quality (12 arcmin field)	$<$ 0.4 arcsec (80% encircled)
Image Quality (42 arcmin field)	$<$ 0.6 arcsec (80% encircled)
Pointing Accuracy	$<$ 2 arcsec rms
Open Loop Tracking	$<$ 0.2 arcsec upto 1 minute, $<$ 0.4 arcsec in 10 minutes
Closed Loop Tracking	$<$ 0.2 arcsec in 1 hour
Rotator Tracking	$<$ 0.2 arcsec in 1 hour

2. DATA FLOW MODEL

We believe the essential aspects of a robotic telescope are best modeled in terms of data flows rather than in terms of control and algorithms. From this point of view a telescope is a complex database system in which an observer makes a request for data they need to answer a certain question or questions. The telescope and its associated instruments will return that data to the astronomer (all be it often many months later), who then has the responsibility of deciding how and if it answers the question he or she originally posed.

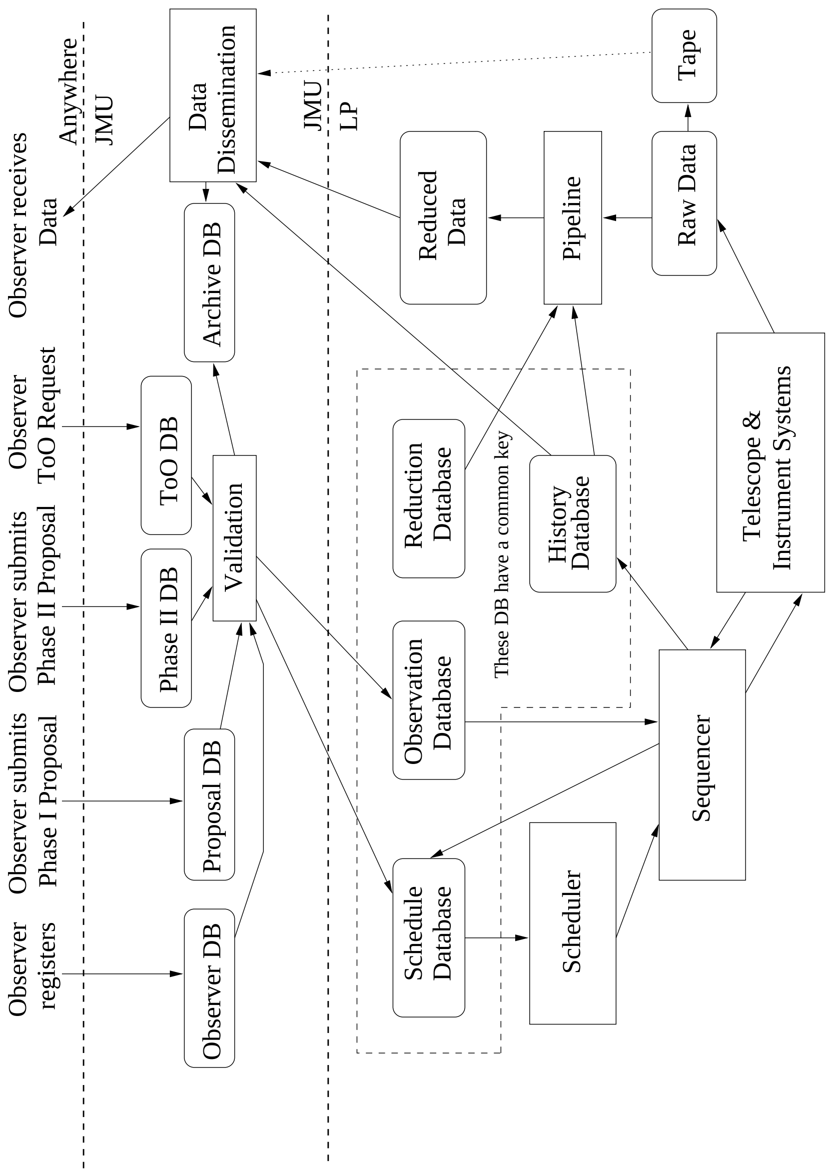

Figure 1. Liverpool Telescope High Level Data flow Model

Figure 1 therefore shows our derived high level data flow model for the Liverpool Telescope robotic systems. The diagram abstracts all of the telescope, A&G, autoguider and instrument systems and their associated subsystems into one unit. That unit receives data from the sequencing task as to what to do, returns data to the sequencing task about successes and failures, and generates raw data from the instruments for subsequent reduction. The communication protocol between the sequencer and the systems below it will be DRAMA.[5] In the diagram data storage areas have rounded edges and data processing and control tasks square edges.

In the following sections we describe the functionality associated with each data storage and processing area.

2.1. Off Site Systems

All telescope users must register to use the telescope and have that registration approved by the project science team according to a set of criteria to be defined. The registration process will be carried out over the Internet using a WWW based form. Each registered user will be issued with a username and password, which will allow them full access to the telescope web site, and therefore the ability to submit both normal (Phase I/II proposals) and Target of Opportunity (ToO) observing requests.

The details of registered observers will be stored in the Observer Database. This will include information such as name, institute, email, postal address. The ability for an observer to modify this information will be provided.

We anticipate two methods of generating observing requests:

2.1.1. 'Normal' Proposals

This method will be used for 'non-ToO' programmes, such as long and short term monitoring, surveys etc. A two phase proposal method is most appropriate here, as it minimizes the effort required by the proposer until the programme has been approved. In Phase I the proposer will provide a scientific and technical case and an estimate of the total observing time and conditions required to make the observations. This phase I proposal will then be passed to the appropriate telescope TAG based on the nationality and/or institution of the PI. Assuming a TAG had made a time allocation to the programme, then the observer will have to carry out Phase II of the proposal process. In this phase complete and final details of the approved observations (e.g. exact coordinates, instrument configuration, exposure times etc.) will be entered for each target.

It is anticipated that the phase I proposal will written using a specially formated text file, probably using the phase I tool currently under development by Gemini. The phase I proposal would be stored in the proposal database, including the full text of the proposal, the scientific and PUST abstracts, the time requested and the time granted.

For phase II, when considerable verification of the input data is necessary we propose to distribute a graphical observation input tool, written as a Java application that interacts via. the JDBC API with the phase II database at JMU. The data would then be re-verified at JMU before being passed to the schedule, observation and reduction databases at the Telescope. This is again similar to the approach taken by Gemini, although for this phase we are of course writing our own observation input tool.

2.1.2. Target of Opportunity Proposals

We have a target timescale of 1 day for implementing Target of Opportunity (ToO) observations. In order to minimize response time, a one phase proposal method is most appropriate here. The observer will have to provide both the scientific case and the detailed target information at the time of the proposal. The proposal will be considered by the Project Science team at JMU, and, if accepted it will after verification be passed straight to the telescope database systems. For this observing mode we will provide a Java servlet web based interface to the ToO database.

2.2. Telescope Database Systems

Data from the proposal systems at JMU will be verified and passed to the schedule, observation and reduction databases at the Liverpool Telescope. These databases have been split so as to reduce the likelihood of conflicts when more than one process wants to access information about a given observation. They will have the information stored in them as follows:

2.2.1. Schedule Database

This database will contain all of the information needed to schedule an observation (which may consist of one or more exposures at one or more pointing positions), e.g. coordinates, required conditions, estimated length, grades etc. It will also contain a status flag which indicates whether the observation is, for example, OK_FOR_SCHEDULING, ON_HOLD, DONE_SUCCESSFULY or DONE_UNSUCCESFULLY.

Most of the data input for this database will come from the proposal databases at JMU. Only one data item will be fed from the sequencer into this database, namely the observation status. This reduces the risk of corruption of this database through simultaneous reading/writing or (more preferably!) performance degradation through locking conflicts.

We defer our discussion of the Scheduler and its algorithm until sections 3 and 4, where it is described in some detail, along with the results of simulations of our chosen approach.

2.2.2. Observation Database

This database accepts input from the proposal databases at JMU and outputs details of an observation on request from the sequencer. These details are those required to actually make the observation, such as coordinates, instrument configurations and exposure times. There is therefore some overlap between the data in the schedule and observation database. This implies either that great care must be taken in ensuring consistency between the two databases or that a relational database system is used to implement both databases as different 'views' of the same master database. The second option is our chosen one as it is more flexible and less prone to synchronization errors.

2.2.3. Reduction Database

This database accepts input from the proposal databases at JMU and outputs details of how the data should be reduced by the pipeline reduction process at the telescopes. It is separated from the other databases to prevent locking conflicts as it is likely that as one observation is being scheduled and executed, previous observations will be being reduced, and it would be unacceptable for observations to be delayed by waiting on reduction (which can be carried out at any time).

2.2.4. History Database

This database takes data from the sequencer about the progress of a particular observation and stores it. That data is then available to the data reduction pipeline so it knows the history of data it is going to reduce. The History database will also store all error messages raised, associating them both with the observation that was being undertaken at the time, and the particular action that was being carried out when the error occurred.

2.3. The Sequencer

The Sequencer acts as the essential link between the database systems and the telescope and instrument systems. It has the following areas of responsibility:

- Start / End of observing session procedures
- Requesting observations from the Scheduler
- Sequencing the actions within an observation
- Initiating Shutdowns in bad weather
- Responding appropriately to errors from telescope and instrument systems

From the above it can be seen that the sequencer is a complex, concurrent system. Detailed analysis and design of this system is currently in progress.

2.4. Pipeline Data Reduction

Pipeline data reduction will be carried out at the telescope. For imaging observations we intend to allow the observer the choice of any combination of the following information:

1. Raw Data
2. Debiased and flat fielded frames.
3. Uncalibrated parametrized object lists for frames.
4. The relative magnitude of a single object in the frame to secondary standards within that frame
5. Photometrically and Astrometrically calibrated parameterized object lists.
6. The apparent magnitude of a single pre-chosen object in the frame with in a standard photometric system.

We note that the combination of flexible scheduling and standardized observing modes resulting from robotic operation means that the implementation of the procedures necessary to produce such highly reduced data are more feasible than with a conventional telescope. For example the complete set of photometric standards that will automatically be obtained regularly throughout each night will allow proper calibration and quality control of the reduced data.

2.5. Data Dissemination and the Archive Database

Both pipeline reduced and raw data will eventually be returned to JMU, and forwarded from there to the observer. That data will also be copied to the archive database, along with the history of that observation from both the proposal databases and the history database at the telescope, and a time after which the data will become public. A web based interface to this archive database will be provided along similar lines to that proposed for Gemini. The archive database will also contain the data necessary for monitoring telescope scientific productivity and performance and the distribution of observing time between various project partners. A graphical interface to the database will therefore be an essential management tool for the telescope.

3. THE SCHEDULER

3.1. What is a good schedule?

For the purposes of a telescope which has observing time shared between a number of partners we have identified three criteria that must be considered when constructing a schedule:

1. **The schedule must be fair.** All of the partners must receive allocations of time on the telescope that reflect both an agreed total percentage for that partner and an equitable distribution of both good and bad observing conditions (seeing, sky brightness etc.) between all of the partners.
2. **The schedule must be efficient.** Observations should be well matched to the observing conditions and the wishes of the Telescope Allocation Group (TAG). In addition the number of observations done in a given period of time should be as high as possible.
3. **The schedule must be sensible.** Observations should not be attempted that are impossible in the current conditions or because of visibility or other constraints.

It is immediately apparent that both within and between the first two requirements above are many conflicting factors. For example if we place priority on matching conditions, this will tend to detract from doing observations based on a scientific priority made by the TAG. We therefore take the view that a good schedule is therefore one that minimizes all of these conflicting factors.

3.2. Optimum vs. dispatch scheduling

Approaches to scheduling can conveniently be divided into two. The first approach attempts to construct an optimum schedule for a substantial period of time (say 1 night of observations). This has the advantage that factors such as minimization of slew times over the whole night can be considered, and such schedulers have been developed.[6] The disadvantage of such a system is that if the observing conditions are changing on a timescale similar or less than the period of time a schedule is constructed for, then the schedule will 'break' and have to be reconstructed. This is generally a computationally intensive task.

The second approach is to use a simple 'dispatch' scheduler. This merely looks to see what the best observation to do at a given time based on the current telescope state and conditions without any attempt to look ahead. This has obvious advantages for a site such as La Palma where conditions throughout a night can change both rapidly and often. Of course in long periods of stable conditions it is likely to be less efficient than a optimum scheduler.

We have chosen to implement a dispatch scheduler for the Liverpool Telescope, although the modular nature of the software system means we could change this at a later date if we felt the necessity. The implementation of such a scheduler in the case of the Liverpool Telescope is slightly more complex than is usual as we do need to look ahead one observation to ensure we will not overrun into a fixed time constraint. Such fixed time constraints can be both observational, for example simultaneous observations with a satellite or of a transit or eclipse, or more mundane, such as the necessity to obtain a sky flat during morning twilight.

In order to decide which observation (n) to carry out we use three rating functions. The first function $f(n)$ is a measure of fairness, the second $e(n)$ is a measure of efficiency and the third, a boolean veto function $v(n)$, is a measure of sensibleness. For an observation consisting of a single pointing/exposure these are combined as:

$$R(n) = f(n).v(n).e(n). \tag{1}$$

For a total of X multiple pointing/exposures (x) within a single observation this becomes

$$R(n) = f_1(n). \prod_{x=1}^{x=X} v_x(n). \frac{\sum_{x=1}^{x=X} e_x(n)}{X}. \quad (2)$$

The observation that is carried out is then simply that with the maximum value of $R(n)$.

3.3. The Fairness Function

The purpose of the fairness function is to evaluate how fair doing a particular observation is based on the requirement of equitable shares of observing conditions between all partners. To do this we first consider what the various conditions are that are to be 'shared-out'.

We consider two variables when determining the quality of a given night, the lunar phase and the atmospheric quality. For each we list a hierarchy. A given observation can be carried out in the level corresponding to those conditions and any higher level. In other words we assume you can, for example, do a programme requiring 'average' seeing in 'best' seeing conditions without *scientific* penalty.

3.3.1. Moon Phase

Since the telescope will be scheduled on a minute-by-minute rather than night-by-night basis, we can split nights into dark and bright sections, rather than using the concept of grey time. Averaged over a year the moon will be above the horizon 50 per-cent of the time at night.

Table 2. Moon Phase Conditions

Name	Definition	Estimated %
Dark	Moon below the Horizon	50
Bright	Moon above the Horizon	50

3.3.2. Atmospheric Conditions

Since seeing and the photometric quality of a night are not independent variables we choose not to separate the two, but instead simply assign the conditions to one of four categories outlined in Table 3 below. This is similar to the approach taken by Gemini.[7] The photometric and clear percentages used to calculate the occurrence of these each set of conditions are based on the ING site survey.[8]

Table 3. Atmospheric Conditions

Name	Description	Estimated %
Excellent	Photometric with upper quartile seeing	15
Good	Photometric with 2nd & 3rd quartile seeing	30
Poor	Photometric w. lower quartile seeing OR Non-photometric w. any seeing	25
Unusable	$>$ 50% cloud or humidity/wind etc. limits exceeded	30

3.3.3. Resulting Constraints on the TAG

In order that fairness is achieved in scheduling, each of the partner TAG's will therefore have to provide observing proposals for the telescope scheduler in the proportions given below. They will even allocate the 'Unusable' time (probably to projects of low priority), as the exact amount of unusable time in any one year will of course vary.

These proportions may have to be revised once the photometric/non-photometric split is better determined during actual telescope operations.

Table 4. Percentages for TAGs to allocate time in

Conditions	Excellent	Good	Poor	Unusable	TOTAL
Dark Moon	7.5%	15.0%	12.5%	15.0%	50%
Bright Moon	7.5%	15.0%	12.5%	15.0%	50%
TOTAL	15%	30%	25%	30%	=100%

3.3.4. Form of the Fairness Function

Since the TAGs will assign their time in the expected ratios of the conditions, then over a sufficient period of time the fairness function should simply have the form

$$f(n) = 1. \tag{3}$$

However in practice the length of time for the fair distribution to converge will probably be limited to one year. Multi-night modeling with the test scheduler (Section 4) has shown this is insufficient time to rely simply on the above fairness function.

Instead we therefore use a fairness function of the form

$$f(n) = \alpha_{p,c}(n), \tag{4}$$

where $\alpha_{p,c}$ is a near-unity constant that has different values for different partners p and requested conditions c. There are therefore eight α values per partner. The values of $\alpha_{p,c}$ will be adjusted monthly according to the formula

$$\alpha_{p,c} = \frac{\text{target\_fraction\_for\_}p, c}{\text{achieved\_fraction\_over\_last\_6\_months}}. \tag{5}$$

3.4. The Efficiency Function

The purpose of the efficiency function is to decide which observation to carry out at any given moment in time considering the factors of scientific priority and best use of observing conditions. This will be done using an efficiency function of the form

$$e(n) = \beta_1 E_1(n) + \beta_2 E_2(n) + \cdots + \beta_Q E_Q(n), \tag{6}$$

where the functions $E_q(n)$ represent the match between the current conditions or requested priority and the actual conditions. Three example E functions are those related to the object height in the sky, the observing conditions and the scientific priority, which we describe in detail below. Other functions we will implement consider such aspects as slew times between the current position and the object, instrumentation configuration and change times, the need for regularity in monitoring observations and programme completion.

3.4.1. Object Height

It is obviously advantageous to observe objects at as low an airmass (i.e. highest in the sky) as possible. This would imply an E function of the form

$$E_1(n) = \frac{\text{current\_altitude\_of\_n}}{90^\circ}. \tag{7}$$

However this will tend to favour objects at around the declination of the observatory site against any others. Therefore a better criteria is to attempt to observe objects at the highest points in their course across the sky. This implies an E function

$$E_2(n) = \frac{\text{current\_altitude\_of\_n}}{\text{transit\_altitude\_of\_n}}, \tag{8}$$

is better suited to this task.

3.4.2. Observing Conditions

As already discussed in Section 3.3.2, we have split observing conditions into 4 levels, which for the purposes of the evaluation function we number from 1 (Unusable) to 4 (Excellent). The veto function V_4 (see below) prevents objects being observed in poorer conditions than is required. However to allow scientific priority to also be a factor in efficiency we do not simply require an exact match of required and observed conditions, but instead use the following function:

$$E_3(n) = \frac{1}{|(\text{required\_condition\_number} - \text{actual\_condition\_number})| + 1}. \tag{9}$$

3.4.3. Scientific Priority

If TAGs wish, scientific priorities for observations may be supplied for use in scheduling. The priorities $p(n)$ from a given TAG will be normalized to ensure that

$$\frac{\sum p(n)}{N} = 0.5, \tag{10}$$

where N is the total number of observations from a given TAG. If the TAG chooses not to assign priorities, then we assign $p(n) = 0.5$ for all of that TAG's observation. In either case the E function is simply

$$E_4(n) = p(n). \tag{11}$$

3.5. The Veto Function

The purpose of the veto function is to prevent observations being carried out that are not possible at the current time due any of a number of boolean (yes=1,no=0) constraints:

$$v(n) = V_1(n) \times V_2(n) \times \cdots \times V_P(n). \tag{12}$$

In this function V_p represents one of a series of P boolean expressions. Some examples of these expressions are given in Table 5.

Table 5. Example Scheduler Veto Functions

V_1	=	1 if (altitude) > (telescope altitude limit)
	=	0 otherwise
V_2	=	0 if length of time to next fixed time observation
		< estimated(slew time + reconfigure time + exposure time)
	=	1 otherwise
V_3	=	0 otherwise
	=	1 if dark time needed and moon up
V_4	=	1 if (conditions) equal to or better than (required conditions)
	=	0 otherwise

4. TEST SCHEDULER

In order to evaluate the scheme proposed above we have written a simplified implementation of it in Perl. This implements the veto functions V_1 and V_4, the efficiency functions E_1, E_2, E_3 and E_4 and the fairness function for 3 partners with equal shares of time. The ability to alter the efficiency function coefficients β_{1-4} is also provided. The scheduler can operate over one or many nights on upto several thousand observations (limited only by processor speed). To evaluate the success of the schedules produced we use both subjective criteria (e.g. plots of Hour Angle versus time through the night), and objective criteria (e.g. fraction of the total number of observations of a given priority done). For the simulations we present here we simply consider scheduling one night out of a total sample

Table 6. Three examples of input parameters to the scheduler and some quantitative measures of the resulting schedules

Object Height	Scientific Priority	Upper-quartile Priority Done	Scatter from HA=0	See Figure
YES ($\beta_2 = 2.0$)	NO ($\beta_4 = 0.0$)	26%	2°	2
NO ($\beta_2 = 0.0$)	YES ($\beta_4 = 2.0$)	78%	57°	3
YES ($\beta_2 = 1.0$)	YES ($\beta_4 = 1.0$)	55%	24°	4

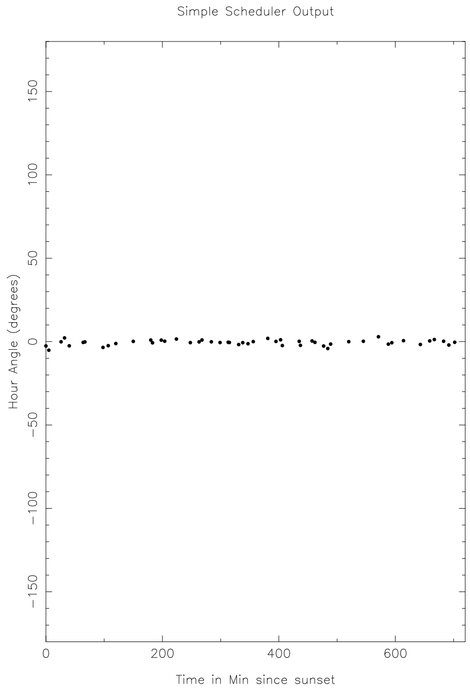

Figure 2. A schedule based solely on object height

of 200 potential observations. The parameters of the input observations to the scheduler are generated randomly according to the expected distributions outlined in Section 3.3.

As an example of the use of this program, consider the effects of varying the efficiency function coefficients. Table 6 lists the parameters that are varied to investigate the trade-off between fractional object height (β_2) and the scientific priority (β_4). Also listed in the table for the resulting schedules are the resulting rms scatter around Hour Angle=0 and the percentage of upper quartile priority observations carried out. In Figures 2, 3 and 4 we plot the Hour Angle versus time plots for the three schedules. It can be seen that by varying the β values we are able to tune the balance between any two scheduling parameters with relative ease.

5. CONCLUSION

The Liverpool Telescope has the potential to do unique scientific projects in the areas of fast response observations, long term monitoring and matching observations to conditions. To take advantage of this potential means we must implement fast and reliable human and software procedures for the management of proposals and telescope data, and scheduling and sequencing software that is both fair and efficient. This paper has discussed how we intend to implement those procedures and software in terms of the data flows present. In addition we have discussed our chosen approach to scheduling the telescope, and presented the results of simulations showing that a relatively simple dispatch scheduler can be tuned to produce schedules that balance conflicting efficiency criteria.

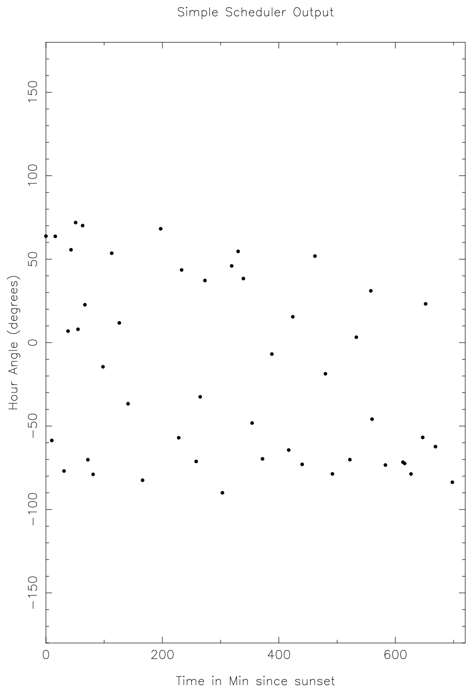

Figure 3. A schedule based solely on scientific priority

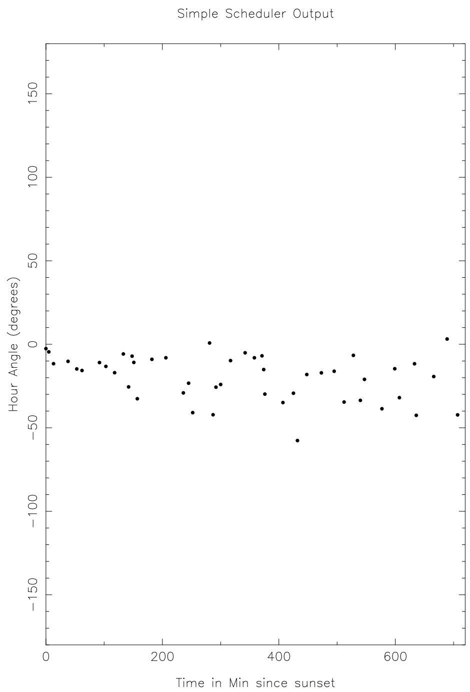

Figure 4. A schedule based on equal weights applied to object height and scientific priority

ACKNOWLEDGEMENTS

We wish to thank Keith Swanson, John Bresina and Will Edgington of the NASA Ames Telescope Scheduling Group, Mark Drummond of Tyecin Systems and Phil Puxley of the Gemini Science Team for much useful advice on telescope scheduling. The NGAT project is partially supported by the European Union under Objective 1 funding for the Merseyside Region. The Observatorio del Roque de los Muchachos (ORM) is operated by the Instituto de Astrofisica de Canarias, Tenerife.

REFERENCES

1. N. M. Parker, "Linking liverpool and la palma," *Spectrum* **10**, pp. 14–17, 1996.
2. M. F. Bode, "Editor's summary," in *Robotic Observatories*, M. F. Bode, ed., pp. 140–146, John Wiley and Sons, 1995.
3. D. L. Crawford, "Some gnat issues," in *Robotic Observatories*, M. F. Bode, ed., pp. 77–84, John Wiley and Sons, 1995.
4. G. Wright, M. C. Mountain, A. Briger, P. N. Daly, J. L. Griffin, and S. K. R. Howat, "Cgs4 experience: two years later," in *Infrared Detectors and Instrumentation*, A. M. Fowler, ed., *Proc. SPIE* **1946**, pp. 547–557, 1993.
5. J. A. Bailey, T. Farrell, and K. Shortridge, "Drama: an environment for distributed instrumentation software," in *Telescope Control Systems*, P. T. Wallace, ed., *Proc. SPIE* **2479**, pp. 62–68, 1995.
6. M. Drummond, J. Bresnia, W. Edginton, K. Swanson, G. Henry, and E. Drascher, "Flexible scehduling of automatic telescopes over the internet," in *Robotic Telescopes: Current Capabilities, Present Developments, and Future Prospects for Automated Astronomy*, G. W. Henry and J. A. Eaton, eds., *ASP Conference Series* **79**, 1995.
7. P. Puxley and T. Boroson, "Observing with a 21st century groundbased telescope - or how to do unique science with the gemini telescopes," in *Optical Telescopes of Today and Tomorrow*, *Proc. SPIE* **2871**, 1996.
8. B. McInnes *Quart. J. Royal. Ast. Soc.* **22**, p. 266, 1981.