Astronomical Data Analysis Software and Systems IX
ASP Conference Series, Vol. 216, 2000
N. Manset, C. Veillet, and D. Crabtree, eds.

Queue Scheduling at the JCMT: the Real-life Experience

R. P. J. Tilanus[1]

Joint Astronomy Centre, 660 N. A'ohoku Pl., Hilo, HI 96720

Abstract. During the past few years the subject of dynamic scheduling has become a prominent topic in discussions about future telescope operations. The term "dynamic scheduling", however, is somewhat misleading since it suggests that the central problem has to do with scheduling of specific observations at the telescope and that by solving that aspect you will successfully implement dynamic scheduling. The JCMT experience is that this is not the case: the majority of our development is geared towards providing an infra-structure which will allow for and support dynamic time allocation and service observing. Perhaps contrary to expectations, most of this development is concerned with what happens *during* and *after* observations take place, not before. A reliable but flexible infra-structure is the true challenge compared to which the details of the scheduling itself are relatively mild. Unlike a number of other observatories, however, the JCMT does not employ staff dedicated to service observing, but instead uses visiting observers to support its dynamic scheduling.

1. Dynamic Scheduling

For a number of reasons ground-based observatories are being pushed towards flexible queue observing or "dynamic scheduling". Dynamic scheduling aims to continuously schedule those projects on the telescope for which the observing conditions and telescope configuration are optimal[2]. In other words, the observations for all projects presently awarded time form a queue from which the ones best fitting the current constraints are selected to be observed. Each principal investigator (P.I.) is expected to provide the observatory with a precise "recipe" for their observations defining the required configuration as well as relevant constraints. Apart from weather conditions, common constraints are e.g. a minimum elevation for the sources or the phase of the moon, but a priority rating from the time-allocation process may also be incorporated.

While seemingly straightforward, this approach forms a dramatic departure from "classical" astronomical observing at most ground-based optical, infra-red and sub-mm observatories. Traditionally the P.I. or other members of the team

[1] Netherlands Organization for Scientific Research (NWO).

[2] The terms flexible scheduling is also in use. I have chosen to use dynamic scheduling throughout this paper.

associated with a project come to the telescope for "their" slot of awarded time and either manage to conduct the intended observations or switch to their fallback projects in case conditions are bad. Since on the whole fallback projects tend not to have the same scientific excellence, quality, and peer review as the primary projects, this situation in effect results in a net loss of observing time to the astronomical community. Dynamic scheduling aims to remedy this situation since often conditions which may not be suitable for one project may be perfectly adequate for others. An example is spectroscopy which does not need photometric sky conditions. A clear implication of dynamic scheduling is, however, that the P.I. will not in general be present at the telescope and that observations will need to be carried out in a serviced mode. While projects can be queued it is impractical to keep a queue of observers readily available on short notice, except in a situation where routine *remote observing* is feasible.

The foremost advantage of dynamic scheduling is that it results in a marked increase in the number of completed projects which require "specific" conditions, as opposed to many projects which are only partially completed. This in turn has an out-of-proportion effect on the number of publications since a not insignificant number of partially completed projects never get completed and/or published. Dynamic scheduling can result in a factor of 3 or larger increase in publications which require rare or exceptional observing conditions and a 10-20% increase in the overall publication rate. In this context, another very significant aspect of dynamic scheduling is that one can prioritize the observation queues to reflect the assessment resulting from the time-allocation or other peer-review process. While the individual P.I. may not appreciate this, it allows the community to bias the observations towards the ones which it believes will have the most scientific value. Finally, the ability to swap projects is very helpful in scheduling around observatory problems such as equipment failures with minimum impact to the overall program.

However, traditionally P.I.s have been very reluctant to hand over their projects for serviced observing because of concerns over the quality of the data and the ability to best achieve the scientific goals of the project. *While dynamic scheduling will maximize the quantity of useful astronomical data this does not automatically imply that it will also achieve the best quality.* Nevertheless, for dynamic scheduling to be a viable option it must be able to routinely deliver as good a quality as is achieved by the average experienced observer. If it fails to achieve this goal, dynamic scheduling will ultimately not be accepted by the astronomical community. For this reason, the majority of the development at the JCMT has concentrated on optimizing and securing data quality rather than scheduling the telescope. Key issues in this respect are that software, procedural, and feedback protocols *must* ensure the quality irrespective of whoever happens to be the observer at the telescope.

It is important to realize that when it adopts dynamic scheduling an observatory "de facto" also assumes a major responsibility for the quality of the data since the observer at the telescope can no longer be held accountable. Assuming that the observations are sufficiently well and unambiguously specified by the P.I., if the observer makes mistakes this clearly can not be charged against the time allocated to the project which happens to be the "victim". In fact, mistakes of this nature imply that either the protocol for serviced mode observing at the telescope is insufficient or in error, or that the observer had been insufficiently briefed or trained. Dynamic scheduling should result in more papers and a higher visibility, but the price paid by the observatory is a degree of liability and the investment needed to minimize this liability by implementing secure protocols.

In this respect, frequent and reliable feedback with the P.I. is an absolute requirement. While the observatory can make sure that the individual observations are of the highest quality, it can not assume responsibility for the project goals as a whole. To do that would require the same in-depth understanding of the problem as the P.I. and would also involve second-guessing the P.I. Often securing data of the required quality is sufficient to achieve the project goals. However, precisely the ability to deal with those situations when this is not the case, i.e. when intermediate results indicate that the observation process needs to be adjusted, is one of the most valued characteristics of traditional ground-based observing. It is essential to retain this aspect of ground-based observing as much as possible within the context of dynamic scheduling and the most obvious way to do this is to cycle observations back to the P.I. as quickly as is reasonable for feedback and possible changes to the observation plan[3].

Dynamic scheduling with P.I. feedback as an integral part makes sure that the control over the overall program remains with the P.I. even while the observatory assumes a degree of control over the individual observations. Ultimately such setup will allow for some creative new possibilities. For instance, the P.I. may request a *checkpoint* in the program at which time further observations will be stopped until the P.I. has had a chance the review the results and adapt the program. More extreme but potentially more rewarding would be to enforce a *mid-way checkpoint* in each project: for many projects, after half the allotted time has been expended it often is quite clear whether or not the science goals will be achieved or not e.g. in terms of detection rate, signal-to-noise etc. Rather than allowing unsuccessful projects to fill their time, it would be more efficient to re-allocate this time to more successful or new projects. In other words, the time allocation itself can become more dynamic and be made contingent upon meeting certain defined criteria at some point after the project has started. While such developments probably won't happen immediately, they indicate that further optimization and evolution of ground-based astronomical observing will be possible in the long run.

2. JCMT Developments

2.1. Boundary Conditions

In the next few sections I would like to describe some development work undertaken in support of dynamic queue scheduling at the JCMT. To do this it

---

[3] Obviously, one can not stop a program after every observation, hence this probably must be done on a "best-efforts-basis" e.g. to allow a response from the P.I. before the next appropriate observing period. Dynamic versus classical static scheduling thus implies a trade-off between optimized conditions and the degree of interactive feedback. "Remote eavesdropping", where the remote P.I. looks at the results and provides feedback but is not in actual control of the telescope, may be useful for certain projects.

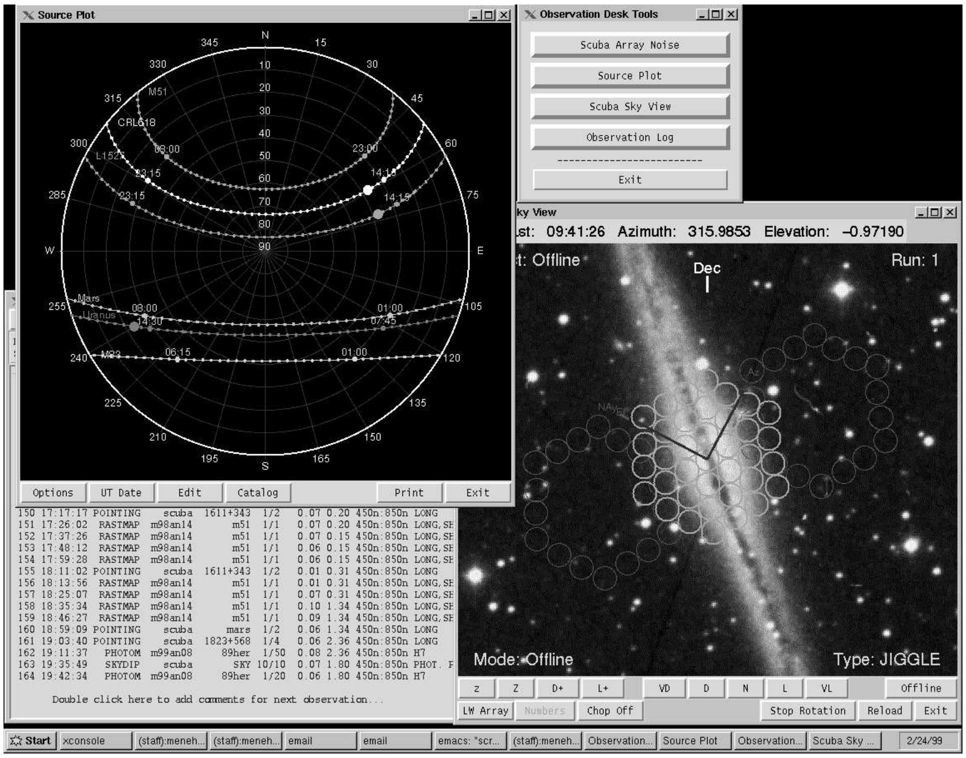

Figure 1. Sample screen of the Observation Desk utilities: Sourceplot (top-left), Obslog (bottom-left), Scusky (bottom-right).

is necessary to briefly sketch the background against which we undertook this work. Firstly, the JCMT is a sub-mm radio-telescope located on a 14,000 ft. volcano on the Big Island of Hawaii. The sub-mm nature of the observations make them very sensitive to weather conditions: 4 weather bands of varying opacity have currently been implemented to guide the choice of projects and conditions often gradually change during the night. The location of the JCMT and the fact that it is not fully automatic requires that at least two persons have to be physically present at the telescope during operation, a requirement that works against remote observing as the norm. Because of financial constraints the JCMT does not employ service astronomers and has to rely on external observers to carry out its observation program. Hence, a number of observers still have to come to the telescope, but they will not necessarily carry out their own program and may also stay longer than their allocation[4]. Observers may not be proficient with all the different observing modes and instruments.

Finally, data rates can be quite high and the instruments do not deliver a readily accessible product that can be used to evaluate the observations. Either substantial hands-on data reduction or heavy number crunching and processing is required to obtain preview images and spectra of sufficient quality upon which to base critical decisions.

The boundary conditions outlined above have guided the developments at the JCMT.

2.2. JCMT Observation Management Project

In order to provide the infra-structure for dynamic scheduling at the JCMT, the Observation Management Project (OMP) was defined to deliver solutions in three areas:

1. Utilities to enable possibly inexperienced observers to easily carry out all types of observations efficiently and to help visualize the telescope operation and configuration,
2. A sophisticated real-time pipeline reduction which delivers high-quality preview images without the observer needing to be familiar with the reduction process,
3. A highly automated administration and communication protocol to reduce overhead for the observer and staff and to ensure reliable feedback with the P.I. This item includes queue management and observation selection tools.

It should be realized that the scope of the project is not trivial at all. Since the observatory now has the requirement to monitor the full process from the initial submission of observation recipes or the observation plan, through the observations, all the way to the data hand-over, the project almost by default implements a "cradle-to-grave" approach. This is a trend common to all observatories which are planning to use dynamic scheduling. However, often the emphasis is on utilities to generate observation plans and the implementation of

[4]In practice at the JCMT visiting observers are allowed to at least partially carry out their own program if the conditions are sufficient because valuable experience gained this way leads to better projects. It also provides an incentive to come to the telescope as an observer.

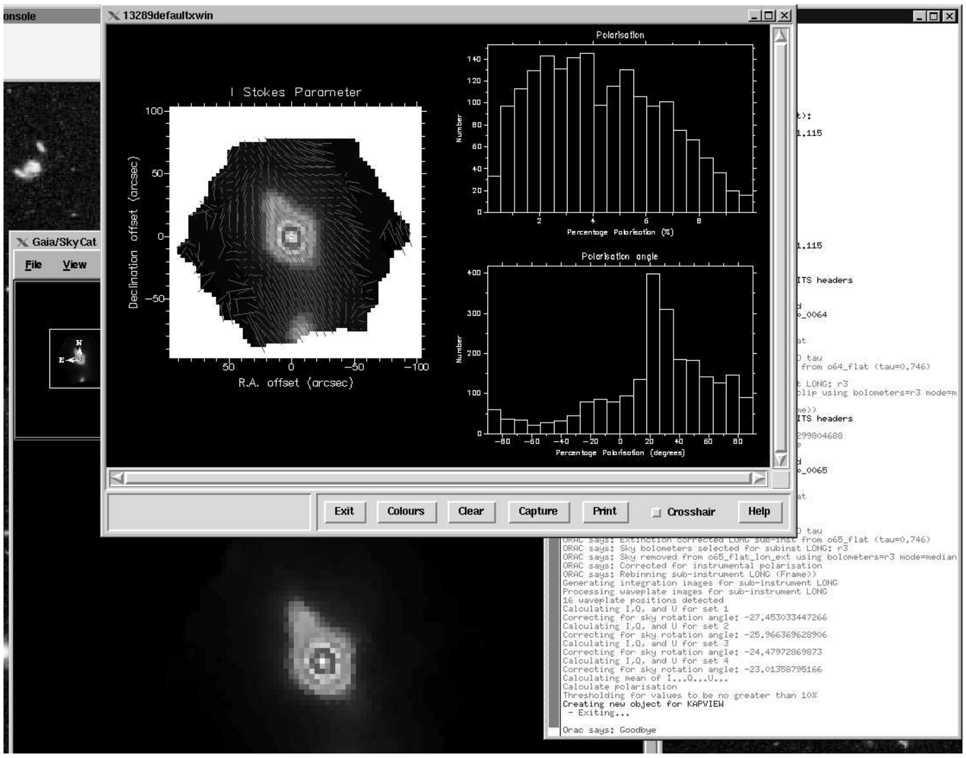

Figure 2. Sample pipeline reduction of polarization data showing the image, polarization vectors, and analysis histograms.

an observations database as the basis of an observing queue, at the expense of an infra-structure needed during and after the observations.

2.3. Observation Desk

The Observation Desk is the suite of utilities which assists observers at the telescope. Currently the following tools have been implemented (Figure 1).

- SOURCEPLOT: a utility to plot source parameters as a function of time, azimuth, and/or elevation. The sources move along their tracks as time goes on.
- SCUSKY: a utility to show the instrument science area and configuration on the sky (currently taken from the Digital Sky Survey). The tool shows the current pointing of the telescope and configuration of the instrument.
- SCUNOISE: a utility to graphically display instrument diagnostics.
- OBSLOG: a utility to create an observation log and which allows the observer to enter a narrative log entry and set quality flags for each data set. These comments and flags are saved in the JCMT data archive along with the data.

Other tools which have been suggested are utilities to more easily compare actual with expected results, to provide better monitoring of error conditions, and a sophisticated tool which continuously characterizes the real-time queue and its observations based on actual progress and expected conditions during observing. Note that all tools can also be run in an off-line mode allowing them to be used for preparation and diagnostics.

The general aim of the Observation Desk is to make observing "simple" and more reliable by providing expert support with specialized utilities.

2.4. Pipeline

The JCMT data reduction pipeline[5] (Jenness & Economou 1999) was implemented using the ORAC-DR backbone (Economou et al. 1999) and has proven to be extremely successful. The pipeline reduces and calibrates individual observations, co-adds compatible data sets (multiple exposures, mosaics, etc.) regardless of the sequence of data acquisition, and derives secondary results when relevant: e.g. histograms of the percentage and angle of polarization for polarization observations (Figure 2).

The pipeline not only frees the observer from the extremely time-consuming task of having to reduce the data while observing, but also obviates the need for the observer to be familiar with the data reduction for every observing mode and instrument. Since it is easily customized, the pipeline is increasingly being used for off-line reduction as well. Because of the data rate and amount of reduction needed we would have had to implement a pipeline irrespective of dynamic scheduling, but the latter is not possible without one.

---

[5] "Guide to using the ORAC-DR SCUBA Data reduction pipeline", `http://www.jach.hawaii.edu/JCMT/software/oracdr/scuba_userguide`.

2.5. Administration and Communications

While utilities such as the Observation Desk and the Pipeline are relatively straightforward, the procedural part, i.e. administrations and communications, is more elusive and contentious. Many telescopes are now adopting a "phased" approach to observation projects:

- *Phase-I*: submission of the observation proposal and time allocation.
- *Phase-II*: preparation and submission of the detailed observation plan and/or observation recipes for approved proposals.
- *Phase-III*: deals with the tracking of individual observations, the status of the overall project, as well as the feedback with the P.I. and data hand-over.
- *Phase-IV*: would involve the tracking and statistics of subsequent publications and use of the data via public archives.

Phase-I and -II have become fairly standard, Phase-III is generally recognized but has received much less attention and Phase-IV is merely a glimmer or thorn in a few eyes depending upon ones feelings towards it.

The Gemini Observation Tool (Wampler et al. 1997) and its incarnation for UKIRT (Bridger et al. 2000) are good examples of a sophisticated Phase-II observation preparation tool. The tool allows the P.I. to design a comprehensive observation plan which is saved in the observatory's database. Besides the astronomical sources, the plan also describes the calibration observations and allows for concepts such as sequencing (observation order) and blocking (observations which should be executed as a single block). Parts of a plan can be extracted and submitted for observing at the telescope. The JCMT currently works with observation recipes, but we plan to create our own version of the Gemini Observation Tool as part of the OMP.

For dynamic scheduling to work well one at least needs a database of the parameters of the queued observations. It does not necessarily need to be different from a bunch of files, but a relational or object-oriented DB is preferred. Observations can then be selected for scheduling at the telescope by applying well-defined filters on the DB. When conditions are predictable, one can try to employ an optimizing scheduler for a long-term plan covering the observation period e.g. the SPIKE scheduler used at STScI for HST observations. However, if conditions are very unpredictable a "just-in-time" browser to select the best few observations given the present constraints is probably sufficient and as efficient.

Having embarked upon dynamic scheduling we have found that currently Phase-III generates by far the largest drain on local resources since we do not yet have a well-designed system in place to handle the administrative overhead associated with observing. Observations need to be checked and must receive an initial sign-off. An account of the observing time needs to be filed and partially finished observations need to be flagged. Previews must generated and forwarded to the P.I. who sends the final sign-off or modifications of the observing plan. Quite a lot of this needs to be achieved on the day or few days following the observations to ensure that no duplicate observations happen, that mistakes are being corrected, and that updates are incorporated in a timely fashion.

To streamline this process we will be looking at action-request software as used by help desks and customer support services. While they may be overkill, these systems incorporate many of the features we need as they allow one to track a "case" through the system. In particular we want to keep track of all communications associated with a project. Not only the fact that emails have been sent, but also the confirmations that they have been received. We want to keep track of projects which may be on hold while waiting on feedback and be alerted if they remain on hold for too long. Time accountings, preview generation, and standard emails to the P.I. need to be generated automatically.

At the moment most of this work is being done by hand by local staff, the "queue managers", using email, text files, and web pages. While in principle this is a possible solution, it is very staff intensive and fairly unreliable because information has to be extracted from several different sources and, e.g., problems arise when people are absent. Ultimately this is not a sustainable mode of operation and needs to be replaced by an automated system as described above.

3. Conclusions

In the previous sections I have described dynamic scheduling and our "real-life" experience with it at the JCMT, and have argued that communication and feedback with the P.I. should be an integral part of the overall scheme. This requirement arises because ground-based astronomy is not about carrying out observations following a fixed set of parameters and which upon completion can be marked off as "done". Instead, ground-based observations are aimed at achieving a scientific goal in a highly interactive environment. In my opinion it would be a mistake to sacrifice this approach for the purpose of dynamic scheduling and to start treating ground-based observatories as zero-altitude satellites[6]. Instead we should aim higher and adopt the best features of three approaches:

1. Use dynamic scheduling to maximize the amount of high-quality astronomical data,
2. Request the rigorous planning of observations associated with space-based telescopes,
3. Adopt an intensive communications and feedback protocol to retain the interactive flavor of ground-based astronomy.

Based upon our experiences with dynamics scheduling we have found that the third item is a necessary ingredient. Lacking a reliable feedback mechanism the JCMT has created the position of "queue managers" who are expected to keep track of the progress of the overall queue and the status of the individual projects. Very quickly it became clear that these queue managers run into complex decisions when they try to assess the observations in the context of the project goals. For instance: looking at images they may find that the object is more or less extended than expected suggesting a change of the mapped region, or signal-to-noise of detection experiments is 3 or 4 sigma possibly suggesting an increase of observing time and fewer objects. Even more basic: a decision whether to continue a partially finished observation versus starting on another

[6]It is my understanding that most orbiting telescopes have rather rigid observing schedules on time-scales of weeks. It may well be that in the future they will be moved to a more interactive mode as well.

target in the project with a higher priority and which has become only recently accessible. While one can adopt the policy to execute observing plans by the letter, the P.I. can rightfully object that issues like these can not blindly be ignored without damaging the project. In my opinion the only solution is to put a P.I. in the position to decide these issues for him/herself and to make feedback an integral part of dynamic scheduling.

Dynamic scheduling has resulted in a measurable increase in the number of completed projects at the JCMT, in particular ones which require exceptional weather conditions. In the course of its implementation we have found that there is an overall tendency to put too much emphasis on tools and utilities and not enough on policies and procedures, in particular as to what happens after observations for a particular project have started. For a successful, efficient, and sustainable implementation of dynamic scheduling the overall policy needs to include feedback with the P.I. as an integral part.

Acknowledgments. I thank the staff of the JAC for a very useful ongoing dialog regarding dynamic scheduling and in particular Frossie Economou, Tim Jenness, John Davies, Andy Adamson, and Ian Robson for their contributions and support of these developments. I also want to thank Mark Holdaway (NRAO), Phil Puxley (Gemini), and Steve Scott (OVRO, Caltech) for their valuable and insightful comments over the years. I thank the users of the JCMT for their understanding and cooperation in trying something different and unproven. Finally, but not least, I thank Andy Adamson and John Davies for their critical comments on and valuable suggestions for this paper.

References

Bridger, A., Wright, G., Tan, M., Pickup, A., Economou, F., Currie, M., Adamson, A., Rees, N., & Purves, M. 2000, this volume, 467

Economou, F., Bridger, A., Wright, G. S., Jenness, T., Currie, M. J., & Adamson, A. 1999, in ASP Conf. Ser., Vol. 172, Astronomical Data Analysis Software and Systems VIII, ed. David M. Mehringer, Raymond L. Plante, & Douglas A. Roberts (San Francisco: ASP), 11

Jenness, T. & Economou, F. 1999, in ASP Conf. Ser., Vol. 172, Astronomical Data Analysis Software and Systems VIII, ed. David M. Mehringer, Raymond L. Plante, & Douglas A. Roberts (San Francisco: ASP), 171

Wampler, S., Gillies, K., Puxley, P. & Walker, S. 1997, Proc. SPIE, 3112, 246, ed. H. Lewis