The JCMT Telescope Management System

Remo P. J. Tilanus,[1] Tim Jenness, Frossie Economou, Steve Cockayne[2]

Joint Astronomy Centre (JAC), 660 N. A'ohoku Place, University Park, Hilo, HI 96720, E-mail: rpt@jach.hawaii.edu

Abstract. Established telescopes often face a challenge when trying to incorporate new software standards and utilities into their existing real-time control system. At the JCMT we have successfully added important new features such as a Relational Database (the Telescope Management System—TMS), an online data Archive, and WWW based utilities to an, in part, 10-year old system. The new functionality was added with remarkably few alterations to the existing system. We are still actively expanding and exploring these new capabilities.

1. Introduction

The introduction of GUI-based applications and control systems, Database Management Systems (DBMS), HTML, and the World Wide Web (WWW) have set new standards in the way users expect to interact with computers and data. Living up to those standards with systems which essentially predate these developments presents a noteworthy challenge, in particular if these systems are not UNIX-based.

The real-time control system at the JCMT uses a VMS-based network where control of (and data from) the telescope is handled by a number of VAX computers. As is traditionally the case at many telescopes an observation results in a file on a disk consisting of a "header" part and "astronomical data" part. In addition to astronomically relevant items the header also contains much information about the hardware status of the telescope. This presents an annoying and very time-consuming complication in case problems arise which need tracing or monitoring over significant time-intervals: hundreds of files may need to be accessed in order to extract the relevant parameter.

In order to address these problems we have introduced at the JCMT a Telescope Management System (TMS): a DBMS-based system which keeps permanently available on-line the same information as in the file headers plus any other information deemed of interest. Another way to express the goal of the TMS is that it aims to provide an accurate image of the generalized telescope status for each instance in time. Ideally, the real-time system would interact with the TMS to optimize the real-time performance since from the history con-

---

[1]Netherlands Foundation for Research in Astronomy (NFRA)

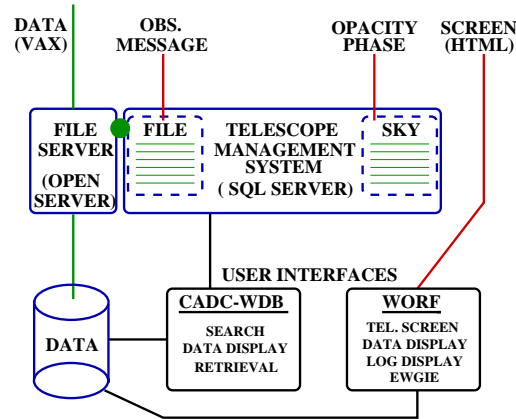[2]Now at Canadian Astronomy Data Centre (CADC)

Figure 1.   The JCMT Telescope Management System (UNIX).

tained in the TMS it should be possible to make reasonably accurate predictions for settings of sub-systems which (slowly) drift in time.

2.   The Telescope Management System

The TMS is based upon Sybase, a relational DBMS, installed on a UNIX host at the JCMT and consists of a set of linked "flat" tables which together contain all the items one wants to keep track of. A key strategy in its design stems from the realization that a telescope in general consists of a number of discrete sub-systems which each change with their own time-constant. Quite often these logical subsystems correspond to physical subsystems of the telescope, but not always: a camera, which otherwise retains its configuration during the night may contain a filter wheel which changes with each observation.

The most efficient and accurate implementation of a TMS is to create a table for each of the (logical) subsystems: when the subsystem changes its configuration only one table needs a new entry, while the others do not change. The telescope can be regarded as providing a continuous, asynchronous stream of information which gets mapped onto the appropriate tables in the TMS. In such a scenario, the astronomical data is just another package delivered, albeit an important one. The complete telescope status can thus be reconstructed from the tables by retrieving records from each table corresponding to the correct time-interval (although a smart use of "indexes" (yuk: "indices" please) rather than the timestamps will in practise be a more efficient approach).

The TMS actually installed at the JCMT represents a compromise between the ideal system described above and speed of implementation. In particular we wanted to add a TMS with minimal changes to the existing real-time system. Most information is read out of the traditional file header (rather than directly from the real-time system) and uploaded to the TMS once the observation is finished. Consequently, the TMS retains the telescope status valid per observation rather than for every instant in time. Upon the completion of the observation and the close of the observation file on the VAX the real-time systems sends a notification to the TMS on the UNIX host which immediately copies the file
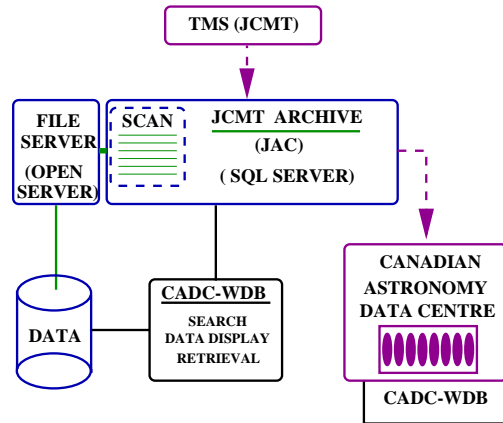
Figure 2.   The JCMT Archive.

into a disk archive and uploads the header information to the DB. An exception are the opacity and seeing monitors (and shortly the weather station), which log directly to the DB independent of observations.

In summary, three components were added to the real-time control system: the notification of the TMS via an observation completion message, the direct logging of the atmospheric conditions, and, thirdly, an HTML dump of the telescope status screen at regular intervals for use with the WWW.

Figure 1 schematically illustrates the TMS on the UNIX host. The atmospheric conditions are logged directly into the corresponding DB tables. The insertion of the observation completion message into a DB table causes a "trigger" to be activated which results in the copying of the data file and the upload of the header information to the DB using a Sybase OpenServer. The OpenServer gives the ability to extend the SQL Server with custom processing, e.g., the transparent delivery of data to SQL Clients even while the data resides on disk rather than within the DB. Another application is the sending of a mail notification if cryostat temperatures exceed a certain limit indicating the need to refill cryogens.

3.   The User Interface

Any standard SQL-based tool can be used to interact with the TMS and retrieve and display information. However, our general user interface is WDB (Rasmussen 1995) developed by ESO and the Canadian Astronomy Data Centre (CADC) as a Web-based archive query tool (Figure 1). It currently lacks the ability to plot any item against any other item but we hope to add that capability shortly using an existing Java applet. Before WDB we used the Starcat interface which was superb in this respect through the use of Xmgr. In order to extend the use of WDB from astronomical archives only to a general interface to (Sybase-based) DBs we use custom data (pre-)viewing[3] and retrieval tools and

---

[3]The TMS shares many of the custom modules with WORF2

added user authentication to safeguard the proprietary nature of some of the information. The use of WDB illustrates a very powerful feature of the TMS: since it is based upon a commercial SQL DBMS very little work is needed to provide adequate and versatile user applications.

4. The JCMT Archive

In addition to the TMS, a Sybase-based Astronomical Archive has been installed at the JAC in Hilo (Figure 2) and CADC. Astronomically relevant header information is extracted from the TMS into a science-oriented archive together with the actual data on disk (and, in the future, CD-ROM). The JCMT Archive has been developed in collaboration with the CADC and also uses WDB as a frontend. An automatically synchronised copy of the Archive tables is maintained at the CADC which is the main entry point for searches involving non-proprietary data.

5. Conclusion and Future Developments

A noteworthy aspect of the new features, such as the TMS and WORF2 (Jenness, Economou, & Tilanus 1997), added to the UNIX-based systems is that a state-of-the-art look-and-feel has been built around an existing (VMS-based) real-time system while requiring the latter only to send a single message to the TMS when an observation has finished. Although currently the TMS is not used by the real-time system itself (a possible future development), it provides the observatory staff and observers with the general diagnostic capabilities of a fully integrated RDBMS. Many features, especially the Web based utilities, required little local development and often took a few days rather than months to implement.

In the future we expect Web based utilities to become more integrated with the actual process of observing at telescopes. In particular the on-line availability of remote data archives and catalogs will make it possible to instantaneously correlate new data with existing information. Examples are marking known source positions on a real-time display, solving for plate-solutions, identification of spectral features and overlaying data from different telescopes. With astronomically relevant information increasingly available at one's "finger-tips" at the telescope observing once again should encompass "observation" rather than just data collection followed by a analysis stage much later.

References

Jenness, T., Economou, F., & Tilanus, R. P. J. 1997, this volume, 401

Rasmussen, B. F. 1995, in ASP Conf. Ser., Vol. 77, Astronomical Data Analysis Software and Systems IV, ed. R. A. Shaw, H. E. Payne, & J. J. E. Hayes (San Francisco: ASP), 72