

Data Management at the UKIRT and JCMT

Tim Jenness^a and Frossie Economou^a

^aJoint Astronomy Centre, 660 N. A‘ohōkū Place, Hilo, HI, 96720, U.S.A.

ABSTRACT

For more than a decade the Joint Astronomy Centre has been developing software tools to simplify observing and make it possible to use the telescopes in many different operational modes. In order to support remote operations the data handling systems need to be in place to allow observation preparation, flexible queue scheduling, data quality pipelines and science archives all to be connected in a data-driven environment. We discuss the history of these developments at UKIRT and JCMT and how the decision to combine software development at both telescopes led each to get features that they could not have justified if they were treated independently.

Keywords: JCMT, UKIRT, eSTAR, Data Handling, Pipelines

1. INTRODUCTION

The Joint Astronomy Centre runs two distinctly different telescopes. The James Clerk Maxwell Telescope (JCMT) is the world’s largest sub-millimetre telescope with a 15-m primary mirror. The United Kingdom Infrared Telescope (UKIRT) is a 3.8-m infrared telescope. Both are sited on Mauna Kea in Hawaii and both telescopes have been operating for many years (UKIRT since 1979 and JCMT since 1987) and developed many legacy systems. In the late 1990s the software groups were merged and code re-use and shared support became an important driver for the future.

2. DATA PIPELINES

In the mid-1990s data reduction software was written directly for the instrument with no intent to re-use when the instrument was retired. At UKIRT CGS4DR^{1,2} delivered near-publication quality data for CGS4^{3,4} and was well regarded at the time. Unfortunately its design was inflexible to change and was highly dependent on the format and headers associated with CGS4 data. At JCMT SCUBA⁵ was delivered in 1996 with a real-time data reduction system running on a VAX/VMS system that was highly integrated with the observing system but had no offline capability at all. Adding new features such as sky subtraction became difficult to debug and the ability to regenerate the map offline became very important as new CPU-intensive map-making algorithms were developed. The resulting package, SURF,⁶ could be used by the observer at the summit to try to do a better job at making the map but shell scripts had to be written to bulk process observations. The JCMT heterodyne systems^{7,8} were a little better off because the correlators (for example AOSC and the DAS) provided a buffer between the users and the frontend but the extent of the data processing was limited and a specialist offline tool (SPECX) reading a niche file format (GSDD^{9,10})* was provided to allow the observer to process their own data in near real-time if they were so inclined.

During this time software development effort in the UK moved from VAX/VMS to Solaris and Digital Unix and budgets became tighter so there was an opportunity to re-implement software systems whilst reducing the cost for delivering data reduction software with new instruments. The ORAC¹¹⁻¹³ project at UKIRT was designed to rewrite the observatory control system to provide instrument-agnostic interfaces for observation preparation, sequencing and data reduction. The data reduction part of the project was labelled ORAC-DR^{14,15} and the aim was to provide an infrastructure for data processing using reusable processing blocks, called primitives, that represent high-level concepts such as “flatfield these images” or “combine into a single image”, and are combined into “recipes”. The system is designed to be completely generic and to ensure this the initial development included support for CGS4 and UFTI¹⁶ at UKIRT and SCUBA at JCMT. The overall design is shown in Fig. 1 1 and consists of the following features:

Further author information: (Send correspondence to T.J.)

T.J.: E-mail: t.jenness@jach.hawaii.edu

*The General Single Dish Data format (GSDD) was an offshoot of a joint MRAO, NRAO, IRAM project to provide a generic data format for single-dish telescopes.

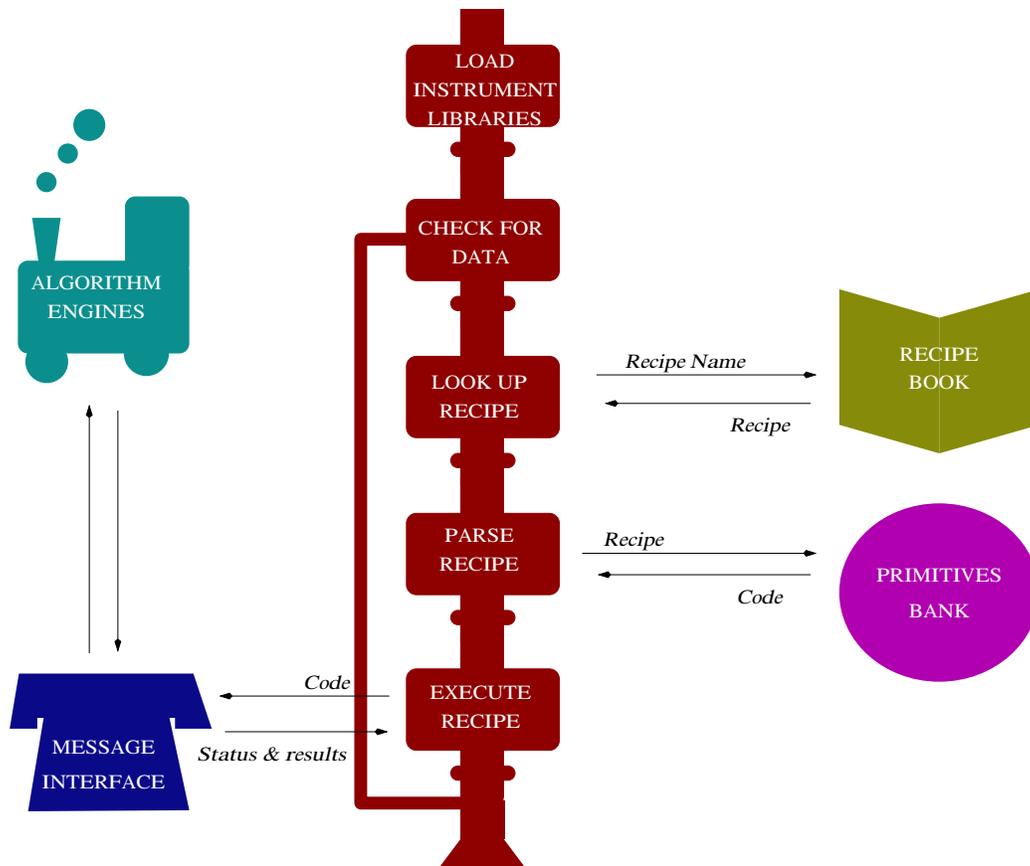


Figure 1. The design of the ORAC-DR pipeline.

- Recipes are designed to be human readable and amenable to being built up by recipe builder software. No code should appear at the recipe level although primitives are allowed to take parameters to control behaviour. Additionally a recipe parameter file can be supplied to the pipeline to control global behaviour without requiring the recipe to be edited.
- Primitives are written in code (in this case Perl) and should implement a simple well-understood concept.
- The parser builds up an executable recipe from the primitives. Primitives are read at run-time and cached. If a primitive changes the recipe will change without having to rerun the pipeline. Additionally the parser automatically inserts status handling, debugging and timing code.
- Low-level pixel processing is implemented by applications in the Starlink software collection^{17–20} controlled using the ADAM message bus.^{7,21} Starlink was chosen because of the robust messaging system and well-defined error handling as well as support for variance propagation and automatic history writing.²² At the time, 1997, IRAF²³ was not as robust and native support for numeric arrays in languages such as Perl or Python was very new. The pipeline architecture allows any messaging system and algorithm engine to be used and DRAMA²⁴ support has been implemented as an alternative to ADAM; demonstrating flexibility.
- Pipeline state is stored in “frame” objects associated with each observation. This class changes depending on the instrument and handles file naming conventions, knowledge of file structures and metadata handling. There are also classes for handling a collection of frames, a “group”, and calibration systems. The calibration classes, for example, know which flatfield or dark is relevant for a particular frame or how extinction correction should be applied.

- As far as possible the primitives using generic metadata rather than using the actual FITS header stored in a particular file. When a file header is read the header is translated[†] and both the original and translated header are stored in the frame object.

The most important design consideration is that the pipeline should never pause to wait for input from the user. All the metadata needed to process the observation must be available in the data headers and they must be correct. Algorithms must also be developed to do as good a job as possible without needing every recipe to be tuned in advance. For example, the ACSIS²⁵ pipeline does not need to be told where the baseline is when doing baseline subtraction. All that is required is that the pipeline is given an idea of whether it is processing broad line, narrow line or line forest data and it picks an appropriate algorithm. The narrow line recipe uses the CUPID²⁶ clump-finding software to first find the emission regions before generating a mask that is used to do high quality baseline subtraction.

In the 13 years since the pipeline was released it has been used for CGS4, IRCAM, UFTI, UIST,²⁷ MICHELLE²⁸ and WFCAM²⁹ on UKIRT;^{30–32} SCUBA, ACSIS and SCUBA-2 on JCMT;^{33–36} and IRIS2³⁷ on the AAT.³² Proof-of-concept recipes have also been developed for ESO and Gemini instruments.^{32,38,39}

3. FLEXIBLE SCHEDULING

Submillimetre observing is very sensitive to the weather conditions and in particular the amount of water in the atmosphere. In the first few years of JCMT observing a “classical” scheduling approach was used where a few shifts were allocated to an observer and they would get their data or not depending on the weather. In some cases they would observe low frequency projects in the best conditions or even observe their own backup project if the weather was poor. It was realised that this approach was not the most efficient use of the telescope and plans were adopted to start a flexible approach to observing where the people wanting the best conditions would be asked to come out to observe but they would be given more nights than they needed and be allowed to choose when to take their own data if they agreed to take data for someone else if the weather was not good enough.⁴⁰ In 1996 JCMT began a flexible observing system⁴¹ that involved multiple projects being allocated blocks of time allowing the flexing inside those blocks. As more experience was gained an implementation of this scheme required the PIs to fill in a form saying what they would like to observe and the observer going through a folder looking for a relevant backup proposal.⁴² Of course this system had a large overhead for everyone and sometimes led to heated email exchanges as the remote PI complained that their data was not taken with as much care as would have been the case if they were present. This was followed by a prototype system at JCMT that stored all the targets to be observed in a database along with weather requirements and allocated priority. No attempt was made to be able to send the results of a query to the telescope so it was still necessary to define the observation at the telescope. Meanwhile, UKIRT had been thinking about queue scheduling for some time⁴³ and were running a service programme based on a PostgreSQL database that could track whether the data had been accepted by the PI.

In June 1998 an initial project plan for a JCMT Observation Management Project⁴⁴ was developed but the full implementation stalled as development effort was transferred to other higher priority projects. By late 2000 it was clear that a single unified system should be adopted by both telescopes to minimise the overhead of supporting two disparate systems (the JCMT and UKIRT software groups had recently merged) and to provide a system that could not only tell you what to observe but also send the observation to the telescope without the observer having to do anything. The JAC Observation Management Project (OMP)⁴⁵ was the result of this merger heavily influenced by the earlier JCMT OMP and the developments at UKIRT delivered as part of the ORAC project. The requirements for the OMP were:

- An observation database populated by the astronomers using an observation preparation tool.⁴⁶ The science programme is split up into scheduling entities known as Minimum Schedulable Blocks (MSBs).
- A tool enabling the database to be queried and the results to be constrained by the current weather conditions and sorted by allocation priority.⁴⁷
- A translation layer that will convert the abstract science description of the MSB into something understood by the telescope observatory control system (OCS).

[†]See the `Astro::FITS::HdrTrans` Perl module at <https://github.com/bradcavanagh/perl-Astro-FITS-HdrTrans>

- A queue that will send translated output to the OCS one observation at a time and realise when an MSB has been completed.⁴⁸ The queue then gives an opportunity for the observer to accept or reject the MSB and associate this with any comment on data quality.
- A collection of dynamically generated web pages (the “feedback system”) to allow the remote observer to check on program completion, to comment on data quality and to retrieve the data files.⁴⁹ This system also lets the support astronomer contact the PI to provide feedback and suggestions on the observing plan.
- Data reduced by the ORAC-DR pipeline should be made available to the PI over the web as thumbnail images of maps or spectra, building on the WOLF system deployed for the older systems.^{50,51}
- Basic scheduling logic to allow MSBs to be disabled if other MSBs were observed first. This was implemented using OR and AND folders. An OR folder could be configured containing a mix of MSBs and AND folders containing MSBs. If N of them were observed the remaining entries would be disabled.
- Importantly for UKIRT the system had to work in purely classical mode if required.

The UKIRT ORAC system had been delivered with an observation preparation tool (OT)¹³ derived from an early release of a (since completely replaced) Gemini OT⁵² so it was decided early on in the project that JCMT should adopt the same system and extend the XML schema with features to support site quality (submm opacity, seeing, moon state), logic folders and MSB folders.

UKIRT tested the full system in March 2002 to do their service programme and the full system was live on JCMT by July 2002 for SCUBA. The project was completed in September 2003. The heterodyne instruments took a bit longer and an interim system was developed to allow the OMP to be used with the legacy interfaces. The translation system for the DAS backend would pop up a list of instructions for the telescope operator to follow (e.g. tune to this frequency, set up a 5 x 5 grid, use the 250 MHz correlator settings) and when they were finished they would click a button to tell the system that the MSB had been observed. This was in use until ACSIS was commissioned in 2005/6.

The OMP system changed the way both telescopes ran and resulted in improved completion statistics of highly ranked projects.^{53,54}

4. ESTAR RAPID RESPONSE

Once we had a data reduction pipeline and a flexible scheduling system it became possible for the telescope to be used in novel ways not previously considered. One of our PIs was interested in fast followup of gamma ray bursts (GRB) and had been talking to the eSTAR project. eSTAR⁵⁵ is a system of autonomous software agents designed for use in robotic telescope networks. The system worked by having one agent monitoring GRB events, for example from SWIFT, and it would then send out an observing request to telescope agents asking whether they could do an immediate followup observation. Any number can respond with a “yes” and the agent would then send the request and wait to be told that the data was available for analysis. The ultimate aim was that the agent would analyse the data from the telescopes and then decide whether further observing was warranted.

By early 2003 the eSTAR system was working for test robotic telescopes but it was realised that a UKIRT telescope running the OMP and ORAC-DR could be made to look identical to a robotic telescope.⁵⁶ The system worked by installing an “embedded agent” at the JAC to handle the eSTAR requests. The PI prepares a science program that has a dummy target position but fully describes the rest of the observing parameters. When the agent receives a request from a particular eSTAR user agent the user information is translated to a local OMP project code and the agent can check to see if time remains on the project for observing and if the source is going to be up in a few hours. If the user agent indicates that the embedded agent should do the observation the template science program is retrieved, the MSB is copied and the relevant target is inserted. UKIRT is not robotic and may well be in the middle of an MSB and so can not instantly slew to the GRB target. What actually happens is that an alert goes off in the control room indicating a time-sensitive observation has entered the observing database and the operator can decide whether to abort the current MSB or continue to observe it.

Once the decision to observe it is taken the operator can locate a relevant guide star and adjust exposure times as necessary. The final piece of the puzzle is that the observing system writes a special header item indicating that the files are associated with an eSTAR trigger and storing the unique eSTAR identifier for the observation. When the data reduction

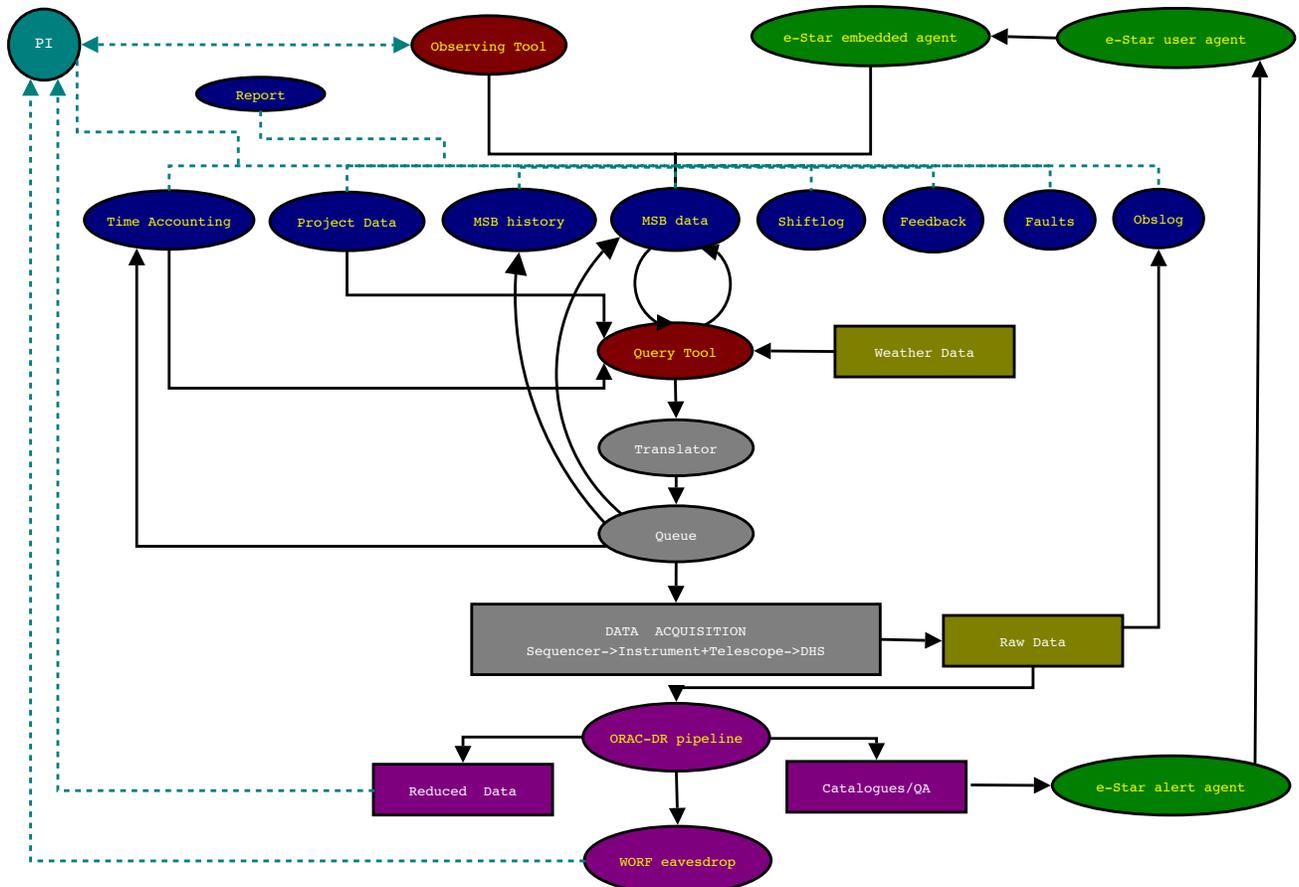


Figure 2. OMP and eSTAR system architecture.

pipeline reduces the data at the summit and generates the mosaic it sends the reduced data to the embedded agent which forwards the result back to the PIs user agent which can then decide whether to schedule a followup.

The first demonstration of eSTAR at UKIRT was carried out in August 2003 and the system triggered an automatic follow up observation of GRB 050716 in 2005.⁵⁷ The observation was queued in the observing system 48 seconds after it was detected by SWIFT but due to operational difficulties the telescope could only observe 56 minutes after initial detection. In 2009 GRB 090423 was detected within 20 minutes of the alert using eSTAR⁵⁸ (the weather was not good so it took a while to open the telescope).

The flow chart showing how the OMP interacts with eSTAR is shown in Fig. 2.

5. DATA ARCHIVING

Given history and the different funding models for UKIRT and JCMT these telescopes have different data archive centres. This section describes each telescope in turn.

5.1 JCMT

The JCMT archives all data at the Canadian Astronomy Data Centre. We use the standard CADC e-Transfer system initially developed for the Gemini Science Archive.⁵⁹ The data are transferred to CADC within a few minutes of the files appearing on disk. The key design issue is that metadata are sent to CADC before the file itself is transferred and CADC reject the file if the metadata does not match a pre-defined data dictionary. This allows file header problems to be spotted quickly and files must be fixed by the telescope support staff before attempting to re-ingest them.

A single JCMT observation can consist of multiple files and with SCUBA-2 eight separate data acquisition computers write out files in parallel every 30 seconds during an observation. These files have a shared control task which allows certain headers to be synchronized to make it simple to work out that they are related files. As each file appears on disk a daemon checks the data directory by looking in a small text file, called a “flag” file, to see if any new files have been created. This flag file is only updated when the data acquisition has finished writing a file and it is ready to be picked up by the data reduction pipeline or archiving software. The archive daemon then reads the header from the file and decides whether this is a new observation or part of an existing observation.

The JCMT data archive consists of four database tables. `COMMON` contains the metadata common to every instrument. As well as the items from the file headers we also store the details of the area of sky covered by the observation. There are then instrument-specific tables which at JCMT consists of a `SCUBA2` table and an `ACISIS` table. These tables can have multiple rows per observation, known as “sub-systems”, because with SCUBA-2 we write data at two wavelengths simultaneously and for the ACSIS correlator we can write out up to four separate spectra from different positions in the 2 GHz band of the frontend. Finally there is a table, `FILES`, containing the names of the files associated with each row in `COMMON` and `SCUBA2/ACISIS`. The rows are linked with an observation identifier (known as an “obsid”) that is unique for each observation and a subsystem identifier (“obsidss”) which is a combination of the obsid and the subsystem number. The obsid takes the form “*instrument_run-number_date-obs*” to enable it to be unique and human-readable. An example obsid might be `scuba2_00071_20100310T101813` with a corresponding obsidss of `scuba2_00071_20100310T101813_450`.

These tables are stored in a Sybase database and automatically replicated at CADC. There is an additional table at the JCMT that tracks the state of a particular file. This state table has states for “file found”, “metadata ingested”, “row replicated”, “copied to e-transfer”, “arrived at CADC” and “error”, making it easy to determine how efficiently the system is transferring data and whether we have a blockage.

For a new observation a row must be placed in each of the relevant tables. If the new file is part of an observation that is already present in `COMMON`, any metadata associated with the end of an observation (such as `DATE-END` or the end elevation or azimuth) is updated and an entry is made in the `FILES` table and the instrument table as appropriate. Every few minutes another daemon checks with CADC to see if the local database has been replicated correctly to CADC. If the `FILES` table at CADC contains a file that is awaiting transfer a soft link is made into the e-transfer directory and the file is passed to the CADC software to be transferred to CADC. If the file is rejected for any reason the rejected file (or its soft link) is written to a rejection directory indicating the reason why (reasons include: a corrupt file, an empty file, non-standard filename, incorrect metadata, or the file is already at CADC and was not marked for replacement).

On a daily basis a further cron job runs at the telescope designed to free up disk space on the data acquisition machines. A file can only be deleted if CADC report that they have successfully ingested it.

5.2 UKIRT

UKIRT used to have two modes of operation. When WFCAM arrived the Cassegrain instruments (CGS4, UIST, UFTI) were used some of the year and WFCAM was used for the rest of the year. Budget cuts and a re-focussing on survey science meant that the Cassegrain instruments were mothballed and UKIRT now uses WFCAM full time in order to complete the UKIDSS survey.⁶⁰ The Cassegrain archive consisted of writing data to tapes and, when disk space became cheaper, storing it on a disk in Hilo. Every semester data tapes were also sent to the archive in Cambridge at the Cambridge Astronomical Survey Unit (CASU) where they were converted to FITS format.

WFCAM generates much more data than the Cassegrain instruments and the archive works slightly differently. The raw data format from WFCAM is in NDF and these files are copied to Hilo and placed on LTO-3 tapes. The raw data are deleted from the summit, oldest first, when the LTO tape has been read back and verified to be correct. On a daily basis a cron job converts the raw data to FITS format using Rice compression and the data are retrieved by CASU. When they have retrieved the data a flag file is written indicating that the cleanup jobs can remove the directory.

Storing the raw data on tape has been very useful over the years. In general we never need to read them but there have been instances where the conversion to FITS has been buggy and we have not realised for a month or two. The data were no longer on disk but could be retrieved from tape, albeit slowly, and reconverted to FITS. The main issue with our raw data archive is the changing media. Initially files were written to LTO-1 and from mid-2008 LTO-3. The LTO standard indicates that an LTO-4 drive does not need to read LTO-1 tapes and given that LTO-5 is the current standard (which will read LTO-3) it is highly likely that the LTO-1 tapes covering 2005 to 2008 will never be read, but of course the data have been fully reduced and published by WFAU in Edinburgh⁶¹ so there is no reason to need to read them now.

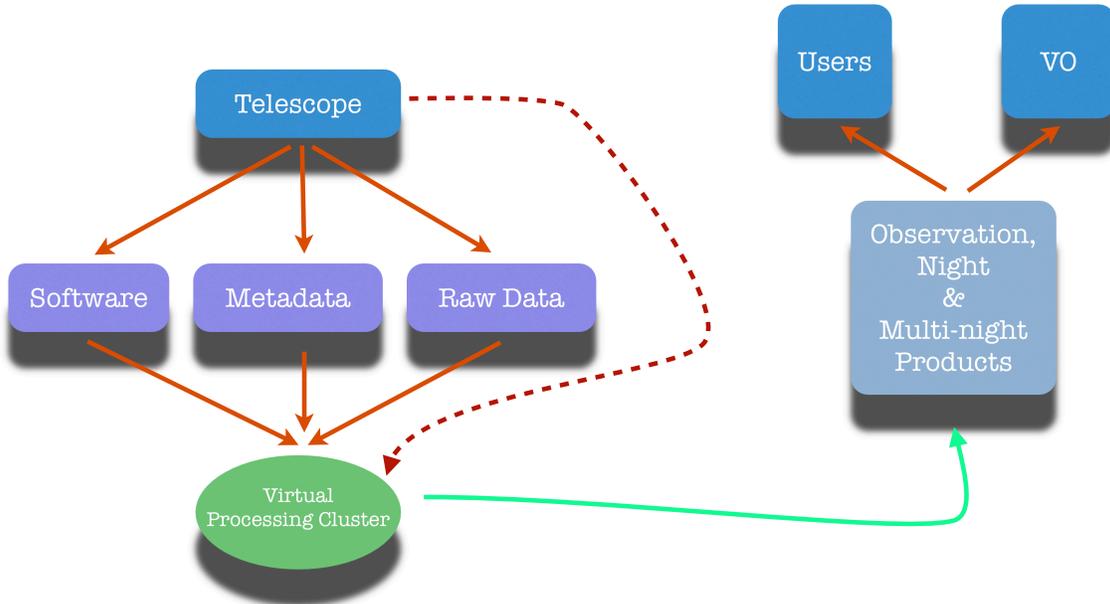


Figure 3. Conceptual diagram of the JSA. The telescope makes raw data, metadata and archive available to the data centre. It then triggers processing on the data centre’s processing nodes. Products are immediately exposed

6. ARCHIVE PROCESSING

The JCMT Science Archive (JSA)^{62–64} is a collaboration between the JAC and CADC to provide calibrated reduced products from the JCMT to the astronomy community as quickly as possible. The JAC is interested in increasing the productivity of the JCMT by providing its community with science-ready data products and by enhancing the value of the JCMT Legacy Survey programme. CADC provide their knowledge of managing large data centres and integrating into the virtual observatory and also learn how to handle datasets that are very different from the more common 2-D optical/IR data.

The JSA is separated into two components with well-defined interfaces. CADC are the “data store” and service requests for data files to be made available to the processing system and allow products from the data processing to be ingested into the reduced products database. Again metadata management is critical with CADC refusing to accept any files that do not follow the rules. CADC also provide the public interface to the data and handle authentication issues and proprietary periods. The JAC are entirely responsible for data processing and for deciding which data files should be processed at CADC.

There are four types of products in the JSA: single observations (“*obs*”), combined observations from a single night, combined observations from a single project and a combination of all relevant public data for a particular field. The last product has not yet been generated although all generated products do become public when the proprietary period expires.

Each job submitted to the processing system is assigned a unique “association identifier” that links the job with the product and is derived from the grouping algorithm used to determine that the submitted files are related. If more observations are taken satisfying a particular query the association identifier will not change when those observations are included in the newly recalculated product. The conceptual diagram of the JSA is show in Fig. 3.

The JSA does not have periods set aside for data releases. Every time data are reprocessed by the JSA pipeline the earlier version is removed and replaced with the current version. Headers in the data files indicate the versions of the software used to generate them (via SHA1 identifiers from the `git`[‡] source code repository) along with the data processing date so it is possible to track which version of a product has been used in a publication. This scheme allows us to re-reduce data when there is an update to the data pipeline. Additionally having this infrastructure makes it very easy to test the data pipeline by doing a bulk reprocessing on a development system.

[‡]<http://git-scm.com/>

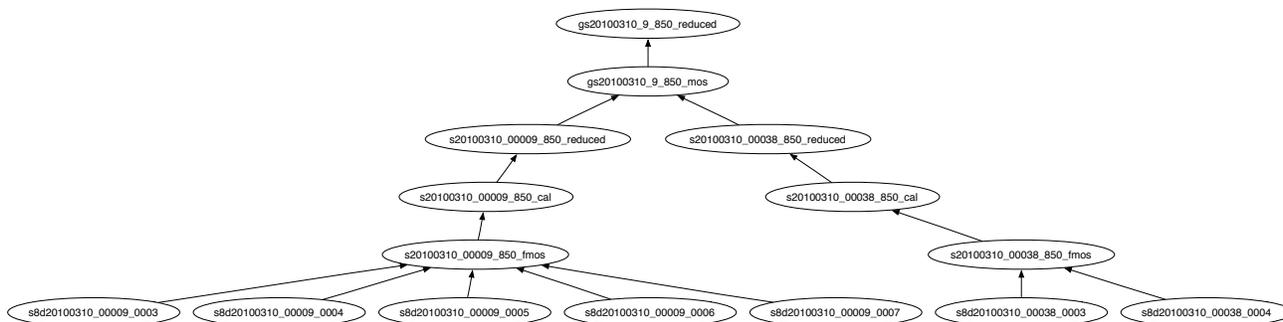


Figure 4. Provenance tree from two SCUBA-2 observations combined into a single mosaicked product. There are multiple data files from each observation.

6.1 Provenance Tracking

One important requirement placed on the data processing system by the science archive is to track provenance. Every product archived in the JSA has to have full knowledge of its parent data files and the operations performed on that data as well as the ability to walk through the information to determine the original raw data files that went into the product. The NDF history mechanism²² only tracks a single path through the provenance so a new system was needed to track full provenance. Additionally, it was decided that modifying each processing step to explicitly propagate provenance was not acceptable so a low-level solution was required that would handle provenance propagation automatically.

Most Starlink applications use the NDG library⁶⁵ to obtain a group of input or output files. The NDG library is a layer on top of NDF and makes use of a system that can register callbacks with the NDF library. When a file is opened the information is retained by the library and when an output file is closed the provenance is propagated to that output file unless the provenance has been explicitly augmented by the application. This allows applications where input data flows to output files to handle provenance transparently but more complex applications can override with specialist knowledge. The SMURF⁶⁶ routines were modified to seed a root provenance with the correct information (the “obsidss”) and this information then flows to the JSA products automatically. An example provenance tree is shown in Fig. 4 where two SCUBA-2 observations are combined by the pipeline into a single mosaicked image.

The pipeline provenance information includes the parent file and the complete command details used to generate the file. The CADC provenance infrastructure requires that the FITS header includes details of the parent provenance and the root ancestor “obsidss”. The parent provenance must refer to a file archived at CADC so it is necessary to compress the provenance before ingestion by removing temporary files from the tree.

7. CONCLUSIONS

The combination of the pipeline, flexible scheduling and science archive enables the astronomer to be involved in the observing to an extent that has been impossible since many telescopes moved to queue-scheduled observing. The reduced data are available within a few hours and the astronomer can adjust their science program and resubmit it before the next observing night begins.

There is also direct feedback from the summit data quality assessment to the archive processing. If the observer marks an observation as questionable or bad the science archive will not combine the data into the combined group. If the data are marked junk, hopefully a rare occurrence, the data will not be processed at all. The science archive pipeline itself does more detailed quality assessment of the data and is able to mark the observation status in the OMP accordingly. This feedback loop allows instrument scientists to get rapid feedback of instrument issues before they become serious.

A fully automated data-driven data reduction pipeline is pivotal to the ongoing development of the data management system at the JAC. All new instruments must be delivered with a DR pipeline based on ORAC-DR and must also use the standard OCS interfaces to allow simple integration into the OMP.

As time passes UKIRT and JCMT have adopted more and more shared infrastructure.⁶⁷ Starting with the release of ORAC-DR in 1998 the JAC has continued to benefit from sharing the software and each telescope has received updates that would not have been possible if we had retained separate software groups.

The JSA brings together the pipeline and flexible scheduling to allow fast feedback to astronomers giving them fully reduced observations within a few hours of the data being taken. Many sites do not have the computing resources required to process or even transfer the 100 GB of data taken each night on a reasonable timescale. This then gives them an immediate opportunity to adjust their science programme for subsequent nights and provides them with multi-night coadds that can be used to jumpstart paper writing.

ACKNOWLEDGMENTS

The James Clerk Maxwell Telescope is operated by the Joint Astronomy Centre on behalf of the Science and Technology Facilities Council of the United Kingdom, the Netherlands Organisation for Scientific Research, and the National Research Council of Canada. The United Kingdom Infrared Telescope is operated by the Joint Astronomy Centre on behalf of the Science and Technology Facilities Council of the United Kingdom.

REFERENCES

- [1] Stewart, J. M., Beard, S. M., Mountain, C. M., Pickup, D. A., and Bridger, A., “Automatic On-line Data Analysis for an Infrared Long Slit Array Spectrometer,” in [*Astronomical Data Analysis Software and Systems I*], D. M. Worrall, C. Biemesderfer, & J. Barnes, ed., *Astronomical Society of the Pacific Conference Series* **25**, 479–+ (1992).
- [2] Daly, P. N., “Porting CGS4DR to Unix,” in [*Astronomical Data Analysis Software and Systems IV*], R. A. Shaw, H. E. Payne, & J. J. E. Hayes, ed., *Astronomical Society of the Pacific Conference Series* **77**, 375–+ (1995).
- [3] Mountain, C. M., Robertson, D. J., Lee, T. J., and Wade, R., “An advanced cooled grating spectrometer for UKIRT,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], D. L. Crawford, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **1235**, 25–33 (July 1990).
- [4] Wright, G. S., Mountain, C. M., Bridger, A., Daly, P. N., Griffin, J. L., and Ramsay Howat, S. K., “CGS4 experience: two years later,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], A. M. Fowler, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **1946**, 547–557 (Oct. 1993).
- [5] Holland, W. S., Robson, E. I., Gear, W. K., Cunningham, C. R., Lightfoot, J. F., Jenness, T., Ivison, R. J., Stevens, J. A., Ade, P. A. R., Griffin, M. J., Duncan, W. D., Murphy, J. A., and Naylor, D. A., “SCUBA: a common-user submillimetre camera operating on the James Clerk Maxwell Telescope,” *MNRAS* **303**, 659–672 (Mar. 1999).
- [6] Jenness, T. and Lightfoot, J. F., “Reducing SCUBA Data at the James Clerk Maxwell Telescope,” in [*Astronomical Data Analysis Software and Systems VII*], R. Albrecht, R. N. Hook, & H. A. Bushouse, ed., *Astronomical Society of the Pacific Conference Series* **145**, 216–+ (1998).
- [7] Bridger, A., Prestage, R., and Fuka, M., “Cooperating software tasks for the control of large telescopes,” *Nuclear Instruments and Methods in Physics Research A* **293**, 148–153 (Aug. 1990).
- [8] Fuka, M. A., Kelly, D., and Prestage, R. M., “Upgrading the JCMT control system to platform independence,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], H. Lewis, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **3351**, 148–156 (May 1998).
- [9] Hewish, A. et al., “Mullard Radio Astronomy Observatory, University of Cambridge (Reports of Observatories),” *QJRAS* **27**, 675–+ (Dec. 1986).
- [10] Jenness, T., Tilanus, R., Meyerdierts, H., and Fairclough, J., “The global section datafile (GSD) access library,” tech. rep., Starlink User Note 229, Starlink Project (1999).
- [11] Bridger, A., Economou, F., Wright, G. S., and Currie, M. J., “Observing control and data reduction at the UKIRT,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], P. J. Quinn, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **3349**, 184–194 (July 1998).
- [12] Bridger, A., Wright, G. S., Economou, F., Tan, M., Currie, M. J., Pickup, D. A., Adamson, A. J., Rees, N. P., Purves, M., and Kackley, R., “ORAC: a modern observing system for UKIRT,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], H. Lewis, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **4009**, 227–238 (June 2000).
- [13] Wright, G. S., Bridger, A. B., Pickup, D. A., Tan, M., Folger, M., Economou, F., Adamson, A. J., Currie, M. J., Rees, N. P., Purves, M., and Kackley, R. D., “An Observer’s View of the ORAC System at UKIRT,” in [*Astronomical Data Analysis Software and Systems X*], F. R. Harnden Jr., F. A. Primini, & H. E. Payne, ed., *Astronomical Society of the Pacific Conference Series* **238**, 137–+ (2001).

- [14] Economou, F., Bridger, A., Wright, G. S., Jenness, T., Currie, M. J., and Adamson, A., “ORAC-DR: Pipelining With Other People’s Code,” in [*Astronomical Data Analysis Software and Systems VIII*], D. M. Mehringer, R. L. Plante, & D. A. Roberts, ed., *Astronomical Society of the Pacific Conference Series* **172**, 11–+ (1999).
- [15] Cavanagh, B., Jenness, T., Economou, F., and Currie, M. J., “The ORAC-DR data reduction pipeline,” *Astronomische Nachrichten* **329**, 295–297 (Mar. 2008).
- [16] Roche, P. F., Lucas, P. W., Mackay, C. D., Etedgui-Atad, E., Hastings, P. R., Bridger, A., Rees, N. P., Leggett, S. K., Davis, C., Holmes, A. R., and Handford, T., “UFTI: the 0.8 - 2.5 μm fast track imager for the UK infrared telescope,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], M. Iye & A. F. M. Moorwood, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **4841**, 901–912 (Mar. 2003).
- [17] Dunford, E. and Terrett, D. L., “Recent Developments in STARLINK,” *Bulletin d’Information du Centre de Donnees Stellaires* **31**, 69–+ (Nov. 1986).
- [18] Warren-Smith, R. F. and Wallace, P. T., “The STARLINK Software Collection,” in [*Astronomical Data Analysis Software and Systems II*], R. J. Hanisch, R. J. V. Brissenden, & J. Barnes, ed., *Astronomical Society of the Pacific Conference Series* **52**, 229–+ (Jan. 1993).
- [19] Wallace, P. T. and Warren-Smith, R. F., “Starlink: Astronomical computing in the United Kingdom,” *Information Handling in Astronomy* **250**, 93–+ (2000).
- [20] Jenness, T., Berry, D. S., Cavanagh, B., Currie, M. J., Draper, P. W., and Economou, F., “Developments in the Starlink Software Collection,” in [*Astronomical Data Analysis Software and Systems XVIII*], D. A. Bohlender, D. Durand, & P. Dowler, ed., *Astronomical Society of the Pacific Conference Series* **411**, 418–+ (Sept. 2009).
- [21] Allan, P. M., “The ADAM software environment,” in [*Astronomical Data Analysis Software and Systems I*], D. M. Worrall, C. Biemesderfer, & J. Barnes, ed., *Astronomical Society of the Pacific Conference Series* **25**, 126–+ (1992).
- [22] Warren-Smith, R. F. and Berry, D. S., “NDF – routines for accessing the extensible N-Dimensional data format,” tech. rep., Starlink User Note 33, Starlink Project (2010).
- [23] Tody, D., “IRAF in the Nineties,” in [*Astronomical Data Analysis Software and Systems II*], R. J. Hanisch, R. J. V. Brissenden, & J. Barnes, ed., *Astronomical Society of the Pacific Conference Series* **52**, 173–+ (Jan. 1993).
- [24] Bailey, J. A., Farrell, T., and Shortridge, K., “DRAMA: an environment for distributed instrumentation software,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], P. T. Wallace, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **2479**, 62–68 (June 1995).
- [25] Buckle, J. V., Hills, R. E., Smith, H., Dent, W. R. F., Bell, G., Curtis, E. I., Dace, R., Gibson, H., Graves, S. F., Leech, J., Richer, J. S., Williamson, R., Withington, S., Yassin, G., Bennett, R., Hastings, P., Laidlaw, I., Lightfoot, J. F., Burgess, T., Dewdney, P. E., Hovey, G., Willis, A. G., Redman, R., Wooff, B., Berry, D. S., Cavanagh, B., Davis, G. R., Dempsey, J., Friberg, P., Jenness, T., Kackley, R., Rees, N. P., Tilanus, R., Walther, C., Zwart, W., Klapwijk, T. M., Kroug, M., and Zijlstra, T., “HARP/ACSIS: a submillimetre spectral imaging system on the James Clerk Maxwell Telescope,” *MNRAS* **399**, 1026–1043 (Oct. 2009).
- [26] Berry, D. S., Reinhold, K., Jenness, T., and Economou, F., “CUPID: A Clump Identification and Analysis Package,” in [*Astronomical Data Analysis Software and Systems XVI*], R. A. Shaw, F. Hill, & D. J. Bell, ed., *Astronomical Society of the Pacific Conference Series* **376**, 425–+ (Oct. 2007).
- [27] Ramsay Howat, S. K., Todd, S., Leggett, S., Davis, C., Strachan, M., Borrowman, A., Ellis, M., Elliot, J., Gostick, D., Kackley, R., and Rippa, M., “The commissioning of and first results from the UIST imager spectrometer,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], A. F. M. Moorwood & M. Iye, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **5492**, 1160–1171 (Sept. 2004).
- [28] Glasse, A. C., Atad-Etedgui, E. I., and Harris, J. W., “Michelle Midinfrared Spectrometer and Imager,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], A. L. Ardeberg, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **2871**, 1197–1203 (Mar. 1997).
- [29] Casali, M., Adamson, A., Alves de Oliveira, C., Almaini, O., Burch, K., Chuter, T., Elliot, J., Folger, M., Foucaud, S., Hambly, N., Hastie, M., Henry, D., Hirst, P., Irwin, M., Ives, D., Lawrence, A., Laidlaw, K., Lee, D., Lewis, J., Lunney, D., McLay, S., Montgomery, D., Pickup, A., Read, M., Rees, N., Robson, I., Sekiguchi, K., Vick, A., Warren, S., and Woodward, B., “The UKIRT wide-field camera,” *A&A* **467**, 777–784 (May 2007).
- [30] Currie, M., Wright, G., Bridger, A., and Economou, F., “Data Reduction of Jittered Infrared Images Using the ORAC Pipeline,” in [*Astronomical Data Analysis Software and Systems VIII*], D. M. Mehringer, R. L. Plante, & D. A. Roberts, ed., *Astronomical Society of the Pacific Conference Series* **172**, 175–+ (1999).

- [31] Economou, F., Jenness, T., Cavanagh, B., Wright, G. S., Bridger, A. B., Kerr, T. H., Hirst, P., and Adamson, A. J., “Infrared Spectroscopy Data Reduction with ORAC-DR,” in [*Astronomical Data Analysis Software and Systems X*], F. R. Harnden Jr., F. A. Primini, & H. E. Payne, ed., *Astronomical Society of the Pacific Conference Series* **238**, 314–+ (2001).
- [32] Cavanagh, B., Hirst, P., Jenness, T., Economou, F., Currie, M. J., Todd, S., and Ryder, S. D., “ORAC-DR: One Pipeline for Multiple Telescopes,” in [*Astronomical Data Analysis Software and Systems XII*], H. E. Payne, R. I. Jedrzejewski, & R. N. Hook, ed., *Astronomical Society of the Pacific Conference Series* **295**, 237–+ (2003).
- [33] Jenness, T. and Economou, F., “The SCUBA Data Reduction Pipeline: ORAC-DR at the JCMT,” in [*Astronomical Data Analysis Software and Systems VIII*], D. M. Mehringer, R. L. Plante, & D. A. Roberts, ed., *Astronomical Society of the Pacific Conference Series* **172**, 171–+ (1999).
- [34] Jenness, T., Lightfoot, J. F., Holland, W. S., Greaves, J. S., and Economou, F., “SCUBA Observing Techniques and Data Reduction Pipeline,” in [*Imaging at Radio through Submillimeter Wavelengths*], J. G. Mangum & S. J. E. Radford, ed., *Astronomical Society of the Pacific Conference Series* **217**, 205–+ (2000).
- [35] Gibb, A. G., Scott, D., Jenness, T., Economou, F., Kelly, B. D., and Holland, W. S., “Design of the SCUBA-2 Quick Look Display and Data Reduction Pipeline,” in [*Astronomical Data Analysis Software and Systems XIV*], P. Shopbell, M. Britton, & R. Ebert, ed., *Astronomical Society of the Pacific Conference Series* **347**, 585–+ (Dec. 2005).
- [36] Jenness, T., Cavanagh, B., Economou, F., and Berry, D. S., “JCMT Science Archive: Advanced Heterodyne Data Products Pipeline,” in [*Astronomical Data Analysis Software and Systems XVII*], R. W. Argyle, P. S. Bunclark, & J. R. Lewis, ed., *Astronomical Society of the Pacific Conference Series* **394**, 565–+ (Aug. 2008).
- [37] Tinney, C. G., Ryder, S. D., Ellis, S. C., Churilov, V., Dawson, J., Smith, G. A., Waller, L., Whittard, J. D., Haynes, R., Lankshear, A., Barton, J. R., Evans, C. J., Shortridge, K., Farrell, T., and Bailey, J., “IRIS2: a working infrared multi-object spectrograph and camera,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], A. F. M. Moorwood & M. Iye, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **5492**, 998–1009 (Sept. 2004).
- [38] Currie, M. J., “Extending ORAC-DR for Offline Processing of ESO, INGRID, and Classic Cam data,” in [*Astronomical Data Analysis Software and Systems (ADASS) XIII*], F. Ochsenbein, M. G. Allen, & D. Egret, ed., *Astronomical Society of the Pacific Conference Series* **314**, 460–+ (July 2004).
- [39] Currie, M. J., “ISAAC Photometric Comparison of ECLIPSE Jitter and the ORAC-DR Equivalent Recipe for ISAAC,” in [*Astronomical Data Analysis Software and Systems XIV*], P. Shopbell, M. Britton, & R. Ebert, ed., *Astronomical Society of the Pacific Conference Series* **347**, 580–+ (Dec. 2005).
- [40] Robson, I., “New Modes of Observing - The Key Issues,” in [*New Observing Modes for the Next Century*], T. Boroson, J. Davies, & I. Robson, ed., *Astronomical Society of the Pacific Conference Series* **87**, 3–+ (1996).
- [41] Padman, R., “Flexible Scheduling at JCMT; an Observer’s View,” in [*New Observing Modes for the Next Century*], T. Boroson, J. Davies, & I. Robson, ed., *Astronomical Society of the Pacific Conference Series* **87**, 168–+ (1996).
- [42] Watt, G. D., “Submillimeter flexible scheduling with the JCMT,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], P. J. Quinn, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **3349**, 126–134 (July 1998).
- [43] Bridger, A. and Wright, G. S., “Queue Scheduled Observing at UKIRT - Evolution not Revolution,” in [*New Observing Modes for the Next Century*], T. Boroson, J. Davies, & I. Robson, ed., *Astronomical Society of the Pacific Conference Series* **87**, 162–+ (1996).
- [44] Tilanus, R. P. J., “Queue Scheduling at the JCMT: the Real-life Experience,” in [*Astronomical Data Analysis Software and Systems IX*], N. Manset, C. Veillet, & D. Crabtree, ed., *Astronomical Society of the Pacific Conference Series* **216**, 101–+ (2000).
- [45] Economou, F., Jenness, T., Tilanus, R. P. J., Hirst, P., Adamson, A. J., Rippa, M., Delorey, K. K., and Isaak, K. G., “Flexible Software for Flexible Scheduling,” in [*Astronomical Data Analysis Software and Systems XI*], D. A. Bohlender, D. Durand, & T. H. Handley, ed., *Astronomical Society of the Pacific Conference Series* **281**, 488–+ (2002).
- [46] Folger, M., Bridger, A., Dent, B., Kelly, D., Adamson, A., Economou, F., Hirst, P., and Jenness, T., “A New Observing Tool for the James Clerk Maxwell Telescope,” in [*Astronomical Data Analysis Software and Systems XI*], D. A. Bohlender, D. Durand, & T. H. Handley, ed., *Astronomical Society of the Pacific Conference Series* **281**, 453–+ (2002).

- [47] Rippa, M., Economou, F., Jenness, T., Hirst, P., Tilanus, R., Delorey, K., and Adamson, A., “The Infinitely Configurable Observation Query Tool,” in [*Astronomical Data Analysis Software and Systems XI*], D. A. Bohlender, D. Durand, & T. H. Handley, ed., *Astronomical Society of the Pacific Conference Series* **281**, 144–+ (2002).
- [48] Kackley, R. D., Rees, N. P., Walther, C., and Jenness, T., “The JCMT observing queue and recipe sequencer,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], H. Lewis & G. Raffi, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **5496**, 718–728 (Sept. 2004).
- [49] Delorey, K. K., Jenness, T., Cavanagh, B., and Economou, F., “When the Observer isn’t there: The OMP Feedback System,” in [*Astronomical Data Analysis Software and Systems (ADASS) XIII*], F. Ochsenbein, M. G. Allen, & D. Egret, ed., *Astronomical Society of the Pacific Conference Series* **314**, 728–+ (July 2004).
- [50] Economou, F., Bridger, A., Daly, P. N., and Wright, G. S., “Remote Eavesdropping via the World Wide Web,” in [*Astronomical Data Analysis Software and Systems V*], G. H. Jacoby & J. Barnes, ed., *Astronomical Society of the Pacific Conference Series* **101**, 384–+ (1996).
- [51] Jenness, T., Economou, F., and Tilanus, R. P. J., “Remote Eavesdropping at the JCMT via the World Wide Web,” in [*Astronomical Data Analysis Software and Systems VI*], G. Hunt & H. Payne, ed., *Astronomical Society of the Pacific Conference Series* **125**, 401–+ (1997).
- [52] Wampler, S., Gillies, K. K., Puxley, P. J., and Walker, S., “Science planning for the Gemini 8-m telescopes,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], H. Lewis, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **3112**, 246–253 (Sept. 1997).
- [53] Robson, I., “Lessons learned from four years of queue-based flexible observing with the James Clerk Maxwell telescope,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], P. J. Quinn, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **4844**, 86–93 (Dec. 2002).
- [54] Adamson, A. J., Tilanus, R. P., Buckle, J., Davis, G. R., Economou, F., Jenness, T., and Delorey, K., “Science returns of flexible scheduling on UKIRT and the JCMT,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], P. J. Quinn & A. Bridger, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **5493**, 24–32 (Sept. 2004).
- [55] Allan, A., Naylor, T., Steele, I. A., Jenness, T., Cavanagh, B., Economou, F., Saunders, E., Adamson, A., Etherton, J., and Mottram, C., “eSTAR: intelligent observing and rapid responses,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], H. Lewis & G. Raffi, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **5496**, 313–322 (Sept. 2004).
- [56] Economou, F., Jenness, T., Cavanagh, B., Adamson, A. J., and Allan, A., “Time-sensitive astronomy in non-robotic telescopes,” *Astronomische Nachrichten* **327**, 788–+ (Sept. 2006).
- [57] Tanvir, N., Lowe, K., Gledhill, T., Levan, A., Rol, E., O’Brien, P., Bannister, N., Carroll, T., Guidorzi, C., Gomboc, A., Monfardini, A., Steele, I., Mundell, C., Mottram, C., Bode, M., Smith, R., and Carter, D., “GRB 050716 - UKIRT identification of candidate afterglow,” *GRB Coordinates Network* **3632**, 1–+ (2005).
- [58] Tanvir, N., Levan, A., Kerr, T., and Wold, T., “GRB 090423: UKIRT k-band observations,” *GRB Coordinates Network* **9202**, 1–+ (2009).
- [59] Melnychuk, G., Damian, A., Goliath, S., Shao, L., Hill, N., Gaudet, S., Bohlender, D., and Aspin, C., “Gemini Science Archive Automatic Data Transfer,” in [*Astronomical Data Analysis Software and Systems XIV*], P. Shopbell, M. Britton, & R. Ebert, ed., *Astronomical Society of the Pacific Conference Series* **347**, 647–+ (Dec. 2005).
- [60] Lawrence, A., Warren, S. J., Almaini, O., Edge, A. C., Hambly, N. C., Jameson, R. F., Lucas, P., Casali, M., Adamson, A., Dye, S., Emerson, J. P., Foucaud, S., Hewett, P., Hirst, P., Hodgkin, S. T., Irwin, M. J., Lodieu, N., McMahon, R. G., Simpson, C., Smail, I., Mortlock, D., and Folger, M., “The UKIRT Infrared Deep Sky Survey (UKIDSS),” *MNRAS* **379**, 1599–1617 (Aug. 2007).
- [61] Hambly, N. C., Collins, R. S., Cross, N. J. G., Mann, R. G., Read, M. A., Sutorius, E. T. W., Bond, I., Bryant, J., Emerson, J. P., Lawrence, A., Rimoldini, L., Stewart, J. M., Williams, P. M., Adamson, A., Hirst, P., Dye, S., and Warren, S. J., “The WFCAM Science Archive,” *MNRAS* **384**, 637–662 (Feb. 2008).
- [62] Gaudet, S., Dowler, P., Goliath, S., and Redman, R., “The JCMT Science Archive and the Virtual Observatory,” in [*Astronomical Data Analysis Software and Systems XVII*], R. W. Argyle, P. S. Bunclark, & J. R. Lewis, ed., *Astronomical Society of the Pacific Conference Series* **394**, 135–+ (Aug. 2008).

- [63] Economou, F., Jenness, T., Chrysostomou, A., Cavanagh, B., Redman, R., and Berry, D. S., “The JCMT Legacy Survey: The challenges of the JCMT Science Archive,” in [*Astronomical Data Analysis Software and Systems XVII*], R. W. Argyle, P. S. Bunclark, & J. R. Lewis, ed., *Astronomical Society of the Pacific Conference Series* **394**, 450–+ (Aug. 2008).
- [64] Economou, F., Jenness, T., Cavanagh, B., Goliath, S., Redman, R., Jenkins, D., Berry, D., Chrysostomou, A., Chapel, B., Dowler, P., Gaudet, S., Ouelette, J., and Schade, D., “The jsa and the grid: How “infinite” computing power enables a new archive model for pi-led observatories,” in [*Astronomical Data Analysis Software and Systems XX*], Evans, I. N., Accomazzi, A., Mink, D. J., and Rots, A. H., eds., *Astronomical Society of the Pacific Conference Series* **TBD**, in press, ASP, San Francisco (2011).
- [65] Berry, D. S., “NDG – routines for accessing groups of NDFs,” tech. rep., Starlink User Note 2, Starlink Project (2011).
- [66] Jenness, T., Berry, D., Chapin, E., Economou, F., Gibb, A., and Scott, D., “SCUBA-2 Data Processing,” *ArXiv e-prints* **arXiv:1011.5876** (Nov. 2010).
- [67] Economou, F., Jenness, T., and Rees, N. P., “Sharing code and support between heterogeneous telescopes: the UKIRT and JCMT joint software projects,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], P. J. Quinn, ed., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **4844**, 321–330 (Dec. 2002).