*Astronomical Data Analysis Software and Systems XV*
*ASP Conference Series, Vol. XXX, 2005*
*C. Gabriel, C. Arviset, D. Ponz and E. Solano, eds.*

# The WFCAM Science Archive

Ross Collins, Nicholas Cross, Nigel Hambly, Robert G. Mann, Mike Read, Eckhard Sutorius, and Peredur Williams

*Wide Field Astronomy Unit, Institute for Astronomy, University of Edinburgh, UK*

Ian Bond

*Institute of Information and Mathematical Sciences, Massey University, New Zealand*

**Abstract.** The WFCAM Science Archive (WSA) will contain the complete set of data created by the WFCAM camera on UKIRT, including the UKIRT Infrared Deep Sky Surveys (UKIDSS). Developed and hosted by the Wide Field Astronomy Unit (WFAU) within the University of Edinburgh's Institute for Astronomy, the WSA is designed as an SQL database with user access through a public web interface. The UKIDSS surveys have now begun and the first data release will become available, after an initial phase of science verification, to all members of the ESO community. Observations will continue over the next seven years, typically producing 100 GBytes of data per night, resulting in an archive of the order of tens of terabytes in size. Consequently, scalability is one of the most important aspects in the design of the database. Here we provide an overview of our hardware and software design to store and access such vast quantities of data efficiently. Development of the archive continues as we prepare to host data from the VISTA surveys and enhance the end-user's experience, including incorporation of Virtual Observatory functionality delivered through AstroGrid.

## 1. Overview

The WSA will eventually contain the complete catalogue of source detections and image products from UKIRT observations with WFCAM. Primarily this will consist of the UKIDSS (Lawrence et al. 2006), which are divided into five survey projects, simply named to describe their purpose: the Large Area Survey (LAS), the Galactic Plane Survey (GPS), the Galactic Clusters Survey (GCS), the Deep Extragalactic Survey (DXS), and the Ultra Deep Survey (UDS). From these surveys we expect to see, on average, 100 GBytes of real-time pipeline-processed data generated per night. The data is transported by tape to the Cambridge Astronomical Survey Unit (CASU), where further science data processing is performed (see Lewis et al., these proceedings). The FITS files containing image and catalogue data generated by Cambridge are transferred to Edinburgh's Wide

Field Astronomy Unit (WFAU) over the JANET network at a typical optimum transfer speed of approximately 12 MBytes/s.

## 2. Hardware Choices

The archive is hosted by 32-bit PCs running Linux. To distribute the computing load we employ a dedicated server for user access separate from the server used to perform the database curation tasks.

The pixel (image) data is stored separately from the catalogue data. For the pixel data we employ a combination of 250 and 400 GByte SATA discs in a RAID5 array configuration. This offers us a fast data read rate of 80 MBytes/s together with disc redundancy, though with a penalty that the writing speed is four times slower. However, this is acceptable as writing should only be performed once per image, whereas reading will occur many times by users accessing the database. Back-ups of the data are not, at present, made locally, as the pixel data do not differ from the copies held at CASU.

The database of catalogue data is stored on a 150 Gbyte Ultra320 SCSI discs in a RAID50 array configuration. This offers data access at a rate of 200 Mbyte/s with just a single channel, and 400 Mbyte/s in full RAID50 mode. Back-ups of the catalogue data are stored on Ultrium-II LTO tapes.

At present we have a total storage capacity of 16 TBytes for pixel data and 11 TBytes for catalogue data. This is sufficient to host the complete set of science verification data. However, the storage capacity will be increased to host the full database.

## 3. Database Management

Learning from the experiences of the SkyServer archive for the Sloan Digital Sky Survey (SDSS), we chose to adopt Microsoft SQL Server as our database management system (DBMS). It has been proven to work well with large astronomical databases and Microsoft are keen to support the development of such large-scale databases.

Every table stored in the database is described by a schema model that defines the table fields, e.g. labels, data types, and natural language descriptions. Hence, there is a schema for every survey, and every product of each survey, such as source lists, neighbouring sources, and associated sources in external databases. The schema can be viewed on the WSA website (see section 5.).

## 4. Software Design

Software has been developed in-house for this project to manage the curation tasks for the archive. Each task is controlled by a Python script that handles such tasks as retrieving the data from CASU and its subsequent ingest into our database. A complete list and description of each of the archive's curation tasks is documented on-line (`http://www.roe.ac.uk/~nch/wfcam`).

The more data-intensive tasks, such as producing JPEG image products, image mosaics etc. are performed by software written in C, developed by both

CASU and WFAU, that utilises the CFITSIO library. We have also developed object-oriented C++ code to generate the catalogue of WSA sources, which is a product of merging the CASU-supplied catalogue of detections from each passband. Presently, only the Python code (and the Java servlets used by the web-based user interface, see section 5.) can directly connect to our database. Therefore, all of our C++ code works with our schema designs to manipulate binary flat-file outgest tables from the database.

### 4.1. Database Connectivity with Python

Interaction with the DBMS by our software is made possible by ODBC, the industry standard API for DBMS connections. The obvious choice for an ODBC implementation to be used by our Python code was mxODBC, for which we were provided with a free 'non-commercial use' license. mxODBC is the most reliable and highest quality ODBC implementation, having been designed in collaboration with the developers of the Python language. As the DBMS is hosted on a Windows PC, and the Python software is served by a Linux PC, the ODBC connection consists of many layers. On the Linux PC the Python code calls mxODBC, which makes use of unixODBC, which in turn uses FreeTDS to connect to the ODBC API on the Microsoft SQL Server PC.

### 4.2. Source Neighbour Tables

For every source in the WSA catalogue we maintain a list of neighbouring sources from both the WSA catalogue and other external catalogues that we host. To produce these *neighbour tables* we utilise a CSIRO-developed catalogue matching code (Abel et al. 2004). The algorithm used by this code, scales as $O(N \log M)$, where $N$ and $M$ are the respective sizes (i.e. number of rows) of the two catalogues, making it the most efficient method for comparing two extremely large catalogues. Using this code, our 2.4 Ghz Pentium 4 PC found source neighbours within a search radius of $10''$ in the $10^9$ row SSA database in just 5 hours. Instead of directly reading from and writing to the database we operate on binary files containing the source data extracted from the database. This outgest process takes 2 hours, including a sort by declination required by the source matching code, and the subsequent ingest of $1.2 \times 10^9$ neighbouring source ID pairs with distances takes 1.5 hours. This C++ code is well documented and follows an object-oriented design methodology that makes it simple to adapt to any system.

## 5. User Access

Users can perform queries on the archive database through web query forms on our website (`http://surveys.roe.ac.uk/wsa/pre/`). Here, a simple form (see figure 1), containing options for common query parameters such as RA and declination ranges, can be used to query the database as well as a form offering the user full control for free-form SQL queries. A cookbook of common SQL queries is provided on-line to help users form their own queries, and a schema browser is also provided listing each of the database field names and formats. Joint queries with other locally hosted databases can be performed, and we also offer the ability to cross-match the WSA catalogue with user-supplied source

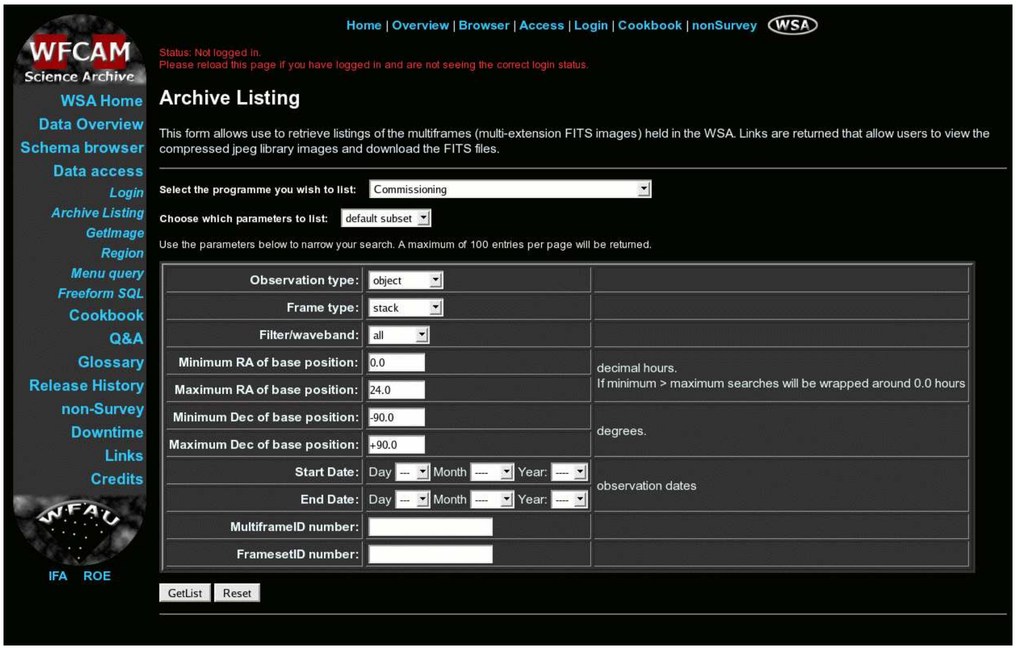


Figure 1. The WFCAM Science Archive query form located at http://surveys.roe.ac.uk/wsa/pre/

lists. These web form generated queries are executed using Java servlets on a Tomcat server that connects to our DBMS via JDBC. The extracted data can then be returned in a variety of formats, including FITS and VOTable. Query products can now also be saved to the AstroGrid's MySpace; our initial integration of the archive with the virtual observatory.

## References

Abel, D. J., Devereux, D., Power, R. A., & Lamb, P. R. 2004, An O(N log M ) algorithm for catalogue matching. Technical Report TR- 04/1846, CSIRO ICT Centre, Canberra, Australia.

Lawrence, A., et al., 2006, in preparation

Lewis, J. R., Irwin, M. J., Bunclark, P., Hodgkin, S. T., Evans, D. W., and McMahon, R. G., these proceedings