

The NOAO Mosaic Data Handling System

D. Tody and F. G. Valdes

NOAO*, P.O. Box 26732, Tucson, AZ 85726, USA

ABSTRACT

The NOAO Mosaic CCD Camera consists of 8 CCDs producing an 8K x 8K format. The *Mosaic Data Handling System* (MDHS) receives the data in real time as it is read out of the detector. A *data feed client* (DFC) receives data packets from the detector system and broadcasts them on a message bus. A *data capture agent* (DCA) receives the data and writes it to a *distributed shared image* (DSIM). The images are displayed as they are read out by a *real-time display* (RTD) and are recorded to disk as *multi-extension format* (MEF) FITS files by the DCA. The DCA triggers a *data reduction agent* (DRA) to do standard pipeline processing and archiving with a graphical user interface (GUI) for user interaction. The MDHS provides a suite of IRAF data reduction tools for the user and the DRA which are tailored for the MEF mosaic data. The tools provide capabilities for quick-look analysis of the data, basic CCD calibration, mosaic reconstruction, and combining of dithered observations into a fully populated (gaps removed) image.

Keywords: data acquisition, data processing, mosaic camera, CCD, IRAF

1. INTRODUCTION

The *NOAO Mosaic Data Handling System* (MDHS) takes data from a mosaic of CCDs and decodes, records, archives, displays, and processes the data. The processing tools are available as an IRAF package for observers to re-reduce the mosaic data at their home institutions using IRAF.

The MDHS is being developed to handle data from the NOAO Mosaic CCD Camera^{1,2}. This camera consists of 8 2048x4096 CCDs producing an 8K x 8K format with minimal gaps (see figure 2). The data produced by the camera consists of 8 or 16 subimages, depending on whether one or two amplifiers are used per CCD. However, the MDHS is not tied to the NOAO Mosaic CCD Camera. It is designed to be portable and applicable to any mosaic of CCDs including single frame CCDs. The entire MDHS may be used by connecting a detector system to a “message bus” (§3) using a well-defined protocol. The data reduction and analysis tools (§9) may also be used separately from the MDHS if the data is recorded in the FITS disk (§4) format used by the software.

A large mosaic such as that produced by the NOAO Mosaic CCD Camera presents many challenges for a data handling system. The use of multiple CCDs requires that data be read out simultaneously from all CCDs, hence the raw data is interleaved as it arrives from the detector and must be “unscrambled” before being written to disk or displayed. The CCDs have different bias and gain characteristics requiring calibration before they can be viewed together on a display. The CCDs are not perfectly aligned and have gaps between the CCDs requiring interpolation, image combination, and dithering. The large field and the use of optical correctors means that field distortions are significant. Combined with the misalignment of the CCDs, this complicates coordinate determination, astrometry, and registration of dithered exposures. Finally, the data set is very large. A powerful computer system and efficient software is required to be able to handle such large formats. Even viewing the data is difficult since an exposure is a composite of a number of smaller images and, at 8Kx8K or 64 megapixels for the NOAO camera, the area is about 50 times that of the typical workstation screen.

The development of the MDHS began in 1996.³ A preliminary but working version of the system has been in use at NOAO with the NOAO Mosaic Camera since June 1997. The current MDHS is able to produce images of high scientific and aesthetic quality despite the engineering grade CCDs used in the present NOAO camera. As of late 1997 work on the full system is still underway. This paper presents the MDHS design and details of the important components, both currently implemented and under development.

E-mail: fvaldes@noao.edu, dtody@noao.edu

*National Optical Astronomy Observatories, operated by the the Association of Universities for Research in Astronomy, Inc. (AURA) under cooperative agreement with the National Science Foundation.

¹<http://www.noao.edu/kpno/mosaic/mosaic.html>

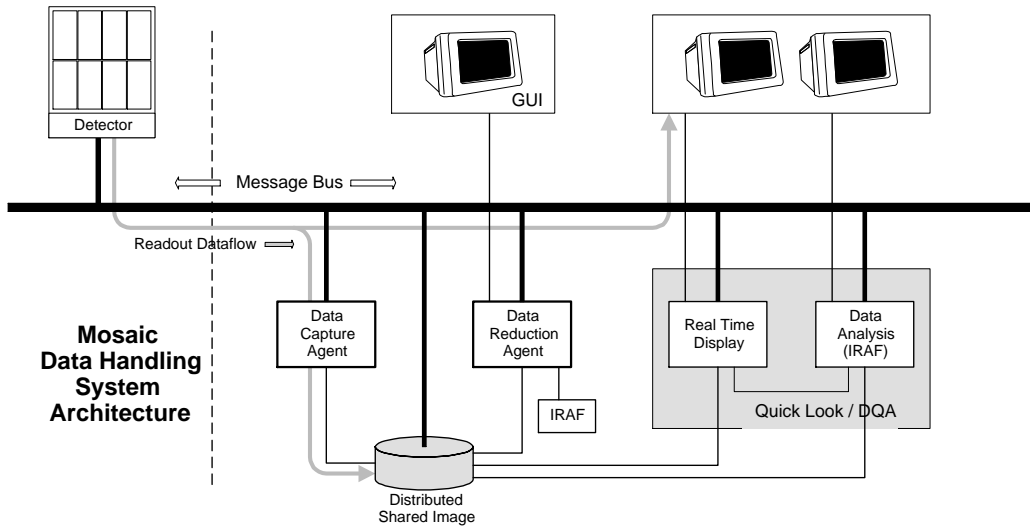


Figure 1. The Mosaic Data Handling System Architecture

2. MDHS ARCHITECTURE

Figure 1 illustrates the major components and data flow of the NOAO Mosaic Data Handling System. A *message bus* connects the various elements of the MDHS. The detector system sends packets of events, descriptive information, and (interleaved) pixel data on the message bus to subscribing clients. For example, the NOAO Mosaic detector system writes data to a set of temporary disk files during a readout. A *data feed client* (DFC) maps the temporary files into memory and reads the data as it is written, translating data packets and writing them to the message bus. The disk files serve as a large FIFO buffer and provide a backup mechanism permitting retransmission should anything go wrong.

The *data capture agent* (DCA) captures the pixel data and observation information packets and creates a *distributed shared image* (DSIM) and a Mosaic *multi-extension format* (MEF) FITS observation file on disk. At the same time a *real-time display* (RTD) accesses the DSIM over the message bus and displays the mosaic exposure during frame readout. Quick-look is provided by the RTD and by IRAF, which can interact with the RTD during and after frame readout. The *data reduction agent* (DRA) directs the post-processing of each observation file, applying standard calibrations and writing the data to tape and to the data archive. The box labeled IRAF constitutes a suite of general and mosaic specific data reduction and analysis tools in IRAF that are used by both the DRA and the user.

3. MESSAGE BUS

The heart of the Mosaic data handling system is the *message bus*⁴ which connects all data system components. The message bus provides flexible and efficient facilities for components to communicate with each other. The message bus (which is a software facility) supports both distributed and parallel computing, connecting multiple host computers or multiple processors on the same host. For example, the MDHS at NOAO currently uses two computers, one for the detector system and one for rest of the MDHS, with a fallback to one computer in the event the second computer fails.

The message bus provides two methods for components to communicate with each other. *Producer/consumer events* allow components to listen for (consume) asynchronous event messages produced and broadcast by other components. *Requests* allow synchronous or asynchronous remote procedure calls (method invocations) to be directed to services or data objects elsewhere on the message bus. Discovery techniques can be used to determine what services are available and to query their methods. Host computers and components can dynamically connect or disconnect from the bus. The bus can automatically start services upon request; or services and other components can be started by external means, connecting to the message bus during startup.

The MDHS uses a custom message bus API which is layered upon some lower level messaging system. At present we are using the Parallel Virtual Machine (PVM) facility; in the future we might use other facilities such as CORBA. The use of a custom API provides isolation from the underlying messaging facility and aids development of a standard framework and set of services for integrating a set of applications.

3.1. DISTRIBUTED SHARED OBJECTS AND DISTRIBUTED SHARED IMAGES

An important class of message bus component is the *distributed shared object* (DSO). DSOs allow data objects to be concurrently accessed by multiple clients. The DSO provides methods for accessing and manipulating the data object, and locking facilities to ensure data consistency. DSOs are distributed, meaning that clients can be on any host or processor connected to the message bus. In the case of the MDHS, the principal DSO is the *distributed shared image* (DSIM) which is used for data capture, to drive the real-time display, and for quick-look interaction from within IRAF. The distributed shared image uses shared memory for efficient concurrent access to the pixel data, and messaging to inform clients of changes to the image.

The current version of the DCA does not implement the full DSIM object. Instead it implements a mapped image file. The observation being captured appears as a valid MEF file (§4) immediately after the configuration information is received from the DFC. This allows applications to examine the file while the readout is in progress. An example of this is the interim display routine that loads a display server frame buffer as the pixel data is recorded giving a simple real-time display capability.

4. MOSAIC DATA STRUCTURE

The mosaic data recorded by the DCA and used by the mosaic data processing tools in IRAF is stored as one *multi-extension format* (MEF) FITS file per exposure. The FITS file contains a *primary header* with no associated data and a number of data *extensions*. The primary header is used to describe the contents of the file and contains *global keyword* information applicable to all the extensions. The extensions include the image data from each amplifier, pixel masks, uncertainty arrays, exposure maps, auxiliary tables, etc. The image data extensions are always present while other information is added at various stages during the reductions.

The CCD image data consists of separate FITS Image Extensions for each amplifier. Each extension has an *extension name* which can be used to refer to the image by IRAF software through the IRAF *FITS Image Kernel*.⁵ For instance to refer to the image data for the third amplifier a reference such as `obs123[im3]` would be used. Pixel masks, uncertainty arrays, and other array type extensions are also accessed by software through the FITS kernel and the extension name syntax.

Information about the exposure is recorded in the global and extension FITS headers. Valdes⁶ describes a methodology for organizing observational information logically and mapping it to FITS header keywords. The NOAO Mosaic data uses a FITS keyword dictionary to formally define and document the header keywords*. See reference 6 for more details on the mosaic data format.

5. DATA CAPTURE AGENT

The *data capture agent* (DCA) is a general purpose network (message bus) data service. It receives messages of various types and, through an event loop, disposes of each message to an appropriate event handler. Examples of message types are control commands, server parameters, readout status, data format configuration, header, and pixel data. The DCA can handle multiple simultaneous readouts from different clients which can be on any host computer connected to the message bus.

Incoming header data is buffered internally within the DCA. Incoming pixel data blocks are unscrambled and written directly to the output MEF image using the DSIM facility. When the readout is finished a table driven keyword translation module (implemented as a configurable TCL script) transforms the input device dependent detector keywords and adds information as necessary to conform to the data format required by the rest of the MDHS. Ultimately a new MEF observation file is written to disk and passed off to the DRA for post-processing.

*<http://iraf.noao.edu/mosaic/imagedef/fitsdic.html>

5.1. DCA User Interface for the NOAO Mosaic Camera

The DCA is started by issuing a command at the host level, or invoked as a service by the message bus. This can occur during login to an observing account, using a window manager menu, or by interactively issuing a user command. The command invocation can include DCA server parameters which can also be set or reset after invocation via messages.

The DCA automatically connects to the message bus when it is started. The runtime operation of the DCA is monitored and controlled by clients via the message bus. In the Mosaic system there are two clients; the data feed client (DFC) and the DCA console client (DCAGUI) (not shown in figure 1). In general any number of clients can access the DCA. Multiple data feed clients or GUIs can be active simultaneously.

The DCA console client has a graphical user interface based on the IRAF Widget Server.⁷ It can send control and parameter messages to the DCA and act on messages broadcast by the DCA, e.g. to display the readout status. Currently the DCAGUI executes an auto-display command for real time display during readout (see Valdes, 1997) and a post-processing command for logging, archiving or other operations once the readout ends.

5.2. The Readout Sequence

A description of what happens during a readout illustrates the functioning of the DCA. An observation readout leading to a MEF observation file consists of a sequence of messages. A readout sequence is initiated by the DFC with a message that includes a unique sequence number. Once a readout has started, each subsequent message associated with that readout must be tagged with the sequence number for that readout (multiple readouts can be simultaneously in progress). Each readout sequence has a separate context identified by its sequence number.

When a readout sequence is initiated the DCA creates a named DSIM. The DFC passes information about the size and structure (such as the number of image extensions) allowing the DCA to configure the DSIM. At this point the DFC can begin sending pixel and header messages. At image creation time the DCA broadcasts a message that the DSIM has been created so that clients like the DCAGUI can access it if desired. Currently the DCAGUI uses this to initiate an interim real-time display.

When the readout is completed the DCA executes a *keyword translation module* (KTM), an externally supplied, interpreted TCL script which converts detector specific information into standard header keywords. After the KTM finishes the DCA and DSO format the output keywords, write them to the image headers, and generation of the new image DSO is complete. The DCA broadcasts a message when the DSO is complete which can be used to trigger post-processing of the data.

5.2.1. Pixel Messages

The pixel data messages contain blocks of raw, unprocessed detector pixel data organized into one or more *streams*, one for each CCD amplifier. Each stream directs pixels to a region of an output image extension. This structure allows the data block to simultaneously contain data for several different regions and the data can be arbitrarily interleaved, encoded, flipped, or aliased. Each block of data is processed asynchronously but the DFC can send a synchronization request periodically to check the output status.

5.2.2. Keyword Messages

The DCA receives header information via the message bus. This information consists of blocks of keywords organized into named, detector specific *keyword groups*. The keywords are stored in keyword databases (an internal random access data structure), one per keyword group. The set of keyword group names is arbitrary. Some examples for the NOAO Mosaic are ICS for instrument control system, TELESCOPE for telescope, and ACEB n for Arcon controller information for controller n .

The Arcon controller system for the NOAO Mosaic consists of a set of Arcon controllers each of which can readout one or more amplifiers. The current system has four controllers each reading out two CCDs using one amplifier per CCD. Thus there can be controller information for the whole system, for each controller, and for each amplifier readout.

A keyword database library handles creation and maintenance of the keyword database. Note that keyword information does not necessarily have to come only from the DFC, the current mode of operation for the NOAO Mosaic. Other schemes are possible.

5.3. Keyword Translation Module

The *keyword translation module* (KTM) is a TCL script called by the DCA at the end of a readout. The purpose of the KTM is to create the keywords for the global and extension headers. The KTM is passed a list of input keyword database descriptors and it returns a list of output keyword database descriptors, one for each output header. The TCL interpreter in the DCA provides custom keyword database extensions which allow manipulation (creating, accessing, searching, etc.) of keyword databases by the TCL script. When the KTM finishes the DCA, via the DSO, writes the keywords in the returned output keyword databases to the output MEF FITS file.

The KTM performs a variety of transformations on the input keywords. A keyword can be copied verbatim if no change is desired. The keyword name or comment may be changed without changing the value. New keywords can be added and default values may be supplied for missing keywords. The KTM can compute new keywords and keyword values from input keywords or other data. Identical keywords in each extension may be merged into a single keyword in the global header. The KTM can detect incorrect or missing keywords and print warnings or errors.

Two examples from the keyword translation module for the NOAO CCD Mosaic follow.

1. The data acquisition system provides the keywords DATE-OBS and UTSHUT giving the UT observation date in the old FITS date format (dd/mm/yy) and the UT of the shutter opening. The KTM converts these to TIME-OBS, MJD-OBS, OBSID, and the new Y2K-compliant FITS date format.
2. The KTM determines on which telescope the Mosaic Camera is being used and writes the celestial coordinate system scale, orientation, and distortion information previously derived from astrometry calibrations for those telescopes. The coordinate reference point keywords, CRVAL1 and CRVAL2, are computed from the telescope right ascension and declination keywords. See §9.2 for more discussion.

6. REAL-TIME DISPLAY AND QUICK-LOOK

The primary function of the *real-time display* (RTD) is to display the Mosaic data in real-time during readout; pixels appear in the display as soon as readout begins. As noted earlier the DCA does not just write to a disk image, it writes to a DSIM. At the same time that the DCA is writing to the DSIM, the RTD is reading from it and displaying the incoming data. The DCA receives an incoming write-pixel request on the message bus from the DFC (or some other client), obtains locks on a set of regions in the output image (e.g. 16 regions for 8 CCDs with 2 amps each), copies the input data to the output regions, and then frees the regions. This causes the DSIM to send messages to all clients, such as the RTD, which want to be informed of changes to the image. The RTD then performs any on-the-fly calibration or other processing and updates the displayed image.

To the user the RTD is an image browser displaying to one or more workstation screens; two in the case of the NOAO Mosaic. One screen shows the full mosaic de-zoomed at 50-to-1. The second screen shows a zoomed up region of the mosaic. Multiple zoom windows can be active on multiple screens.

The RTD is not just a real-time display, it is a fully functional image viewer with extensive built-in functions for quick-look image analysis. Additional functionality (possibly quite extensive) can be added via a dynamic “plug-in” facility, which allows users or projects to easily customize the display or tailor it for existing data systems. Finally, extensive image analysis is available via IRAF or any other external image analysis system which interfaces with the RTD and DSIM. IRAF sees the DSIM as if it were a conventional disk image, allowing any IRAF task to be used. This allows tight integration of IRAF quick-look or analysis tasks with the RTD. It is even possible to use an IRAF task to operate upon the incoming image during readout, before readout has completed.

The RTD described above is under development. Meanwhile an interim display tool is available to display MEF image data as a single pixel matrix in a standard image display viewer such as Ximtool. It includes real-time capabilities to display the MEF data being written by the DCA while a readout is in progress. The DCAGUI calls this tool when the DCA begins writing the MEF file. Related tools allow users to interact with the displayed mosaic exposure (even during readout) to evaluate focus and to do quick-look analysis such as PSF fitting, statistics, graphics, and celestial coordinate measurements.

7. DATA REDUCTION AGENT

The MDHS includes a full pipeline data reduction capability, plus facilities for taping, archiving, viewing and managing the data set. All data reduction is performed by IRAF tasks under the direction of the *data reduction agent* (DRA). The DRA is driven by a device dependent script, hence is user configurable and easily adapted for new instrument configurations.

The DRA automatically or by user command operates on the observation files. It is a high level task with a sophisticated graphical user interface (GUI). It communicates with other processes to perform pipeline calibrations, reductions, data quality assessment, archiving and possibly other functions. These processes are IRAF tasks which access the observation files through the IRAF *FITS Image Kernel*.⁵

The DRA is a continuously running event-driven process. The events which trigger the above functions are when the DCA finishes writing an observation to disk and when the user initiates an action via the graphical user interface. The first case provides automatic processing and archiving. The second case allows the user to perform manual calibrations or initiate recalibrations of the automatic processing. Reprocessing would be done when additional or improved calibration data become available. For example, the automatic processing can proceed using calibration data from the start of the night, and recalibration can be done after additional calibrations at the end of the night are obtained.

The pipeline calibration, reduction, and quality assessment are defined by *recipes* selected from a list of recipes. A recipe is basically a *macro* or *script* that is executed on a specified disk file or set of disk files. Pipeline calibration consists of the standard CCD calibration operations, combining sequences of calibration exposures with scaling and bad pixel rejection and other data reduction operations (§9).

8. ARCHIVE AGENT

The archive agent is responsible for archiving mosaic data. The NOAO MDHS currently uses the NOAO *Save The Bits*⁸ system which archives the FITS header and enters the MEF observation file produced by the DCA into a queue which eventually saves it to tape. Since the Mosaic data format is already an archival FITS format it requires only small additions to the data header before spooling the disk file directly to tape. The archive agent is currently triggered by the DCAGUI when the DCA signals completion of the disk file.

9. DATA REDUCTION

The MDHS includes complete data reduction capabilities. The data reduction tools are implemented as an IRAF package called `mscred`.^{9,10} These tools are accessed either through the DRA or the IRAF command language (CL). They operate on the raw Mosaic MEF FITS files producing new or modified MEF files or mosaiced single images.

The reduction of CCD mosaic camera data can be divided into four stages. The first is the basic calibration of the individual CCDs. This stage is very similar to reducing data from single CCD exposures except that the calibration operations are repeated for all the CCDs in the mosaic. The only significant difference is that any scaling of an exposure, such as normalizing a flat field calibration, must be done uniformly over all the CCDs.

The second stage is calibrating the *world coordinate systems* (WCS) of each exposure. The WCS in this case means the mapping between the pixel coordinates in each mosaic piece and celestial coordinates on the sky. In effect this step registers the exposures to a common sky coordinate system and corrects for variations due to instrumental flexure and atmospheric refraction.

The third stage is resampling the individual mosaic exposures into geometrically correct images of the sky. Since creating a single image from a single mosaic exposure is of marginal value over using the unresampled data, the fourth stage of combining multiple, offset (“dithered”) exposures follows. The combining requires calibrating flux variations between the exposures due to transparency and sky brightness changes, and registering and combining the overlapping regions with gaps and bad data excluded. The final result is a deep single image with gaps and bad pixels removed (see figure 2).

9.1. Basic CCD Calibration

Basic CCD instrumental calibrations consist of correcting each pixel for electronic bias levels, zero level exposure patterns, dark counts, and response variations. This is done on the individual CCD images in the MEF image file. The response calibration involves not only sensitivity variations between pixels in the same CCD but also response differences between the CCDs and illumination changes from a variable pixel scale.

Electronic bias levels are removed using overscan data in the pixel readout from the detector for all CCD lines. The overscan values in the raw data are also removed (trimmed) after use. The zero level exposure patterns are removed by combining a number of shutter closed, zero exposure time exposures and subtracting this from dark count, flat field, and object exposures. Dark count corrections are performed by combining a number of shutter closed exposures with comparable exposure times to the data to be calibrated, scaling if the exposure times are not the same, and subtracting from the flat field and object exposures.

Response calibrations consist of three steps. One is using flat field exposures combined into a single calibration and divided into the object exposures. Since the flat field exposures may not illuminate the detector in the same way as the sky exposures, the second step uses unregistered object exposures. When these are combined without registration and using rejection techniques to excluded contributions from objects in the fields, a sky flat field is produced. The sky flat is divided into the object exposures.

Sky flat field calibrations are quite important with mosaics in order to get the different CCDs to match when making single images. For the common observational method of taking multiple dithered exposures, sky flat calibration is critical but, provided the fields do not contain very large objects, the dithered data itself, over the course of the observing, can be used for making a sky flat field.

The flat field calibrations have the effect of making the counts in the sky pixels uniform. However, if the projected size of the pixels on the sky varies due to the optics, the true sky counts will vary in proportion to the pixel area on the sky. This effect has often been ignored or forgotten with smaller single CCD formats. But the large field of view provided by a mosaic and the optics required to provide it can lead to a significant variation. This effect is quite significant with the NOAO 4-meter telescopes and is also present in the KPNO 0.9-meter to a smaller degree.

This photometric effect is calibrated by dividing each pixel by its sky area relative to some fiducial pixel. As discussed more fully below (§9.2), the MDHS data includes a coordinate system mapping between the pixels and celestial coordinates. This mapping is used to compute the area of the sky covered by each pixel.

The pixel area correction may be performed directly on the original pixels of each mosaic piece or during the geometric resampling when a single image is produced from the mosaic pieces. One would typically only apply the correction to the individual mosaic pieces if the individual CCD images are used for photometric study.

Besides the above basic calibrations there are some additional operations which may be performed during the calibrations. Bad pixels identified by a *bad pixel mask* may be cosmetically removed by interpolation from nearby good pixels. Saturated pixels may be detected and flagged in the mask. Pixel uncertainty information may be derived from the detector characteristics and associated with the data. The various calibration operations propagate these uncertainties.

9.2. COORDINATE SYSTEM CALIBRATION

The geometric correction and combining of dithered exposures depends on having an accurate *world coordinate system* (WCS) for each exposure. A WCS is defined in the image header for each mosaic piece. The WCS consists of three parts: 1) a reference point matching a point in the exposure to a point on the sky, 2) a “plate solution” mapping pixel offsets to arc second offsets from the reference point and 3) a “projection function” applied to the arc second offsets to produce celestial coordinates.

The important point about this WCS definition is that the the plate solution is independent of position on the sky (ignoring refraction and flexure effects) and so can be the same for all exposures. Only the sky coordinate of the reference point needs to be set for each exposure. The plate solution is determined from astrometry fields taken with the camera. The `mscred` package provides tools for deriving the plate solution for each piece of the mosaic.

For the NOAO Mosaic Camera and probably for most mosaic cameras a plate solution can be determined in advance by the instrument support personnel. The MDHS then allows for this default WCS to be added to the data as it is acquired. This can be done by the detector system supplying the information on the message bus or by the

DCA adding it through the KTM. For the NOAO Mosaic Camera the latter is done with the reference point on the sky set from the telescope coordinates provided by the detector system. The accuracy of the WCS for the NOAO Mosaic Camera has been studied by Davis.¹¹

The coordinate calibration steps consists of setting the coordinate reference point more accurately than the telescope system may provide, and determining a low order (linear) correction of the plate solution for origin shift, rotation, and scale change keeping the higher order distortion terms fixed. The coordinate system reference point can be set using a star in the exposure with known coordinates for an absolute calibration or, more commonly, all the exposures in a dither pattern can be adjusted to the coordinate of one star in one exposure as given by the default WCS. A display oriented tool is provided to allow pointing at objects in the exposure and measuring the coordinate from the WCS or specifying the coordinate that it should have for a zero point calibration.

The linear correction, which is needed to correct for atmospheric refraction during a series of dithered exposures, is performed using a list of celestial coordinates for objects in the field. The coordinates might be from a catalog or measured from one exposure of a dither sequence to which the other exposures are matched. Since the WCS of the exposures are fairly accurate (especially if a first reference point adjustment has been made based on one object) the software determines where in each exposure the objects in the coordinate list will appear and finds an accurate pixel position for the object. With the set of pixel positions and celestial coordinates a global rotation, scale change, and origin shift can be determined and used to update the WCS for each mosaic piece.

9.3. MOSAICING

Up until this point the mosaic data from the individual amplifiers and CCDs have been kept separate in the MEF file. The observer may analyze the individual calibrated pieces of the exposure. This avoids resampling (interpolation) of the pixel data. However, the observer may wish to put the pieces of the mosaic into a single large image with the CCD pieces in their proper relation and with optical distortions removed.

This is done by defining a uniform raster of pixels on the sky. In other words, what an image of the sky would be with an ideal large detector of equal sized pixels in a simple raster grid using a distortionless telescope. If multiple exposures are to be combined the image grid is chosen to be uniform over all the exposures so that the pixels of the resampled images will be precisely registered where they overlap.

The resampling consists of mapping each pixel in the mosaiced image to the region of the exposure in one of the mosaic pieces and determining a pixel value from that region. A choice of pixel estimation methods is provided from simple linear interpolation, to polynomial interpolation, sinc interpolation, or the currently popular “drizzle” method. A bad pixel mask for the mosaiced image is created by identifying which pixels have contributions from bad pixels (given in the bad pixel masks for the exposure) in the original mosaic pieces.

As noted earlier (§9.1) the flat field calibration distorts the photometric information because it does not take into account variations in the projected size of the pixels on the sky. This can be corrected in the individual pieces if desired. But if this is not done then the resampling step can include the pixel size correction since it knows the sky area mapping between the original pixels and the mosaiced pixels.

9.4. COMBINING DITHERED OBSERVATIONS

A CCD mosaic exposure usually includes gaps between the CCDs where the sky is not observed. To fill in these gaps multiple offset exposures are taken in a “dither” pattern and the exposures are then combined. This process has the additional valuable advantages of allowing for removal of bad pixels in the CCDs and removal of cosmic rays.

The basic CCD calibrations (§9.1) insure that the pixels across the exposure have a uniform flux scale. However, this calibration does not account for changes in the flux scale between exposures in the dither pattern. These changes are due to variations in the sky brightness (a zero point change) and in the sky transparency (a linear scale change). In order to combine the dithered exposures they must be calibrated to the same flux scale. This is quite important because even small differences will result in visible artifacts in the final combined image.

To calibrate the flux scale a list of coordinates of common objects in the field of the dithered exposures is created. This is currently done by marking positions in one of the exposures using an image display. In the future automatic detection algorithms can be added to do this. Simple photometry is performed consisting of summing the pixel values in a box centered on each coordinate and in an surrounding square annulus. Note that it does not matter what is in




Figure 2. NOAO Mosaic CCD Camera pictures of M33 taken at the KPNO Mayall 4-meter. The left picture is a single 5 minute exposure showing the mosaic format of 8 CCDs each 2Kx4K for a full format of 8Kx8K. The right picture is a fully processed final image which is a combination of 5 dithered exposures. The processing was done using the MDHS data reduction software.

the box or in the outer annulus provided the same image of the sky is measured in each of the dithered exposures. All that is needed is a reasonable range of total counts and a good number of points.

The photometry is matched between the exposures to determine a set of relative offsets and linear scales. An algorithm is used that considers the relationships between all pairs of exposures, rejects deviant points, and combines these into a consistent set of scaling factors. The scale factors are recorded in the image headers for use during the combining step; i.e. the images themselves are not rescaled before the combining step.

The WCS calibration (§9.2) and mosaicing to a uniform sky grid (§??) produce images which are exactly registered over the full field of view. A combined image which contains all the dithered exposures is created. At each point in the combined image all exposures having pixels at that point are considered. The pixels are scaled using the scaling parameters previously determined to put them on a common flux scale. Any pixels marked as having contributions from bad pixels in the original CCD observations are excluded. The good pixels are used to identify cosmic rays using any of a number of clipping algorithms. Finally all the remaining pixels are combined as an average or median and the result recorded in the final image.

Figure 2 shows pictures from the NOAO Mosaic Camera. These pictures were processed as described above. The left picture has been through the mosaicing step so the gaps are geometrically correct. The right picture shows the success of combining dithered exposures to eliminate the gaps and cosmetic defects.

10. CONCLUSION

A version of the MDHS is in regular use at two telescopes on Kitt Peak with the NOAO CCD Mosaic Camera. This version includes the message bus, the data capture agent, and the data reduction facilities. Work is continuing on the message bus, the distributed shared object facility, the real-time display, and the data reduction agent. In addition the data reduction facilities will be enhanced to better support pixel masks, uncertainty information, astrometry, automatic detection of objects, and catalog access.

The components of the current MDHS are available for outside evaluation and use. Since users are routinely using the NOAO Mosaic CCD Camera the IRAF-based mosaic reduction software `mscred` is available to them as a standard IRAF external package for IRAF version 2.11.1 and later.

The primary function of the MDHS is to process data from the NOAO Mosaic CCD Camera. The significance of the project is much greater however. The MDHS itself is applicable to any type of data and when completed will be used for general data acquisition within NOAO and at other observatories. The message bus, DSO, and plug-in image display (RTD) technology used by the MDHS is being developed as a more general facility for use in the IRAF system and by other projects.

This work is supported in part by grants from the NASA Astrophysics Data Program and from the NASA Applied Information Systems Research program.

REFERENCES

1. T. Boroson, R. Reed, W. Wong, and M. Lesser, "Development of a 8192x8192 CCD mosaic imager," in *Instrumentation in Astronomy VIII*, D. L. Crawford, ed., *Proc. SPIE* **2198**, pp. 877–885, 1994.
2. R. Reed, G. Muller, T. Armandroff, T. Boroson, and G. Jacoby, "What is better than an 8192x8192 CCD mosaic imager: two mosaic imagers, one for KPNO and one for CTIO," in *Optical Astronomy Instrumentation, Proc. SPIE* **3355**, 1998.
3. D. Tody, "The data handling system for the NOAO Mosaic," in *Astronomical Data Analysis Software and Systems VI*, G. Hunt, ed., *ASP Conf. Series* **125**, pp. 451–464, ASP, San Francisco, 1997.
4. D. Tody, "Open IRAF message bus and distributed object technology," in *Astronomical Data Analysis Software and Systems VII*, R. Albrecht, ed., *ASP Conf. Series*, ASP, San Francisco, 1998.
5. N. Zarate and P. Greenfield, "A FITS image extension kernel for IRAF," in *Astronomical Data Analysis Software and Systems V*, G. H. Jacoby, ed., *ASP Conf. Series* **101**, pp. 331–334, ASP, San Francisco, 1996.
6. F. Valdes, "Data format for the NOAO Mosaic," in *Astronomical Data Analysis Software and Systems VI*, G. Hunt, ed., *ASP Conf. Series* **125**, pp. 459–462, ASP, San Francisco, 1997.
7. D. Tody, "The IRAF Widget Server," in *Astronomical Data Analysis Software and Systems IV*, R. A. Shaw, ed., *ASP Conf. Series* **77**, pp. 89–98, ASP, San Francisco, 1997.
8. R. Seaman, "NOAO/IRAF's save the bits, a pragmatic data archive," in *Astronomical Data Analysis Software and Systems III*, D. R. Crabtree, ed., *ASP Conf. Series* **61**, pp. 119–122, ASP, San Francisco, 1994.
9. F. Valdes, "IRAF data reduction software for the NOAO Mosaic," in *Astronomical Data Analysis Software and Systems VI*, G. Hunt, ed., *ASP Conf. Series* **125**, pp. 455–458, ASP, San Francisco, 1997.
10. F. G. Valdes, "The IRAF mosaic data reduction package," in *Astronomical Data Analysis Software and Systems VII*, R. Albrecht, ed., *ASP Conf. Series*, ASP, San Francisco, 1998.
11. L. E. Davis, "Astrometry with the NOAO Mosaic," in *Astronomical Data Analysis Software and Systems VII*, R. Albrecht, ed., *ASP Conf. Series*, ASP, San Francisco, 1998.