The CCD imaging systems for DEIMOS

Chris Wright, Robert Kibrick, Barry Alcott, D. K. Gilmore, Terry Pfister, Dave Cowley
UCO/Lick Observatory, University of California, Santa Cruz, CA 95064 USA

ABSTRACT

The DEep Imaging Multi-Object Spectrograph (DEIMOS) images with an 8K x 8K science mosaic composed of eight 2K x 4K MIT/Lincoln Lab (MIT/LL) CCDs. It also incorporates two 1200 x 600 Orbit Semiconductor CCDs for active, close-loop flexure compensation. The science mosaic CCD controller system reads out all eight science CCDs in 40 seconds while maintaining the low noise floor of the MIT/Lincoln Lab CCDs. The flexure compensation (FC) CCD controller reads out the FC CCDs several times per minute during science mosaic exposures. The science mosaic CCD controller and the FC CCD controller are located on the electronics ring of DEIMOS. Both the MIT/Lincoln Lab CCDs and the Orbit flexure compensation CCDs and their associated cabling and printed circuit boards are housed together in the same detector vessel that is approximately 10 feet away from the electronics ring.

Each CCD controller has a modular hardware design and is based on the San Diego State University (SDSU) Generation 2 (SDSU-2) CCD controller. Provisions have been made to the SDSU-2 video board to accommodate external CCD preamplifiers that are located at the detector vessel. Additional circuitry has been incorporated in the CCD controllers to allow the readback of all clocks and bias voltages for up to eight CCDs, to allow up to 10 temperature monitor and control points of the mosaic, and to allow full-time monitoring of power supplies and proper power supply sequencing. Software control features of the CCD controllers are: software selection between multiple mosaic readout modes, readout speeds, selectable gains, ramped parallel clocks to eliminate spurious charge on the CCDs, constant temperature monitoring and control of each CCD within the mosaic, proper sequencing of the bias voltages of the CCD output MOSFETs, and anti-blooming operation of the science mosaic.

We cover both the hardware and software highlights of both of these CCD controller systems as well as their respective performance.

Keywords: CCD controller, CCD readout software control, SDSU-2, DEIMOS

1. GENERAL OVERVIEW

To properly convey the relationship between all components within DEIMOS' CCD imaging systems, a generalized description of the overall architecture is necessary.

The principal components of DEIMOS' imaging systems are the CCD detector vessel which houses both the 8K x 8K science mosaic and the two 1200 x 600 flexure compensation CCDs, the liquid nitrogen (LN2) can that provides the cold strap connection to the mosaic, the three electronics boxes that are mounted onto the Detector Vessel, and the two CCD controllers that are cabled to their corresponding electronic boxes. The imaging systems are shown schematically in Figure 1 and pictorially in Figure 2.

Instrument Design and Performance for Optical/Infrared Ground-based Telescopes,
Masanori Iye, Alan F. M. Moorwood, Editors, Proceedings of SPIE Vol. 4841 (2003)

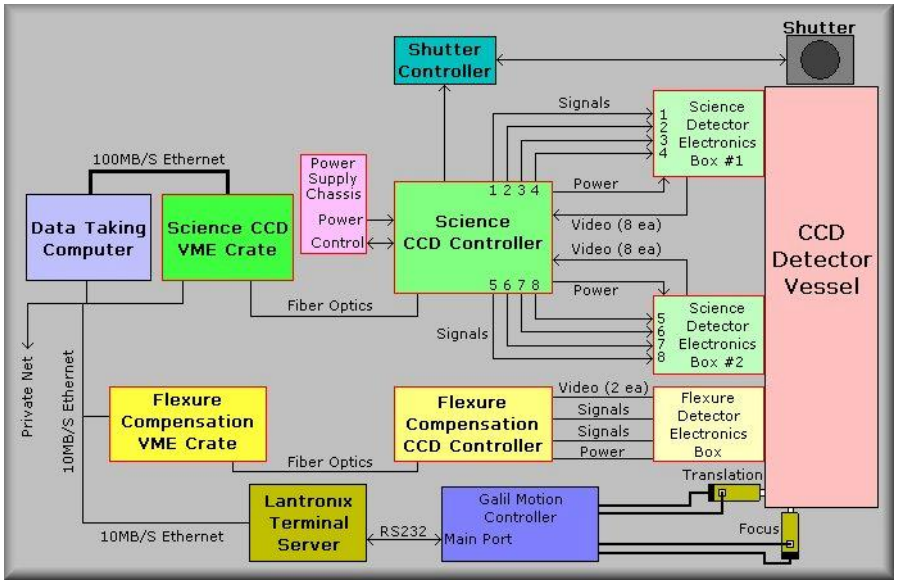

Figure 1. Block diagram of the CCD imaging systems within DEIMOS

Figure 2. DEIMOS' CCD imaging systems in the IR lab at Lick Observatory.
The science mosaic controller is on the left side of the table, the FC CCD controller is on the right, and the power supply for the mosaic controller is between the two. The detector vessel and its attached electronics boxes are in the foreground and the ion pump controllers are on the lower shelf of the table.

The science mosaic consists of eight 2K x 4K MIT/LL CCDs with 15 micron pixels. These CCDs are 45 microns thick for enhanced red response and have an optimized AR coating. Each CCD was individually aligned to the mosaic and shimmed for height achieving a total flatness for the science mosaic of 10.3 microns rms. As shown in Figure 3, each CCD is mounted onto an aluminum nitride package that is attached to a molybdenum block from which its silver cold straps are attached. The cold straps are connected to a thermal copper spider that sits directly below the CCD array which in turn is connected to the cold strap of the LN2 can. Located on each CCD's set of cold straps is a clamping pad where a heater resistor and temperature diode are located. As their names imply, the resistor and diode for each CCD are used to regulate the operating temperature of each CCD individually. Due to thermal cross-coupling across the mosaic and also slight variations in each CCD's individual silver cooling straps, there are slight temperature offsets ranging from three to five degrees C from each CCD to the average mosaic temperature. However, the temperature for each CCD is maintained within +/- 0.25 C of that CCD's temperature set point, as is the average mosaic temperature.

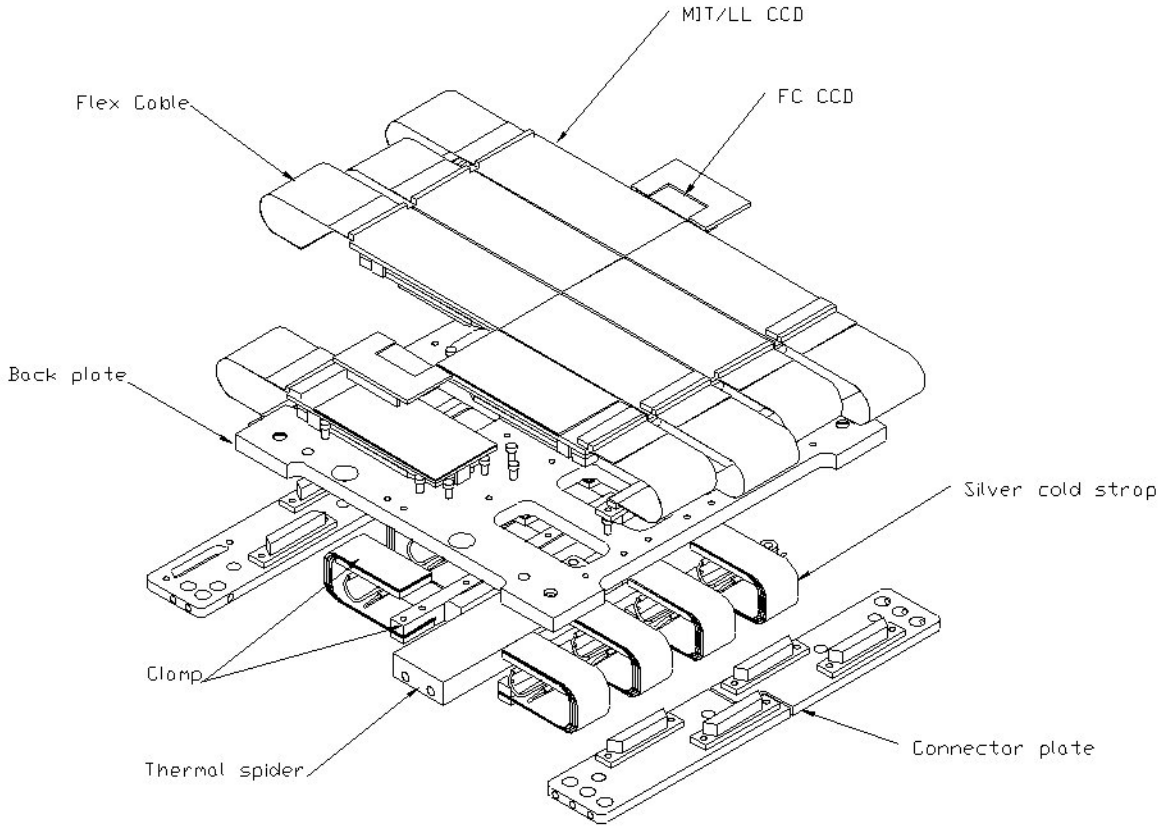

Figure 3. Science Mosaic Assembly

Two 1200 x 600 Orbit Semiconductor devices comprise the flexure compensation CCDs that are used for active, close-loop flexure compensation of DEIMOS' science mosaic images.[1] Each FC CCD is located on either side of the science mosaic and also mounted to aluminum nitride packages. The FC CCDs however are not thermally connected to the cooling copper spider of the science mosaic so their operating temperature is not regulated. The complete and assembled +mosaic is shown in Figure 4.

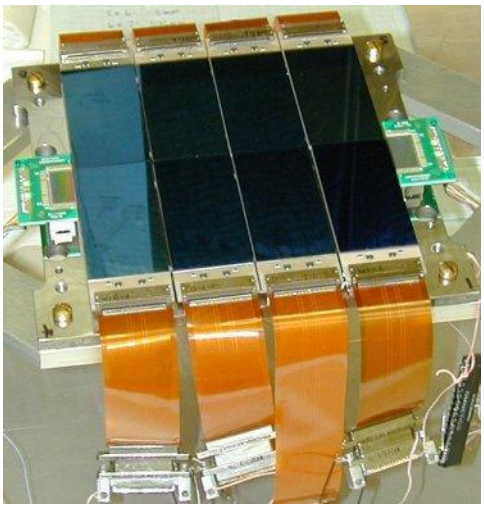

Figure 4. Science Mosaic with FC CCDs located on the left and right sides

The LN2 can consists of an inner liquid nitrogen holding can that is constructed of aluminum, mirror polished both inside and outside, and thermally isolated from the outer stainless steel vacuum can by G10 standoffs and .030" thick stainless straps. A 0.75" copper cold finger, screwed into a boss in the side of the inner can and protruding through the side of the vacuum can, attaches to an identical copper cold finger on the side of the detector vessel via a flexible coupling of pure silver straps. This cold finger attachment is thermally protected from the environment by a separate vacuum housing. The LN2 cooling system has a useable capacity of 24.1 liters and a hold time on Keck II of over 34 hours with a regulated mosaic temperature of –115 C.

The electronics boxes house the CCD preamplifiers and also contain analog switches that allow complete isolation of all signals to each CCD from its respective CCD controller. The science mosaic requires two electronics boxes, each one housing the electronics for four CCDs. The FC CCDs require only one electronics box that is mounted on the rear of the detector vessel. Power is provided to each electronics box from a 17-pin cable originating from the respective CCD controller. A 61-pin cable for each CCD passes all of its clocks and bias voltages from the associated CCD controller to the CCD's corresponding electronics box where the signals are passed through the analog switches and into the detector vessel to the CCD. Each CCD's video output is passed out of the detector vessel to its preamplifier, the output of which is then cabled to the CCD controller via RG-174 coax. Shown in Figure 5 is the DEIMOS dewar with the two science mosaic electronics boxes mounted on either side. Also shown is the internal wiring and board arrangement of one of the electronics boxes.

Each CCD controller consists of a set of SDSU-2 CCD controller boards and a set of UCO/Lick Observatory boards; the controllers generate all of the signals required for each CCD and perform the per pixel signal processing. More detail of the controller design follows in Section 2. The power supply for the science mosaic CCD controller is housed in a separate crate due to the large physical size of the power supplies required by the mosaic controller.

The VME crates provide the link between each CCD controller and instrument computer; they serve as an image buffer and protocol converter (see Sect. 3.2). Each of the two VME crates consists of a Motorola MVME2304 CPU card, a Chrislin 256MB Memory board, and a SDSU-2 VME fiber interface board. They connect to their respective CCD controller via a duplex fiber optic cable that is routed through the telescope and instrument cable wraps. All commands to the CCD controllers and status, telemetry, and image pixel data from the controllers flow via these fiber-optic cables.

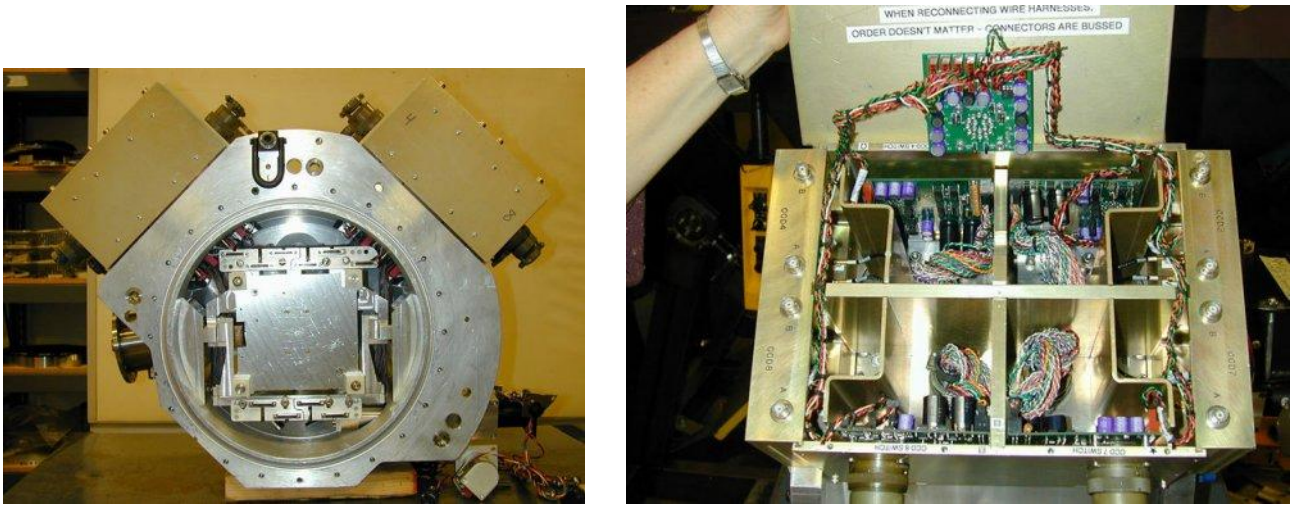

Figure 5. left, DEIMOS dewar with the science mosaic electronic boxes mounted on top; right, electronics box internals

2. ELECTRONIC HARDWARE

2.1 CCD connections within the Detector Vessel

The main impetus behind the arrangement of the CCD connections within the detector vessel is to provide a modular architecture for the CCD mosaic electronics, in which the printed circuit boards, connectors, and cabling are broken down on a per-CCD-basis, thus eliminating any cross-wiring or poor-grounding issues, and ensuring ease of assembly and trouble-shooting.

To facilitate the routing of all signals from both DEIMOS' science mosaic and FC CCDs to the "outside world" via the detector vessel's hermetic connectors, a connection scheme was developed to minimize the amount of individual wiring for each CCD. Each CCD has its own dedicated Hermetic Connector board and CCD Interconnect board from which it receives all of its signals. Connectors on these two boards support a daughter-board arrangement, in which the CCD Interconnect board connects to and sits on top of the Hermetic Connector board. The Hermetic Connector board is soldered directly to a 61-pin hermetic connector on the wall of the detector vessel and passes the SDSU-2 Controller signals to the CCD Interconnect board. The purpose of each CCD Interconnect board is to create the one-to-one mapping of the incoming clock and bias signals to the CCD's input pins. There is one type of CCD Interconnection board for the MIT/Lincoln Labs CCDs and another type for the FC CCDs. Additional filtering for the CCD bias levels and the AC-coupling capacitors for the two CCD video outputs also reside on the CCD Interconnect board. Two SMB connectors on the board pass the two CCD video outputs to hermetic SMA connectors that reside on the detector vessel's walls. Because the hermetic SMA connectors' shields are not isolated from the wall of the detector vessel and therefore at a chassis ground potential, a third hermetic SMA connector is utilized to pass the signal ground of the CCD to the CCD's preamplifier.

The detector vessel printed circuit boards are made from a Teflon-like dielectric, RO4003 from Rogers Microwave Products, which passes the NASA outgassing test. Wiring harnesses and board-to-board connectors that also pass the outgassing test are custom made by the Omnetics Connector Corporation. Each CCD is connected to its associated CCD Interconnect board by a 36 AWG wire harness. The two video outputs from each CCD are also part of the CCD's wire harness but instead are passed through 40 AWG coax cable to insure proper shielding and to eliminate crosstalk between CCD video outputs. Figure 6 shows an earlier version of the detector vessel's FC plate with its associated internal CCD printed circuit boards and wiring harnesses.

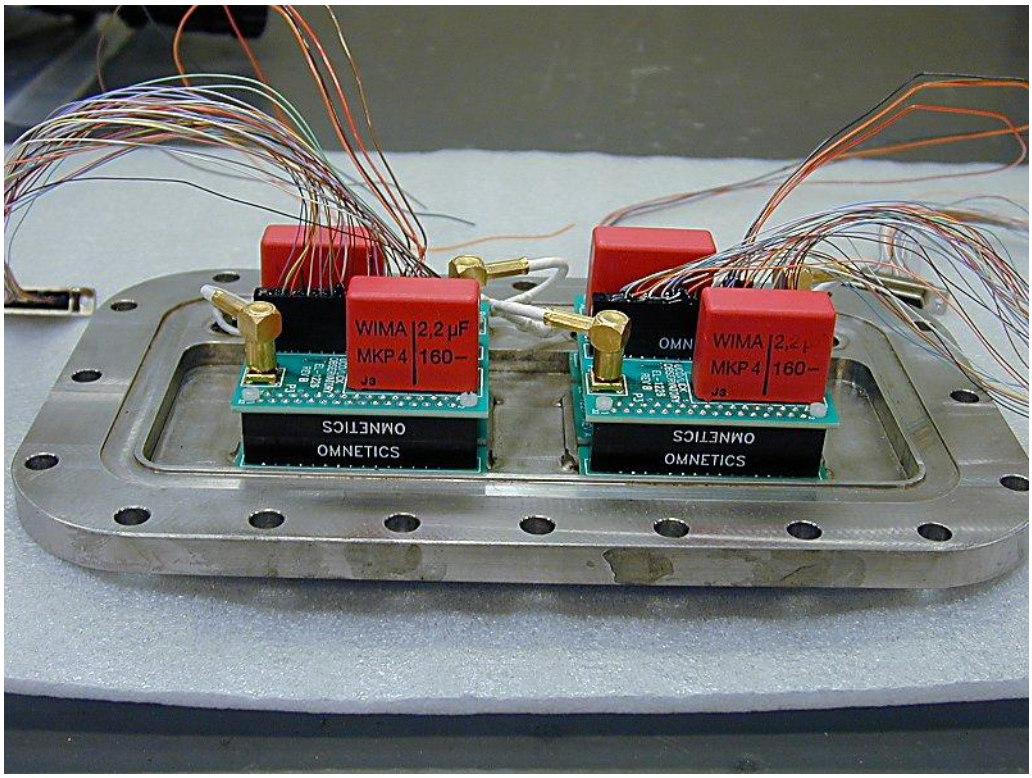

Figure 6. Earlier version of the detector vessel's FC plate indicating the internal circuit boards & CCD wiring harness. The harness is attached to the CCD Interconnect board, which piggy-backs onto the Hermetic Connector board.

2.2 Electronic Boxes

The UCO/Lick Observatory printed circuit boards that reside in the electronics boxes for both the science mosaic and the FC CCDs are the Analog Switch board, the CCD preamplifier, and the Power Filter board.

The Analog Switch board contains semiconductor switches that allow complete isolation of all signals to each CCD from its respective CCD controller. The Analog Switch board receives its CCD signals from the CCD controller via a 61-pin twisted, shielded pair cable. The board is soldered directly onto the back of the 61-pin connector and is attached to the electronics box wall. This eliminates any hand wiring of the 61-pin cable connector to the board, keeping the signals in a clean and shielded environment. All signals to the CCD are switched, including the CCD's Vdd voltage which is switched using discrete MOSFETs. The discrete design of the Vdd switches uses the existing +16, -16, and +5 supplies and removes the need for additional supply voltages to be passed to the electronics box. The CCD bias voltages are also filtered on this board. The Analog Switch board passes the "switched" CCD signals to the associated CCD's hermetic 61-pin connector on the detector vessel wall via a wiring harness composed of twisted pairs (CCD signal with ground).

The CCD preamplifier board, as expected, amplifies the CCD's AC-coupled video output. The preamplifier is a discrete design allowing easy altering of the response and sensitivity of the amplifier for a given CCD and has an output stage capable of driving long lengths of coax cable to maintain the fast per pixel processing times required by DEIMOS. Because the physical distance from each CCD's video output to its associated preamplifier input is rather long (10 or 18 inches) within the detector vessel, compensation of the preamplifiers was necessary and the discrete design facilitated the effort. For the science mosaic of MIT/LL CCDs, each preamplifier's gain was set to 5.5 to accommodate the CCDs' higher sensitivity. The FC CCD preamplifiers have a gain of 28.4. With the added compensation for the long length of coax from each of the CCDs' video outputs, both types of preamplifiers still have respectable settling times of .01% within 500 nanoseconds.

The Power Filter board receives the power (+16, -16, and +5 volts) from the CCD controller and distributes it to all of the boards within the electronics box. Each power input line is filtered with a pi filter and then sent to each board via a separate cable. The Power Filter board is also soldered directly onto the back of its connector, a 17-pin cable connector, and is mounted on the wall of the electronics box.

2.3 CCD Controller

At the heart of each CCD controller is a set of SDSU-2 boards that, under software control, generate the clocking waveforms and bias voltages for each CCD, perform the correlated double-sampling (CDS) of each CCD's preamplifier video output, digitize the result, and transmit the pixel data to the controller's dedicated VME crate. Additionally, the board set controls the timing of the exposures and the shutter and also performs a variety of housekeeping tasks.

Because both the science mosaic and the FC CCD preamplifiers are housed in the detector vessel's electronic boxes, which are approximately 10 feet away from both of the CCD controllers, provisions to the SDSU-2 Video Processing board were needed to accommodate an external preamplifier input. Bob Leach was able to provide such provisions to his SDSU-2 Video board as well as accomodating our additional circuitry modifications and allowed us to order this "Lick-modified" Video board directly. The resultant Lick-modified SDSU-2 Video Processing board is the standard SDSU-2 Video Processing board but with some different components installed and several other components removed.

Each CCD controller also contains a set of UCO/Lick Observatory boards that provide additional functionality and modularity to the overall CCD controller architecture allowing configurations for a variety of CCDs. The following is a list of the UCO/Lick Observatory CCD controller board set:

1.The Utility Support board interfaces the I/O signals, both digital and analog, from the SDSU-2 Utility board and provides the various connectors needed to distribute those signals. It contains additional heater and temperature diode circuitry to support up to eight CCDs of temperature regulation. This board also provides the necessary analog multiplexers and amplifiers to pass any one CCD clock waveform or bias voltage to the SDSU-2 Utility board for software readback. The scaled and multiplexed bias voltages are received from the 61-pin Interconnect board, while the multiplexed clock voltages are received from the respective SMA outputs of each SDSU-2 Clock Generation board. The board provides an optically isolated path for the various shutter control and status lines between the SDSU-2 Utility board to the CCD shutter controller. Software status signals from the SDSU-2 Utility board and the Lick Power Monitor board (see 7.) are also passed through the Utility Support board and optically isolated as well.

2.The 61-Pin Cable Interconnect board acts as a collection point for each CCD's clock and bias signals since these originate from two different SDSU-2 boards (i.e., the Clock Generation and Video Processing boards). This board is mounted on the back panel of the CCD controller and provides a 61-pin connector to mate with the 61-pin cable that connects the CCD's signal path from the CCD controller to the associated electronics box. There is one 61-pin Cable Interconnect board per CCD. A set of jumpers on each board allow for configuration of the CCD's bias voltages for either single or dual amplifier mode and another set of jumpers for configuring the CCD clocks for frame transfer. These jumpers work in concert with cabling from the Clock Cable and Bias Cable Interconnect boards to allow maximum flexibility in configuring the CCD controller for a given CCD and mosaic configuration. This board also contains the scaling resistor networks, amplifiers, and multiplexers to pass any one of CCD's bias voltages back to the Utility Support board for software readback.

3.The Power and Miscellaneous Signal board also mounts on the back panel of the CCD controller and is the source for the 17-pin power cable that connects to each electronics box. Power from the CCD controller is passed through this board and to two dedicated 17-pin connectors. Each 17-pin connector also contains "universal" non-CCD specific temperature diode and heater resistor lines originating from the SDSU-2 Utility board. This board passes the shutter status and control lines from the Utility Support board to the Shutter Controller via an external DB-15 connector.

4.The Clock Cable Interconnect Type I board has a set of DB-37 connectors and it's panel mounts to the front of the CCD controller chassis. This board's function is to connect to each SDSU-2 Clock Generation board via a single 37-pin ribbon cable and pass half of that Clock Generation board's 24 clock lines to each of two 61-Pin Cable Interconnect boards. Thus, each SDSU-2 Clock Generation board provides the clocks for two CCDs.

5.The Clock Cable Interconnect Type II board has a single of DB-37 connector and it's panel mounts to the front of the CCD controller chassis. This board's function is to connect to a SDSU-2 Clock Generation board via a single 37-pin ribbon cable and send a set of three or four additional clocks to each CCD's 61-pin Cable Interconnect board. Typically these additional clock lines would be used for frame transfer operation of a CCD.

6.The Bias Cable Interconnect board has a set of DB-25 connectors and it's panel mounts to the front of the CCD controller chassis. This board's function is to connect to each SDSU-2 Video board via a single 25-pin ribbon cable and send that video board's bias voltages to either one 61-pin Cable Interconnect board for dual amplifier operation of the associated CCD or to two 61-Pin Cable Interconnect boards for single amplifier operation of two CCDs. This again allows flexibility in configuring the CCD controller for a given configuration of devices.

7.The Power Monitor board is mounted in the power supply chassis and it monitors the voltage levels of the power supplies, monitors the AC power to the CCD controller, and accepts inputs from the SDSU-2 Utility board and front panel power switches. Using these inputs, an Altera programmable logic chip determines if the +16V, -16V, and +36V supplies are within tolerance and in what order they may be supplied to the VME style backplane that the SDSU-2 board set plugs into. The board allows the +/-16V to be supplied to the backplane only if all of the power supply voltages are above a predetermined level and the software control of the SDSU-2 Utility board has enabled them. Only then is the 36V supply (from which each SDSU-2 Video board derives a CCD's Vdd voltage) connected to the backplane. If any of the power supplies drop in voltage, the Power Monitor board opens all CCD analog switches that reside in the electronics boxes, disconnecting all CCDs from their clocks and bias voltages, and communicates the condition to the Utility board.

The science mosaic CCD controller shown in Figure 7 consists of a 7U chassis with a 16-slot 3U VME backplane. The following boards reside in the science mosaic CCD controller: SDSU-2 boards - one Timing board, four Clock Generation boards, eight Lick-modified Video Processing boards, one Utility board; UCO/Lick Observatory boards – one Utility Support board, eight 61-pin Cable Interconnect boards, one Power and Miscellaneous Signal board, two Clock Cable Interconnect Type I boards, two Bias Cable Interconnect boards, and one Power Monitor board. The two video outputs from each MIT/Lincoln Labs CCD of the mosaic are connected to one dual preamplifier board within the associated electronics box. Since the SDSU-2 Video Processing board is a dual channel board, the two preamplifiers' outputs are cabled directly to one dedicated SDSU-2 Video Processing board per CCD, thus allowing either single amplifier or dual amplifier readout of the science mosaic.

Figure 7. Science Mosaic CCD controller

The complexity and size of the science mosaic CCD Controller dictated that an external power supply box would be needed. This box was built into a 4U size chassis. The box contains a 12 ampere +5V supply, a 9 ampere +16V supply, a 9 ampere -16V supply, and a 800 milliampere 36V supply. The Power Monitor board is located in the power supply chassis and, because it was not designed to handle the connection of high amperage supplies to the CCD controller backplane, a set of relays are employed to disconnect or switch these supplies from the backplane. The power supply

chassis connects to the CCD controller via two large cables. Each conductor for the power supplies is 14 AWG wire to ensure adequate current carrying capabilities. To retain a low impedance connection of the chassis grounds, a 000 AWG wire is used to connect the CCD power supply chassis to the CCD controller chassis. All power supply sense lines are 18 AWG and the Power Monitor control and status signal wires are 20 AWG. The power supply connects to a constant on-line UPS to further protect the system against power transients and dropouts.

The flexure compensation CCD controller consists of a 4U chassis with a 6-slot 3U VME backplane. The following boards reside in the FC CCD controller: SDSU-2 boards - one Timing board, one Clock Generation board, one Video Processing board, one Utility board; UCO/Lick Observatory boards – one Utility Support board, two 61-pin Cable Interconnect boards, one Power and Miscellaneous Signal board, one Clock Cable Interconnect Type-1 board, one Clock Cable Interconnect Type-2 board, one Bias Cable Interconnect board, and one Power Monitor board. Each FC CCD is used only in the single amplifier readout mode so only one CCD preamplifier's video output is passed to one channel of the SDSU-2 Video Processing board, requiring only one Video Processing board for the FC CCD controller. Because the FC CCD controller houses only two CCD's worth of electronics, the current requirements of the controller's power supplies are considerably less than those of the science mosaic CCD controller, and therefore their physical size is reduced, allowing them to fit internally within the CCD controller. The Power Monitor board for this controller handles all the switching of power supply voltages to the controller's backplane.

2.4 CCD Shutter

As the DEIMOS shutter is located directly in front of the dewar window, a conventional commercial shutter could not be used because the optical beam is too large at that point. Instead, DEIMOS uses a dual-bladed shutter. To close the shutter, one of the two blades is inserted into the beam, while to open it, both blades are withdrawn. Alternate blades are used for opening and closing so as to ensure equal exposure times across the full extent of the mosaic.

Each of the two shutter blades is driven by a Bimba air cylinder. Compressed air is applied either to one end of the cylinder to drive its attached blade into the beam or to the opposite end of the cylinder to withdraw the blade; the air is switched from one end of the cylinder to the other via a latching air solenoid. The air solenoids (one for each cylinder) are in turn controlled by a dedicated shutter controller unit. The shutter controller monitors the position of the shutter blades via Hall-effect limit switches located at each end of travel on each of the two air cylinders. The shutter controller receives a single TTL-level command input (high to command the shutter open, low to command it closed) from the science CCD controller, and via logic encoded in a PAL chip, determines what commands to send to the air solenoids in order to move the appropriate blade to achieve the requested shutter state.

2.5 Design Considerations

Due to the sharp rise time and fall times (7 nanoseconds) of the SDSU-2 clocks and the fast settling times needed to achieve our per pixel processing time of 5 to 7 microseconds, the CCD imaging systems are treated as true high speed, high accuracacy electronic systems and thus incorporate the required proper grounding, shielding, and cabling within their hardware designs.

Great attention was made to the routing and return of the ground currents from the various electronic printed circuit boards and their associated cabling within each CCD's signal path. All clock and bias voltages grounds for each CCD were returned directly from their "originating" SDSU-2 Video Processing board or Clock Generation board and are not allowed to deviate along other "paths". These clock and bias voltage grounds become each CCD's signal ground which is isolated from any other CCD's signal ground except at the CCD controller's backplane, which ties all of the CCD signal grounds together. Except for one "star-ground" connection strap that ties all the CCD controller's power supply grounds to the chassis of the CCD controller (and hence DEIMOS itself), all grounds (e.g., CCD signal, analog, and digital) are isolated from DEIMOS' chassis throughout the entire imaging system. Hence a true Faraday cage, at DEIMOS' chassis ground potential, encloses and shields the CCD controller, the cabling between the controller and its associated electronics box(es), and the detector vessel. This insures that no hostile environmental noise or EMI will couple into any CCD's signal or ground return path. Careful printed circuit board layout, incorporating the isolation of certain grounds and signals and also the particular routing and shielding of signals within each circuit board's layer, was employed on all the UCO/Lick Observatory CCD controller boards and the electronics box boards. Where needed some

digital control signals (e.g., shutter control, Power Monitor signals from the science mosaic CCD controller power supply box) were optically isolated to insure the proper return of ground currents within the system.

Cabling was quite important to the design of DEIMOS' imaging systems to maintain the desired grounding scheme and provide proper shielding and isolation of each of the CCD's signals. As mentioned earlier in Section 2.3 each CCD's set of bias voltages and clocks are passed through its 61-pin cable Interconnect board, through a 61-pin cable to its Analog Switch board residing in its associated electronics box. The 61-pin cable consists of a Belden #1518B cable that has two 61-pin MS-3116 (24-61S) connectors attached to each end. The Belden cable is a high-flex snake cable that consists of 32 pairs of #24 AWG wire. Each pair is twisted and individually shielded, and the cable has an overall foil shield surrounding all 32 pairs. This overall foil shield is connected via external banana jacks to the chassis of the CCD controller and the electronic box, maintaining the Faraday chassis ground connection. Each CCD signal (clock or bias voltage, or spare signals) are twisted with a signal ground, thus there is a maximum of 30 signals that can be passed from the CCD controller to the electronics box. Jumpers on the 61-pin Cable Interconnect board and the Analog Switch board allow the grounding of unused signals or spares.

Because each 61-pin cable is quite flexible, we discovered after extended periods of cable installation and removal that the solder joints of the cable's individual wires to the 61-pin connectors suffered a lot of strain, resulting in some wires being broken off from their connector. We then devised and machined special cable connector stiffeners that attach to the 61-pin connectors and also created a better wire and connector assembly technique. Combined they provide the rigidity and strain relief needed to protect the cable's individual wires.

Of note on an architectural basis, as briefly alluded to in Section 2.1, the CCD imaging systems were designed such that each CCD is electrically independent from any other CCD; there are no shared clocks or bias voltages between any CCD within the science mosaic. Each CCD's amplifier has its own set of indivual voltages which allows the software to simply shut down a particular CCD's output amplifier (either one or both) without affecting the other CCDs in the mosaic. This proved to be quite beneficial in our early engineering mosaic testing, where we encountered CCDs with glowing output amplifiers or other devices that simply weren't operable, yet we could still read out the other working devices in the mosaic.

3. SOFTWARE CONTROL

The software associated with the DEIMOS CCD subsystem is implemented in several layers that are distributed across different processors (e.g., digital signal processors, RISC processors on single board computers, etc.). These processors are in turn distributed between those components of the subsystem that reside within the instrument and those located off of the telescope (see Figure 1). The software (and the hardware upon which it runs) can be broken down as follows:

3.1 High-level software

The high level software runs on the data taking computer (currently a Sun SPARC Ultra-Enterprise 450 with 4 CPUs, 3.25 GB of RAM, and about 212 GB of RAID disk) which is located in the Keck-II computer room and which includes software for: capturing CCD mosaic image pixel streams and relevant FITS header information (e.g., exposure times, mosaic geometry, telescope coordinates, etc.) and writing these CCD mosaic images to disk as multi-HDU format FITS files; displaying images and providing quick-look analysis; graphical user interfaces (GUIs) for specifying exposure parameters, initiating exposures, and monitoring their progress; and optional command-line, keyword-based control of major CCD subsystem functions for support of image acquisition and diagnostic scripts. This software is implemented using C and Tcl/Tk.

CCD mosaic pixel streams are captured, de-scrambled, combined with their relevant FITS header information and formatted into a multi-HDU FITS file in shared memory by the lickserv2 application. This FITS file contains definitions for several world coordinate systems (WCS) that provide a mapping between DEIMOS science mosaic CCD pixels and various focal plane and sky coordinate systems.

As the image reads out of the mosaic CCD detectors, lickserv2 formats the image into the shared memory FITS file and periodically triggers the ds9 image display tool (see below) to paint the newly-received pixels onto its display in real-

time. Since the definition for each WCS mapping is contained in that FITS file, ds9 is able to map the image display's cursor position into any of the coordinate systems so defined. Once the complete image is received, lickserv2 invokes the write_image application, which writes the mosaic image to a locally-attached RAID disk. The write_image application also appends to the FITS file a series of FITS table extensions that provide a complete definition of the geometry of the slitmask used to generate the image contained in that FITS file.

The ds9 image display program[2] is used to display DEIMOS science mosaic images directly as they are read out of the detectors, and can also display DEIMOS images from disk. The real-time, shared-memory image display extensions to ds9 were developed as a collaboration between UCO/Lick Observatory and the ds9 authors at Harvard. An image statistics extension to ds9 (statsbox) was also developed at UCO/Lick, and provides image centroids and other basic statistics for any user-specified region of the displayed image.

A Tcl/Tk-based GUI, based on the dashboard software suite developed by Clarke[3], provides the observer with control and status monitoring of all aspects of the DEIMOS instrument. Control of exposures for the science mosaic is integrated into this GUI, and the most frequently-used functions (e.g., setting exposure times and types, object names, etc.) and corresponding status displays (e.g., remaining exposure time, detector temperatures, shutter status) are directly available from the top-level screen. Less frequently-used functions (e.g., selection of readout amplifiers, CCD binning, etc.) are provided via a CCD-detail pop-up window. A separate Tcl/Tk-based GUI is provided for control of and display of status from the flexure compensation control loop. Both GUIs interact with the lower-level CCD software via a series of keyword libraries, which enable reading, writing, and monitoring of individual keywords that are mapped to specific functions and parameters of the various subsystems that comprise DEIMOS.

All major functions of both the science mosaic and flexure compensation CCD systems can also be controlled via keywords using the Keck Tasking Library (KTL)[4] that is common to all Keck instruments, including the two telescopes. Under this model, each individual parameter (e.g., the elapsed time of the current exposure) of a given subsystem (e.g., the DEIMOS science CCD mosaic) is mapped to a keyword that conforms to the naming conventions for FITS keywords (i.e., alphanumeric names up to 8 characters long)[5]. Each subsystem is named as a KTL service (e.g., the DEIMOS science CCD mosaic is the deiccd service, while the Keck II Telescope control system is the dcs2 service). The keywords for a given service are implemented as a sharable library that can be dynamically linked to by application software. A common set of generic routines (that are shared across all Keck instruments and both telescopes) allow the values of keywords for each KTL service to be read, written, or dynamically monitored. The same set of KTL keyword-based operations is used to write GUIs and to implement scripts.

Scripts can be implemented using any of the Unix command shells (e.g., sh, csh, tcsh, bash, etc.) or with Tcl/Tk, and can be used for any batched or automated sequences of commands that would be tedious to carry out using an interactive GUI. Such scripts and GUIs inter-operate seamlessly, since both interact with the underlying subsystems via the same set of keyword operations; operations performed via scripts are immediately visible on the status displays of the GUI, and functions performed via the GUIs are equally accessible to scripted routines. This scripting capability proved extremely valuable during DEIMOS test and integration, and enabled various performance and diagnostic tests to be automated and run overnight while staff were asleep[6].

3.2 Mid-level software

The mid-level software is implemented entirely in C and runs on the respective VME crates (see Fig. 1) under the VxWorks operating system. This level of software is mostly invisible to the observer, since the observers never log into the VME crates nor do they interact directly with any of the routines running on the VME crates. All observer interaction with the routines on the VME crates is via the keyword interface. This same VME crate architecture has been successfully used by all of the other Keck optical instruments: the High Resolution Echelle Spectrometer (HIRES), the Low-Resolution Imaging Spectrograph (LRIS), and the Echellette Spectrograph and Imager (ESI).

The science VME crates have sufficient RAM (256 MB) to buffer a complete DEIMOS mosaic image. As the mosaic detectors are read out, the pixel stream flows over the fiber optic cable from the CCD controller, and the SDSU-2 VME fiber interface board effectively transfers the pixels into VME memory via DMA. Simultaneously, the Motorola MVME2304 single board computer in the VME crate packages the incoming pixel stream into packets which are then

transmitted to the data taking computer via the 100-Mbit/sec DEIMOS private Ethernet network, where they are received by the lickserv2 process and displayed in real time by the ds9 image display process.

By buffering the image in the VME crate, we avoid installing any custom hardware in the data taking computer, thus ensuring that it remains a generic Unix box that can easily be upgraded in the future as technology improves. This model also ensures that no pixels get dropped should the data taking computer become overloaded by observer-initiated applications.[7] Thus the VME crate can be viewed as an image buffer and protocol converter. It also performs various house keeping functions (e.g., mosaic temperature control, system loop-back diagnostics) that were too difficult in the limited memory and programming environment of the digital signal processors used in the CCD controllers.

3.3 Low-level software

The low-level software runs on the timing and utility boards of the SDSU-2 CCD controllers[8] for the science mosaic and FC systems. It is implemented in DSP-56000 assembly language, and runs from EEPROMs that have very limited program space. Accordingly, only those functions that require rigid and precise high-speed timing (e.g., CCD clock waveform generation, CCD video processing and A/D conversion) or precise voltage control (e.g., CCD bias voltage generation) are performed here, along with some ancillary house keeping functions (e.g., control of the shutter and the dewar heater resistor voltages).

The low-level DSP software that runs in the timing and utility boards is derived from the version that the board vendor shipped. However, to operate the DEIMOS CCD mosaic, significant customizations were implemented by UCO/Lick to support a variety of new capabilities. These include support for: multiple readout speeds, multiple readout configurations (dual amplifier, single amplifier left, single amplifier right), multiple mosaic configurations, ramped parallel clocks (to eliminate spurious charge generation), antiblooming clocks, 8 channels of temperature diodes and heater resistors, precise timing for dual-bladed pneumatic shutters, MIT/LL CCD-specific power-up/power-down sequences, and extensive loopback diagnostics for all clock and bias voltages. While space does not permit detailed description of all of these customizations, some are particularly distinctive and are described below.

3.3.1 Loopback diagnostics

Given the very significant cost and time required to obtain the eight CCDs in the DEIMOS science mosaic, considerable effort was made to ensure that improper voltages were not applied to the mosaic. When the CCD controller is first powered up, the analog switches inside the dewar electronics boxes are commanded open (thus isolating the mosaic CCDs from the controller), and the analog power supplies are disconnected from the backplane into which the various SDSU-2 boards are mounted. The output of the various analog supplies (+16, -16, and +36 volts) are checked via a set of voltage dividers and analog multiplexers which feed an analog-to-digital converter on the SDSU utility board. Only if the supplies are within specification are their voltages then applied to the backplane. Next, the various digital-to-analog (D/A) converters on the SDSU-2 Clock Generation and Video Processing boards are commanded to the specified voltages for each of the respective clocks and bias signals. Each clock and bias voltage is then individually checked via a series of voltage dividers and analog multiplexers on the 61-pin cable interconnect boards which are in turn multiplexed via the utility support board and then fed into the A/D on the utility board. For the science mosaic the loopback diagnostics are quite extensive and involve the checking of several hundred different signals.

If any clock or bias voltage is outside its specified range, the power-up sequence is aborted, the analog switches in the dewar electronics remain open, and the mosaic remains isolated from the controller. Otherwise, if voltages for all of the clock and bias signals are within specification, the analog switches are closed, and then all power supplies, clock, and bias voltages are again individually checked to see if any voltage goes out of specification once the load of the CCD is attached. In addition to these software-based checks, the power monitor board is constantly monitoring the output of the various power supplies, and will open the analog switches in the dewar electronic boxes in the event that any of the power supply voltages drops below specification.

3.3.2 MIT/LL CCD-specific power sequencing

In addition to these power-up diagnostics, it proved necessary to add power-sequencing logic for several of the bias and clock voltages due to the novel characteristics of the output amplifiers on the MIT/LL CCDID 20 devices.[9] These very-low-noise amplifiers contain some extremely tiny structures (much smaller than those we have encountered on any of the

CCDs we have previously worked with) which are easily damaged if the output drain and reset drain voltages are not sequenced in the proper order on power up or power down. The sequence recommended by MIT/LL is that reset drain be powered up prior to output drain and that output drain be powered down before reset drain. In addition, our tests at UCO/Lick indicated that for at least some subset of the CCDID 20 devices, powering up reset gate prior to reset drain could potentially cause damage. Accordingly, we have implemented an extremely conservative power-up sequence designed to minimize the risk of any transient currents damaging the delicate structure of these amplifiers.

3.3.3 Precise timing for dual-bladed pneumatic shutter

Although early "on the bench" tests of the DEIMOS shutter and shutter controller showed that the shutter blade opening and closing times were essentially equal and repeatable at the level of a few milliseconds, these tests were only conducted at one orientation with respect to gravity; later tests conducted once the spectrograph was fully assembled established that the blade opening and closing times varied with the physical position angle of the instrument due to the varying gravitational loads on the blades and cylinders.

In addition, while the latency of the air solenoids (i.e., the delay between when the solenoid is commanded to switch the compressed air from one of its output ports to the other and when that switch actually occurs) is quite repeatable if the solenoids are fired frequently (as was the case during those early tests), the latency becomes quite variable if the solenoids are left idle for more than a few minutes; the longer the solenoids sit idle, the greater their latency. In cases where the shutter has been idle for several hours, these latencies have been observed to be as high as several hundred milliseconds. (The cause of this variable latency has not yet been established, although we suspect it may reflect some sort of variable friction between the magnetic puck, used to switch the air flow, and the barrel through which the puck moves, perhaps due to asymmetries between the barrel and puck coupled either with contaminants or inhomogeneities in any lubricants used within the barrel.)

As a result of these two problems, the exposure timing accuracy of the DEIMOS shutter is affected by both the physical position angle of the instrument and the interval since the shutter last changed state. In many cases (especially for short exposures), the shutter fails to meet the DEIMOS exposure timing specification of 1% accuracy. For these reasons, air solenoid operated cylinders should not be relied on for such time-critical operations. Rather, alternate devices (such as linear servo motors) that are capable of providing more reliable timing should be considered for driving dual-bladed shutters like the one in DEIMOS.

To partially compensate for this deficiency in the DEIMOS shutter, custom logic was added to the DSP software for the utility board in the DEIMOS science mosaic CCD controller. This software monitors the signals from the Hall-effect limit switches located at the ends of each of the two air cylinders that drive the shutter blades. By precisely measuring the intervals between state transitions of these limit switches relative to the timing of shutter commands, the software is able to accurately measure the blade motion times and air solenoids latencies; using these measurements in conjunction with the commanded exposure time, the software derives a measure (accurate to 1 ms.) of the actual exposure time that was achieved. All of these measurements are available to the higher-level software as keywords and are recorded in the FITS header for each exposure. Thus, while they achieved exposure time may not be within 1% of the requested value, the achieved value is recorded to better than 1% accuracy.

4. DESIGN ISSUES

Two significant issues arose during the design and implementation of the DEIMOS science mosaic CCD controller: 1) whether one or two controllers should be used to operate the science mosaic, and 2) what design would enable rapid reconfiguration between single amplifier and dual amplifier readout modes.

4.1 One controller or two?

Based on experience with the first generation SDSU controllers (which used a non-terminated backplane), there were concerns that problems with signal reflections and propagation delays along the backplane might prevent 8 video processing boards, 4 clock generation boards, a timing board, and utility board from all operating reliably over such a long (14-slot) backplane. As a result, we considered the possibility of building two separate CCD controllers, each of which would be used to read out half of the mosaic.

This approach was used in the initial implementation of the 8-CCD mosaic (UH8K) built by the University of Hawaii Institute for Astronomy; that implementation used two separate first-generation SDSU controllers, each of which read out 4 CCDs. However, since the clocking waveforms generated by these two controllers were not synchronized, the two halves of the mosaic could not be read out simultaneously because crosstalk between the two asynchronous sets of waveforms resulted in unacceptable readout noise. Each half of that mosaic had to be read out separately, thus doubling the readout time.[10]

The second-generation SDSU controllers were designed to overcome both of these problems. Compared to the first-generation SDSU boards, the second-generation boards all use faster, higher-power tri-state driver chips to drive signals onto the backplane. Also, the backplane is terminated. As a result, signal reflections and propagation delays are significantly reduced, and a much larger number of boards can reliably inter-operate on the same backplane. In addition, the second-generation timing boards include provisions for synchronizing the system clocks of two separate controllers; that is a necessary precondition for being able to read out both halves of a mosaic simultaneously using two separate controllers.[*] Accordingly, the second-generation SDSU architecture should, in theory, be able to support either a single or dual controller model for reading out an 8 CCD mosaic.

For a variety of reasons, we choose the single controller model. First the hardware costs were lower, since we needed fewer chassis, backplanes, power supplies, boards, and long fiber optic cables. Second, the software was less complex, since the entire mosaic was accessible from a single controller, and no elaborate schemes were needed to synchronize the instruction streams of two separate controllers. Third, we only had to debug the wiring on one chassis, not two. Fourth, at the time the choice between the two models had to be made, the timing board clock synchronization scheme was very new and not extensively tested, so the single controller model seemed less risky.

The DEIMOS single-controller system has proven quite reliable. However, now that others have demonstrated that the clock synchronization schemes of the second-generation timing boards work reliably and can achieve waveform synchronization at the level of 6 nanoseconds[11], a dual-controller model is worth considering for operating large CCD mosaics using second-generation SDSU electronics. First, one-of-a-kind systems like the DEIMOS controller are more difficult to maintain, since having two identical units between which suspect parts can be swapped greatly simplifies trouble-shooting. Also, two smaller, modular units may be easier to package and fit into an instrument than one large, monolithic controller. Third, although the dual-controller model requires that twice as much fiber optic cable be run through the instrument and telescope cable wraps (one pair of fibers for each controller), that additional fiber provides greater bandwidth, which may prove useful as the capabilities of high-speed, low-noise CCD output amplifiers improve.

4.2 Optimizing re-configurability of readout modes

At the time the DEIMOS science mosaic controller was designed, there was considerable uncertainty regarding which type of CCD (Lick/Orbit, MIT/LL, EEV) would be used for the mosaic and little assurance that all 8 devices would have two working amplifiers. Given that uncertainty, we were reluctant to incur the full expense of a dual-amp per CCD (i.e., 16 channel) readout system until we received the actual CCDs for the mosaic.

4.2.1 Initial system: single-amplifier per CCD

Although we designed the controller chassis with enough backplane slots to accommodate enough boards for a dual-amp per CCD system, we initially purchased only 4 SDSU-2 Video Processing (VPROC) boards, which was just enough to support single amp readout from each of 8 CCDs; each dual-channel VPROC board handled the video signals from two CCDs. Modular cabling between the Bias Cable Interconnect boards and the 61-pin Cable Interconnect boards (CIB), along with jumpers on the CIB boards, enabled us to supply bias voltages to all 8 CCDs using only the four VPROC boards. Both the initial engineering-grade mosaic and the first of the two DEIMOS science mosaics (the "blue" mosaic) were operated in this mode for many months of engineering tests. Once it was established that we had 8 science-grade CCDs all having two working amplifiers, we then ordered 4 more VPROC boards. When they arrived, all that was needed to convert the controller hardware from single- to dual-amp-per-CCD readout was to plug the new VPROC

[*] To achieve true synchronization of clocking waveforms and video processing between two separate controllers, not only must their system clocks be synchronized but their respective instruction streams must be synchronized as well, and that involves added hardware and software complexity on both controllers.

boards into the backplane, shuffle some modular cables between the Bias Cable Interconnect and the CIB boards, and change some jumpers on the CIB boards. No other hardware changes were needed.

4.2.2 Software-select between single- and dual-amplifier readout modes

Once we had reconfigured the controller hardware for dual-amplifier per CCD operation, it was initially assumed that we would always operate in that mode. However, as a contingency in case problems later developed with any of the CCD output amplifiers (or with the connections to those amplifiers), a software-selectable fall-back to single-amplifier readout mode was implemented. This fall-back is accomplished without any need for hardware reconfiguration; switching between single- and dual-amp readout is accomplished simply by changing the value of a keyword.

This fall-back mode has proven extremely valuable, since after the dewar was shipped to Mauna Kea, installed in the instrument and thermally cycled, a mechanical connection problem developed inside the detector vessel which resulted in the loss of the video signal from the A amplifier of CCD 5. Because there was not enough time in the commissioning schedule for removal of the detector vessel and repair of the faulty connection, DEIMOS was commissioned using this software-selectable single-amplifier readout mode.

We are currently testing some customization of the PAL logic on the SDSU-2 timing board provided by Bob Leach. That customization will enable the timing board not only to transmit pixels from all video channels within a specified range of channel numbers but to selectively transmit only pixels from odd-numbered or only from even-numbered video channels. Along with corresponding changes to the DSP code for the timing board and SDSU-2 VME fiber interface board, these changes should enable us to nearly double the speed of our single-amplifier readout mode by more effectively utilizing the bandwidth of the fiber optic cabling between the controller and its VME crate. These changes are now scheduled for installation in the instrument in mid-August 2002.

5. PERFORMANCE

The performance of DEIMOS' science mosaic imaging system is tabulated in Table 1. The science mosaic was read-out in single-amplifier mode at a per pixel processing time of 7 microseconds in low-gain mode at a regulated mosaic temperature of –115 C.

CCD	output amplifier	Gain(e/dn)	Noise (e)	dark current (e/pixel/hr)
1	B	1.24	2.23	4.19
2	B	1.20	2.02	3.46
3	B	1.27	2.27	4.03
4	B	1.27	2.48	3.80
5	B	1.26	2.79	4.71
6	B	1.25	2.36	4.28
7	B	1.25	2.46	3.33
8	B	1.25	2.40	3.69

Parallel CTE, all devices: 0.999999 (effectively 1.0 in our measurements)

Serial CTE, all devices: 0.999998

Table 1. Science Mosaic performance in DEIMOS on Keck 2 at a regulated mosaic temperature of –115 C

Overall the performance we have achieved from the Lick-modifed SDSU-2 Video Processing board has been quite good. We did observe however that there is slight crosstalk between the two video channels on the video board as a result of the printed circuit board layout. The magnitude of this crosstalk is on the order of 3 DN for a full scale input applied to one of the channel's A/D converter. However due to the loss of the video signal from the A amplifier of CCD 5, as mentioned in Section 4.2.4, we are now running the science mosaic in the single amplifier mode, utilizing only one video channel from each of the eight SDSU-2 Video Processing boards, inherently eliminating this crosstalk issue. When reading only the top half of the mosaic (CCDs 1 to 4) the dual amplifier readout can be utilized when crosstalk is not of a concern for the desired imaging performance.

ACKNOWLEDGMENTS

We thank Bob Leach from Astronomical Research Cameras, Inc. for his assistance in the development of a specialized video processing board that incorporated both his prerequisites and our signal processing requirements and for his help in debugging some of our particular timing issues and needs.

DEIMOS was started in 1992 with a Facilities and Instrumentation grant from the National Science Foundatation (ARI92-14621) in the amount of $1.79 million. The California Association for Research in Astronomy has so far contributed approximately $7.2 million directly to the project, not including the additional costs of CARA liaison and alterations that were needed for the Keck 2 telescope. UCO/Lick Observatory has contributed over $1 million so far in manpower.

The authors also wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Mauna Kea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain.

REFERENCES

1. R. Kibrick *et al*., "A comparison of open versus closed loop flexure compensation systems for two Keck optical imaging spectrographs: ESI and DEIMOS", in *Instrument Design and Performance for Optical/Infrared Ground-Based Telescopes*, M. Iye, ed., Proc. SPIE 4841-155, 2002.
2. http://hea-www.harvard.edu/RD/ds9.
3. D. Clarke, Dashboard: a Knowledge-Based, Real-Time Control Panel, in Proceedings: 1997 Usenix Tcl/Tk Conference, Boston, Mass., 1997.
4. W. F. Lupton and A. R. Conrad, "The Keck Task Library (KTL)", in Astronomical Data Analysis and Software Systems II, R. J. Hanish, R. J. V. Brissenden, and J. Barnes, eds., ASP Conference Series 52, pp. 315-320, 1993.
5. A. R. Conrad and W. F. Lupton, "The Keck Keyword Layer", in Astronomical Data Analysis and Software Systems II, R. J. Hanish, R. J. V. Brissenden, and J. Barnes, eds., ASP Conference Series 52, pp. 203-207, 1993.
6. S. Faber et al., "DEIMOS spectrograph for the Keck II Telescope: integration and testing", in Instrument Design and Performance for Optical/Infrared Ground-Based Telescopes, M. Iye, ed., Proc. SPIE 4841, 2002.
7. R. I. Kibrick, R. J. Stover, and A. R. Conrad, "CCD Data Acquisition Systems at Lick and Keck Observatory, in Astronomical Data Analysis and Software Systems" II, R. J. Hanish, R. J. V. Brissenden, and J. Barnes, eds., ASP Conference Series 52, pp. 277-288, 1993.
8. R. Leach, F. Beale, and J. Eriksen, "New Generation CCD controller requirements and an example: the San Diego State University generation II controller", in Optical Astronomical Instrumentation, S. D'Ororico, ed., Proc. SPIE 3355, pp. 512-519, 1998.
9. Burke, B. E. et al., in Optical Detectors for Astronomy II: State-of-the-Art at the Turn of the Millenium, ed. P. Amico & J. W. Beletic, (Boston: Kluwer), pp 187, 2000.
10. J. Cuillandre, G. Luppino, B. Starr, and S. Isani, "Performance of the CFH12K: A 12K by 8K CCD mosaic camera for the CFHT prime focus", in Optical and IR Telescope Instrumentation and Detectors, M. Iye and A. F. M. Moorwood, eds., SPIE 4008, 1010-1021, 2000.
11. B. Starr, G. Luppino, J. Cuillandre, S. Isani, "Design of the CFH12K: 12k by 8k CCD mosaic camera for the CFHT Prime Focus", in Optical and IR Telescope Instrumentation and Detectors, M. Iye and A. F. M. Moorwood, eds., SPIE 4008, 1022-1033, 2000.