Smithsonian Astrophysical Observatory & Steward Observatory, The University of Arizona

Smithsonian Institution &
The University of Arizona®

End of Quarter Summary

April - June 2014

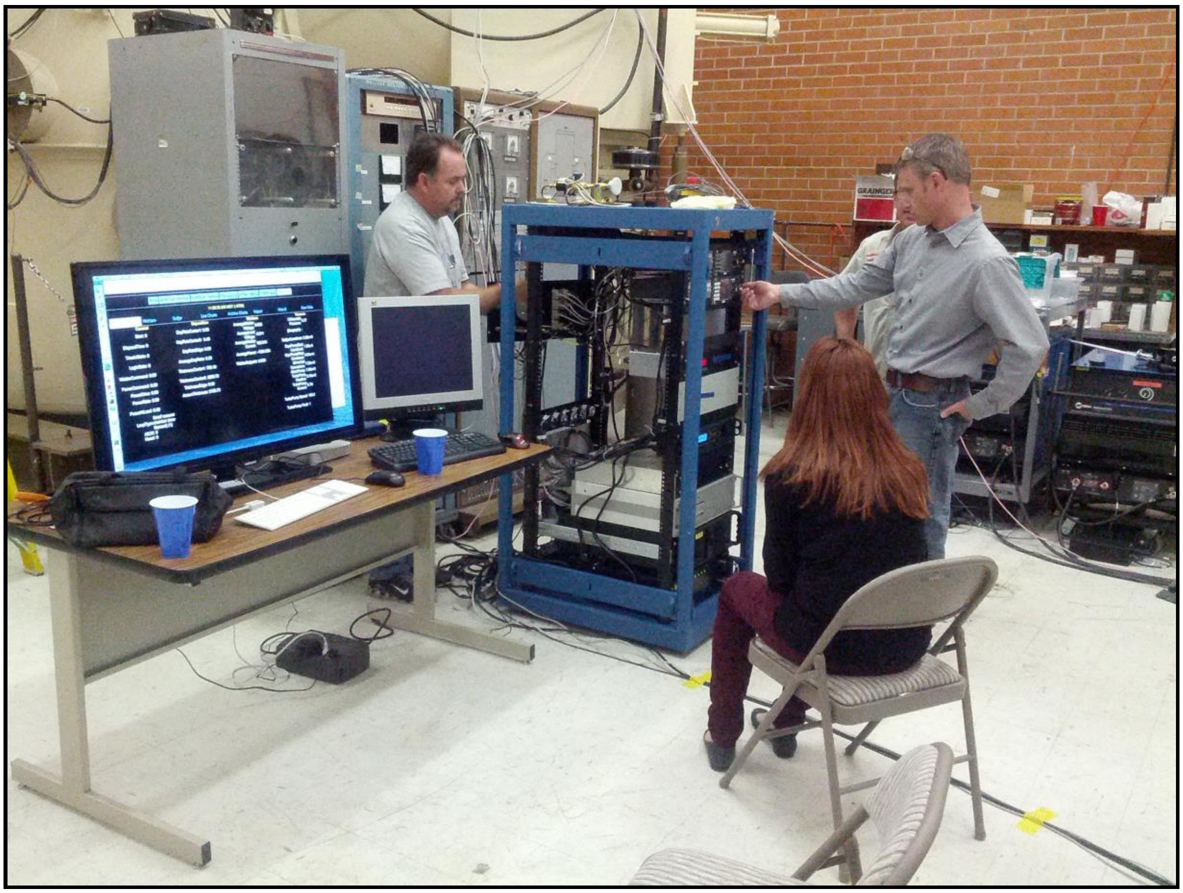

Left to right: C. Knop, D. Gerber (sitting), and W. Goble are shown testing equipment at the Sunnyside facility in preparation for the primary mirror aluminization.

MMT Observatory Activities

Our Quarterly Summary Reports are organized using the same work breakdown structure (WBS) as used in the annual Program Plan. This WBS includes a major category with several subcategories listed under it. In general, many specific activities might fall a tier or two below that. The WBS will be modified as needed in future reports.

Administrative

Staffing

Telephone interviews were conducted in April for the Electronic Technician, Sr. position, followed by site visits to the MMTO with three candidates. Will Bettis accepted the position and started on June 9.

Strategic Planning

The final report from the MMTO Strategic Planning Committee was sent to C. Alcock and B. Januzzi on May 12.

Reports and Publications

There were 15 peer-reviewed MMT-related publications and four non MMT-related publications during this reporting period. No technical memoranda or reports were generated. See the listing of publications in Appendix I, p. 32.

Presentations and Conferences

G. Williams and J. Hinz attended the 224th American Astronomical Society meeting held June 1-5 in Boston, MA.

Six MMTO staff members (D. Blanco, D. Clark, J. Di Miceli, D. Gibson, R. Ortiz, and G. Williams) attended the SPIE Astronomical Telescopes and Instrumentation 2014 conference held June 22-27 in Montreal, Canada, resulting in the following paper and oral presentation:

Software framework for the upcoming MMT Observatory primary mirror re-aluminization

> *J. D. Gibson*, D. Clark, and D. Porter (italics denotes presenter)
> *Proc. SPIE*, **9152**, 91521K-14 (Software and Cyberinfrastructure for Astronomy III conference proceedings)

In addition, D. Blanco presented a poster paper entitled "Test and Metrology of Large Optics" at the SPIE conference. This work was independent of the MMTO.

D. Blanco participated in a design review held June 18-19 at the Stanford Linear Accelerator Center regarding the L3 dewar window and filters for LSST (Large Synoptic Survey Telescope).

Safety

In May, T. Gerl assumed the position of chief safety officer from C. Knop, who will continue as the campus office safety officer. As more safety issues are likely to occur on the mountain, it was decided that the chief safety officer should be located there.

Training

J. Di Miceli attended Wilderness First Responder training in May.

On June 17 several MMTO mountain staff attended a CPR/AED/First Aid refresher training offered by FLWO for Smithsonian and MMTO staff.

Safety Inspections

A Smithsonian METR (Management Evaluation and Technical Review) team visited FLWO and MMTO in early May. A METR safety inspection of FLWO and MMTO was conducted by OSHEM personnel on June 3-5.

Minor safety concerns were noted during the OSHEM METR safety inspection. All items were addressed soon after, including the purchase of safety equipment to aid in moving or removing injured people from the summit.

Procedures and Protocols

As a result of the OSHEM METR safety inspection, the following policies and procedures are being developed or revised (where noted): loss of power, fire shutdown (revised), lightning shutdown (revised), air closure of the front shutters, training requirements and periodicity, as well as chemical storage and handling (revised).

The new annual chemical inventory and MSDS/SDS binder review is near completion. Removal of unused or expired chemicals and acids from the summit was completed.

Miscellaneous

Labels listing emergency contact information were placed on all mountain phones.

Primary Mirror

The mirror was CO_2 cleaned on May 7, removing a layer of dust that had accumulated in the dust storms of the prior quarter. A hand-held dust monitor was ordered and will be delivered after the unit is modified for operation at 8500' elevation.

Problems with raising the mirror occurred on May 18 and again on June 7. The first incident was unusual in that the mirror overshot and ended up about 2 mm higher than its normal raised position. The second incident was traced to operator error; the operator started the mirror raise and then moved the telescope in elevation before the mirror was in position. This was unusual in that the primary mirror controller reported problems during the raise, then cleared and reported that the mirror was in position when it was not. The only indicator that something was amiss was that we could not acquire a star. Lowering and re-raising the mirror cleared the problem.

Coating & Aluminization

An Aluminization Readiness Review was held April 22-23. The external review panel was comprised of representatives from other observatories experienced with mirror coating systems similar to ours. The panel members were: John Hill (LBTO), Frank Perez (Magellan Telescope), Mark Klaene (Apache Obs.), Bill DeGroff (Lowell Obs.), and Byron Smith (formerly with Discovery Channel Telescope). The panel toured the MMT and aluminization equipment at the basecamp on the first day of the review and heard presentations by the MMTO staff. On the second day the panel met at Steward Observatory for continued presentations by staff and discussions, followed by a private session during which they wrote their report.

A webpage was created by the software group for the Aluminization Readiness Review to allow easy sharing of documents and presentations between the staff and the review panel members. The page was created with Twitter Bootstrap and remains online at http://www.mmto.org/aluminization/ for archival purposes. A screen capture of this custom documentation and presentation portal, listing the various topics of discussion, is shown below in Figure 1.

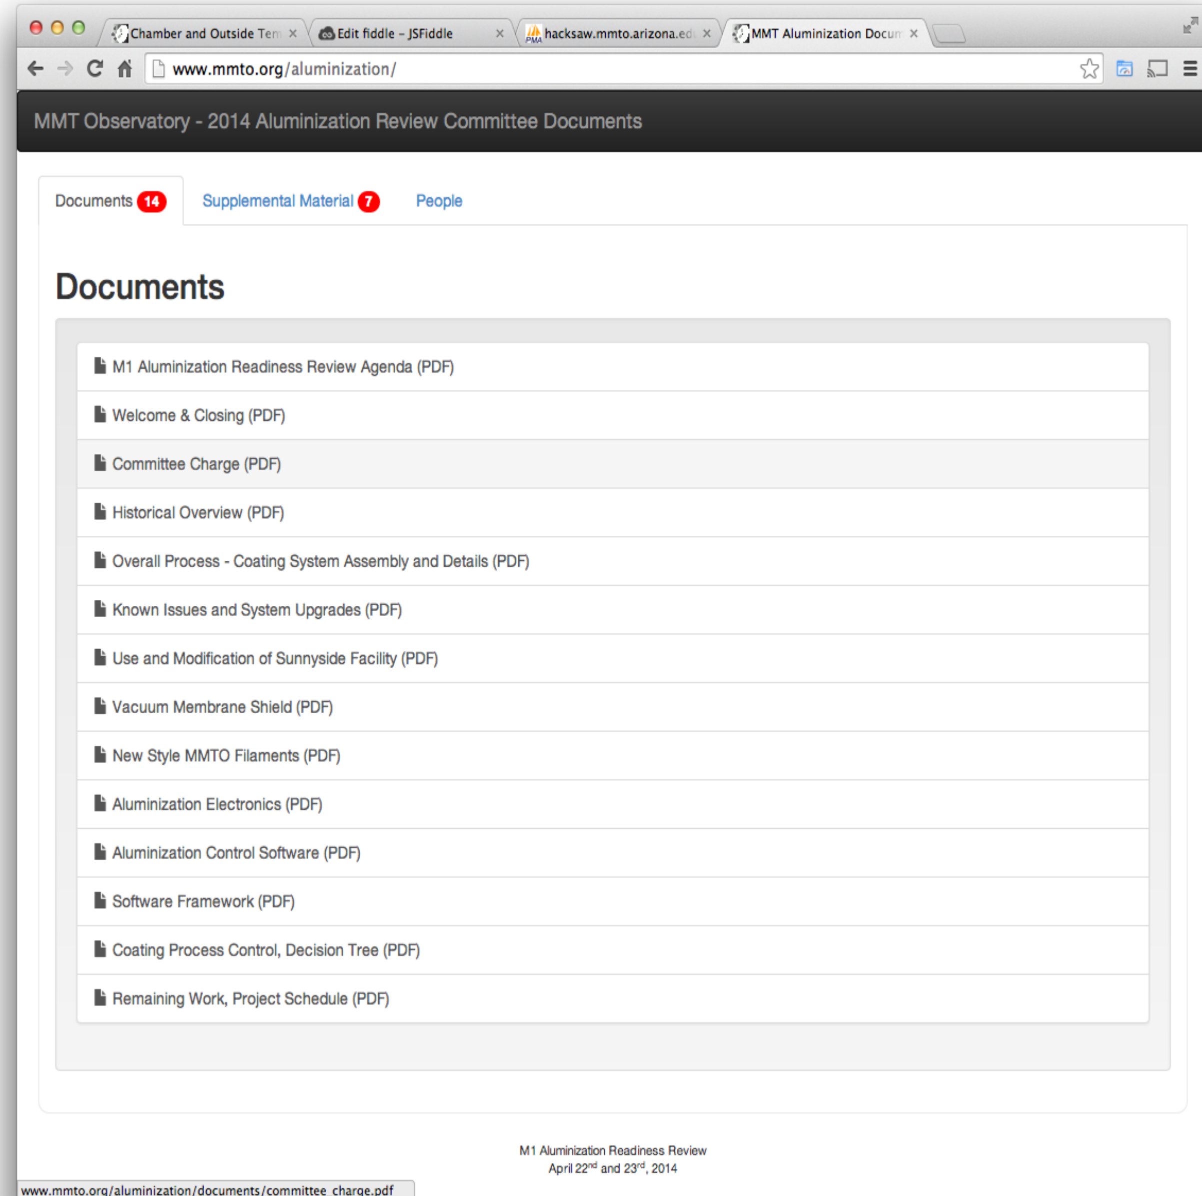

Figure 1. Web-based portal for the MMTO Aluminization Readiness Review. This portal allows easy access to aluminization documents, supplemental material, and contact information for individuals involved in the aluminization.

Figure 2. An informal working lunch during the Aluminization Readiness Review held April 22-23, 2014. Shown from left to right are D. Clark, D. Gibson (standing in background), and J. Hill, a reviewer.

The feedback from the review panel was overall positive. However, they had several recommendations for changes and processes to help in attaining a successful automated realuminization.

Due to these recommendations and the time needed to implement them and other changes, G. Williams proposed on May 23 (after discussions with senior staff) to postpone re-aluminization to 2015. The MMT Council agreed. Recent reflectivity measurements had also indicated that the coating hadn't degraded enough to require a new coating this year.

The new style filaments and automated control system testing was placed on hold for most of April in order to prepare for the Readiness Review. The review preparations provided the aluminization group the opportunity to take a critical look at the coating process and to thoroughly review the data gathered during the first round of testing at the Sunnyside vacuum facility. Thirteen presentations were developed that focused on the process changes implemented to improve the process reliability and to reduce the risk of blemishes during the next coating. The review panel provided numerous suggestions for improvement, and the entire review process helped illustrate advantages of postponing the aluminization until the summer of 2015. A response to the panel's report has been written and is currently waiting for internal review before being more widely circulated.

After the review was completed, another round of testing began at Sunnyside's small (18") and large (84") chambers. Functionality testing was performed after the two process control racks were merged into a single rack. Testing then focused on process optimization and understanding. In addition to a number of small chamber tests, three large chamber tests were completed: two using the new style filaments and one using the old style filaments. Some conclusions from this round of testing are:

- The location of the deposition monitor crystal in the large chamber was relocated and confirmed to estimate the MMTO bell jar reasonably well. For the May testing, 900 angstroms of detected coating thickness corresponded to 0.2 grams of aluminum evaporated from each filament. These numbers are similar to the values reported after the successful 2001 coating, but the thickness in the large chamber is slightly less (900 angstroms instead of 980 angstroms).
- In the present configuration, the new style filaments require at least 50% more power at the chamber feed-throughs than the old style filaments for comparable deposition rates. The large chamber test with the old style filaments used 6.6 kW per welder at the load, while the tests using the new style filaments used 10 kW per welder at the load.
- No drips were observed from the new style filaments in either the small or large chamber testing. Three to four drips were witnessed from the old style filaments during both of the small chamber tests, and two drips were detected during the single large chamber test.
- Based on before and after coating mass measurements, roughly 80% of the aluminum load (not including the aluminum wire under the filament clamps) can potentially be evaporated if sufficient power is applied to the filament.
- Although reusing filaments for a second 1000 angstrom coating was possible in the small chamber, both attempts to reuse filaments in the large chamber have been unsuccessful.

In an effort to reduce the power required to drive the new style filaments, small chamber testing was initiated with new style filaments that had been reduced in overall length. Unfortunately, due to manufacturing limitations, altering the overall length of the filament is the only feasible option for changing the amount of power required. Using the Sunnyside small chamber, new style filaments (overall length 42 inches) were modified and tested with the following overall lengths: 30, 21.75, and 15.5 inches. The shortest filament tested, 15.5 inches, is the minimum length of this style of filament that will deliver the amount of aluminum necessary for a 1000 angstrom coating in the MMTO bell jar. Early data analysis indicates the shorter new style filaments can be safely used in the bell jar but will still require 25% more power than the old style filaments to achieve similar deposition rates. Once the data on the shortened filaments are fully reduced, a batch of filaments will be ordered and tested in the large chamber. This should complete the major testing effort at Sunnyside.

Testing of the new mirror coating system electronics and software in both the large and small coating chambers at Sunnyside was conducted during this reporting period. The testing showed that a major repackaging of the aluminization electronics was required to consolidate the system cabling and clean up electrical issues with the connections of the many different kinds of equipment connected to the data-acquisition (DAQ) electronics.

Two racks, one holding the DAQ electronics and control system PC, and the other with all the vacuum monitoring equipment, were united in a single rack recovered from storage at the basecamp. The rear panel on this "new" rack was modified to swing out on a hinge for service, and to connect all the system sensor cables to the display electronics such that the field cabling can be connected to well-marked connectors on the rear panel without having to reach inside the rack. All the input signals to the rack DAQ boards were completely rewired using two terminal strips installed on the side of the rack, under the side covers. The terminal strips allow measurement of signals throughout the system without opening up the internal electronic chassis boxes.

The DAC (digital-to-analog) card and wiring was modified to add thickness relay closure inputs from the deposition crystal sensors and to clean up the digital signal inputs from the PC. The thickness relay closures allow for redundant shutdown of the coating process when the accumulated coating thickness reaches a setpoint entered into the deposition monitor boxes, along with other criteria the software can use to end the coating process. A newer version of the DAC board was designed to add a watchdog timer on the welder enable signal. If the controller stops updating the DAC, the welders will be disabled for safety.

A new deposition crystal meter was installed in the rack, resulting in a total of three deposition sensors for the mirror coating: two installed in the center area and one near the mirror edge. We also re-worked the isolation chassis cards and the new isolation panel used to provide electrical isolation between the DAQ PC and the vacuum equipment. Additional space on the rack is slated to hold the Austin Scientific cold cathode display panel, which is not yet completed. Once completed, it will also need to be verified against the cold cathode reading chart with the actual cryopump connections.

During work at Sunnyside, the turbopump controller on the small chamber failed. The 24V failure signal voltage ceased to be present when the controller was put into an obvious failure mode by disconnecting the pump from the controller. A burning smell was present and the system was quickly turned off. The manufacturer was contacted, and they verified that, while the mode we had put it in was an approved testing method, the smell was not normal. The broken controller was exchanged by the manufacturer for a new controller and it was received within 7 days. The new controller was tested and the failure signal now displays as it should on the new equipment rack.

We discovered that the standard way of enabling the welders at Sunnyside had the unfortunate consequence of causing the AC supply generator to drop out of regulation. The relay closure for the welder-enable signal is normally chained in series such that high and low welder pairs are offset in time by the few tens of milliseconds that their relay closures take to operate. When testing in the large chamber, the inrush current is too high for the generator to maintain regulation, so its output drops out and the supply power to the welders is lost. We used time-delay relay cards retro-fitted to the welder interface boxes to delay the second welder-enable by approximately 8 seconds, long enough to allow the generator to recover from the inrush of the first welder load going live. We can now consistently test the full large-chamber load without losing power.

The testing performed at Sunnyside allowed us to make considerable progress on production of fully automated deposition-control software. The deposition controller, written in C with both hand-coded and automatically-generated C from Simulink, has a flexible interface to other code written in Node.js that provides the necessary graphical user interface (GUI) and networking interfaces for user I/O and telemetry using Redis, a fast in-memory NoSQL data server.

Aluminization GUIs

Figure 3 below shows testing in the large chamber at the Sunnyside test facility. These tests have been critical in evaluating new hardware and software for the upcoming re-aluminization.

Figure 3. Large chamber testing at the Sunnyside facility. Shown from left to right are: R. Ortiz, G. Rosenbaum (in background), D. Clark, J.T. Williams, and J. Di Miceli. The aluminization software GUIs are shown in the lower left monitor. A close-up of the glowing filaments is shown on the monitor in the upper center of the photo. As the temperature increases, these filaments become white-hot. Control is automated during the tests. Staff mostly watches the monitors as the automated aluminization process proceeds, as the majority of the work is in preparation for the test.

Work continued on the aluminization GUIs. An "Aluminization chart overlay" GUI (Figure 4) was developed that allows two datasets to be overlain, via a drag-and-drop method, onto a single chart for rapid visual comparison. The page loads two different "shots" into two charts. The bottom chart can then be dragged over the upper chart. As the lower chart moves over the upper chart, the transparency changes so that both datasets can be seen. This makes it very easy to see the differences in the data. This GUI was implemented using JavaScript and the jQuery HighCharts JavaScript library.

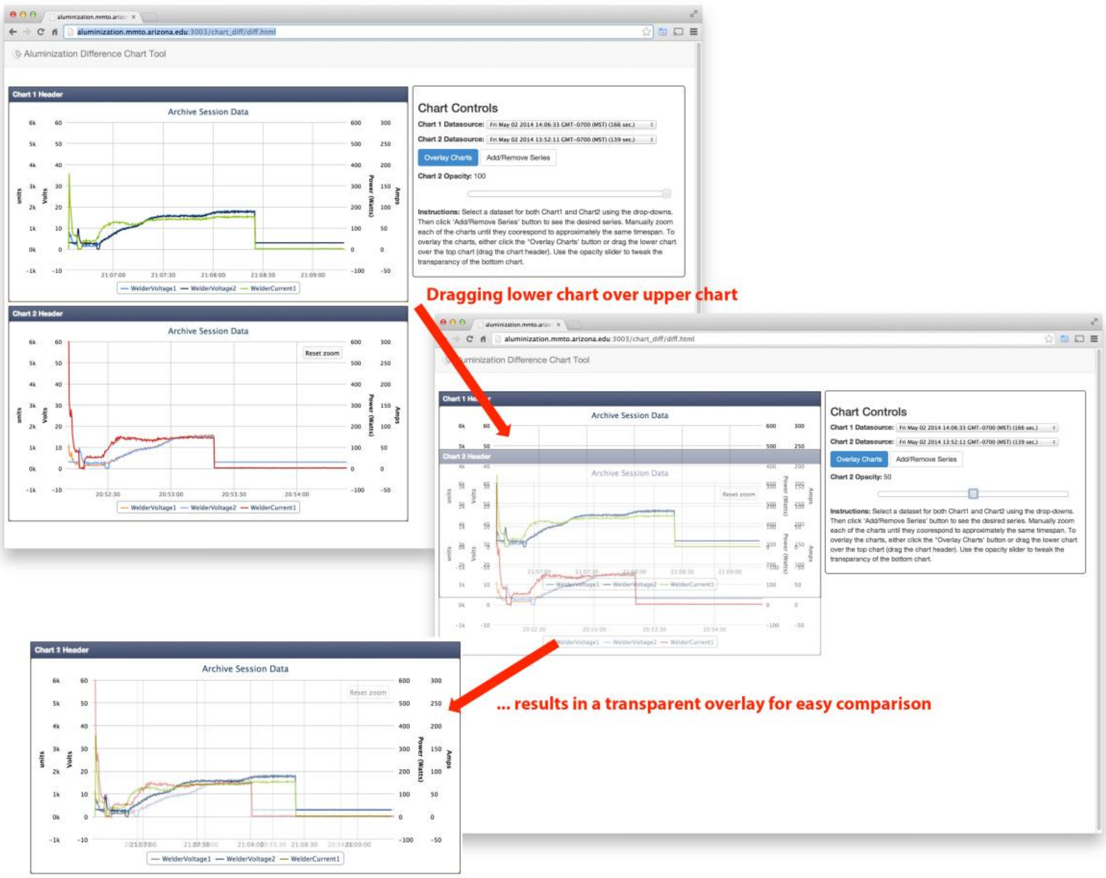

Figure 4. The "Aluminization chart overlay" plotting tool was developed that allows two datasets to be overlain, via a drag-and-drop method, onto a single chart for rapid visual comparison. The immediate availability of charted data after tests and the ability to compare results between different tests has been invaluable during development of hardware and software for re-aluminization.

A new "sessions" MySQL relational database, independent of the main "aluminization" logging MySQL database, was developed for defining distinct aluminization test sessions. The logging server begins a new session when the "start" button is pressed on the main aluminization GUI page. When the "stop" button is pressed, the logging server writes a new row into this "sessions" database. The sessions are used in the archive chart plot webpage to populate a drop-down list of available test sessions. When one of the sessions is selected, the session ID is sent back to the server and the server does a query on that session's ID, gets the start and end timestamps, then queries the log database and gets all points between the two timestamps. Depending on the total elapsed logged interval for a session and the current user settings, data are sparsely sampled so that no more than approximately 800 values are returned at any time. The archive chart allows zooming in to drill down into finer log detail. When the chart is zoomed, it sends the user's new chart start/end timestamps and another database query is performed, again such that no more than 800 values are returned.

Ventilation and Thermal Systems

Automatic purge valves were installed in the air line, located in the pit, that purges the compressed air supply every 45 minutes for the mirror support system. These valves will prevent water accumulation in the lowest points of the air system piping.

We are investigating methods to move the blower start and stop hardware to a WiFi-enabled relay. This should help prevent accidental auto starts (from RF interference) and remove non-safety related equipment from the 26V interlock system.

Actuators

Originally developed for use with the f/9 secondary mirror, a simple Python GUI currently exists to inject known wavefront aberrations into the primary mirror when testing changes in the wavefront sensing software. In June, a version of the software was developed to do the same mechanics with the f/5 configuration. This will be extremely valuable when testing off-axis wavefront sensing currently in development for MAESTRO and Binospec.

Secondary Mirrors

f/15

Work continues on the adaptive optics power supply. The front panel meter firmware was modified, and the meter now properly displays the supply voltages and currents. Minor adjustments to the code are being made to ensure that the display matches a known output. New connections were found for the connection between the power supply and the output connectors. A narrow terminal lug from Burndy, available from local distributors, works well.

Testing continues on a daily basis. Initial testing of the current sensor cards has had positive results. Mounting of the sensor card to the output cable is being evaluated, with large heat shrink tubing being the leading option.

Hexapods

Elcol measurements were taken for both f/9 and f/5 configurations during Maintenance and Engineering (M&E) nights. The slopes used to determine the relative elcol corrections as a function of elevation have remained stable on multiple year timescales.

The hexapod logs were searched to aid in understanding startup problems being experienced. These problems are still being investigated.

Optics Support Structure (OSS)

Nothing to report.

Pointing and Tracking

An analysis of the pointing model for the MMT was begun. We are specifically examining the long-baseline variations in the pointing model, striving to find and understand any residuals between the pointing data and the pointing model in use at the telescope. We have verified the setup of Tpoint software, have begun a complete documentation of the methodology of using the software, and have begun to look at the variations between pointing runs. There has been no indication of strong pointing errors at the telescope, but having detailed documentation on how to recreate the pointing model will aid in the event of future changes to the system.

Science Instruments

f/9 Instrumentation

The prototype for the Blue and Red Channel exposure time calculator has been completed in Python. We are actively developing the PHP/HTML front end for the software users. Once this frontline is completed, we will enter an alpha mode of testing the code.

We began creating updated throughput plots for users. These plots compare the measured counts/s/pixel of a standard star to the standard flux values. Several users have indicated they would like these types of plots in addition to an exposure time calculator.

The f/9 instruments were on the MMT for 43% of the available nights from April 1 through June 30. Approximately 66% of those nights were scheduled with the Blue Channel Spectrograph, 10% with Red Channel, and 10% with SPOL. M&E was scheduled for 13% of the time. Of the total 338.1 hours allocated for f/9 observations, 63.7 hours (19%) were lost to weather conditions. Instrument, facility, and telescope problems accounted for less than 1% of lost time. Blue Channel lost 19% of its time to poor weather, with Red Channel losing 12%, and SPOL losing 29%.

Although there is an issue with the chip in the Red Channel detector, M. Lesser at ITL was able to tweak it adequately so that it could be used for its scheduled four-night run in May.

f/5 Instrumentation

In terms of hours for SAO f/5 instruments, we operated for 346 of the 412 hours scheduled. Approximately 16% of the scheduled time was lost this quarter, mostly due to weather. A cascade of issues caused an early termination of the Hecto positioner run in early April. The instrument did not operate for five of the scheduled nights and the time was transferred to Blue Channel. Issues included a spurious fault error from the R1 X amplifier after one fiber move, and a failure to close on the part of the R2 gripper mechanism. No root cause for these errors was found. More critically, two R2 P axis following errors occurred, which we attributed to the actuator mechanism starting to fail. Non-instrument losses of observing time included a little over half a night on June 1 when the generator fuel pump relay failed. We also lost a short time to telescope oscillations on April 13 and to an improper mirror raise on June 7.

With the above Hecto positioner instrument issues, the service mission that usually takes place in July or August was moved to early May. During the mission, the gimbal encoder voltage levels were noted and later readjusted to the currently specified target values. After separating the positioner, axes were homed and we found that the P2 home position, which had been erratic for over a year, had changed significantly to nearly the same value we had established last summer. The primary change in the situation was the use of the long laboratory cables and use of the 5-foot test cables. We suspected there was an issue with at least one signal on either the drive arc cables or the 25-foot instrument cables. We were not able to find a problem in the initial cable tests. The actuator that had the folioing error was replaced, and its mounting nut was machined to have the home signal in the correct location. The servo parameters were adjusted for the new actuator and the x and y axis homes were adjusted to match the coordinate system of the other robot. The 25-foot cable was rechecked and one of the P2 home signal lines was found to be shorted to the shield. This cable and three other similar cables were replaced. The fiber chain was inspected and the conditions of fibers at the transition box were noted. Accumulated debris (leaves and bugs) was cleaned from the edge of the Wide Field Corrector. When the instrument was mounted on the telescope, the T1 axis failed to stabilize at the end of a couple of configurations. Its servo parameters were adjusted a couple of days into the run.

Hectochelle and Hectospec gathered data on 40 nights this quarter. A total of 564 object exposures were taken of 157 fields on these nights. An additional 1,960 calibration exposures (bias, comp, flat, dark, and sky) were also obtained. Due to weather issues, there were two nights on which no object exposures were obtained and an additional two nights where data were obtained on only a single field.

This quarter saw the start of split nights between spec and chelle. This allowed us to more efficiently use the allocated nights because one instrument had early targets and the other had later targets. The split nights mean that the operators are gathering additional calibration exposures each day. Also, there is a small delay at the time of the camera switch due to moving the shoe and to the camera reaching optimum temperature.

SWIRC data were gathered on eight nights this quarter, one of which only resulted in dark and sky images due to clouds. A total of 5,141 object images were taken along with 1,317 calibration images.

MMTCam gathered data on 16 nights this quarter, with 475 object images taken as well as 824 calibration and engineering images (bias, flat, dark and focus). Some of the calibration data were taken on a half dozen nights when the camera did not take object images, since darks and biases can be taken on cloudy nights. Note that the MMTCam is usually scheduled for a short time during the night with the rest of the time allocated to Hecto.

The WaveFront Sensor (WFS) operated relatively well this quarter, although there were a couple of issues at the beginning where axes failed to complete a move. However, the operator was able to recover. The cause of the problem was not determined, although the unit was exercised in the lab after the run. The problem has not recurred. The homing speed of the focus axis has been slowed for enhanced reliability, and the predicted time has been adjusted in the code to reflect the new homing speed.

One of the dome calibration lamp boxes stopped responding late in the quarter. The battery in an NVRam chip appears to have died after a decade. Replacement chips have been ordered, and a work-around allowed us to get the unit back online at the end of the run.

MAESTRO was originally scheduled for three nights in mid-April. However, issues were reported and the nights were instead given to M&E with the f/9 secondary mirror.

f/15 Instrumentation

There were no Natural Guide Star (NGS) adaptive optics observing runs during this reporting period.

An effort to diagnose WFS camera issues was begun. Cables and electronic cards were tested in an attempt to understand whether any hardware components were contributing to the ongoing WFS camera freezing issues.

The current plan is to replace the nearly 80-foot long SCSI cable with a fiber optic cable. A new Linux machine was created to test the EDT fiber converter card that will reside in the SciMeasure camera electronics box in the AO topbox. A PCI framegrabber card will then be inserted into the current PCR to read in the image data. Work on the fiber converter is ongoing.

Topboxes and Wavefront Sensors (WFS)

f/9 Topbox

Testing of the new f/9 WFS SBIG camera began in April and continued in May. The camera operates properly in the f/9 topbox, and software to expose and readout the camera was developed and tested successfully. However, the software that derives the wavefront errors, based on the spot pattern observed with the new camera, does not converge to remove the wavefront error. Many calibration images with known input aberrations were taken to aid with debugging the wavefront correction code off-sky.

Natural Guide Star (NGS) Topbox

Maintenance was performed on the NGS topbox to clean up and repair various electrical connections and optical components, making the system more reliable as well as easier to service.

Facilities

Main Enclosure

The fluorescent light fixture was replaced in the yoke room. New emergency lights were installed in various locations throughout the building.

General Infrastructure

A systematic investigation of the compressed air system was started in June and has initially focused on documenting its present condition. As part of this assessment, a hygrometer has been ordered to help determine the effectiveness of the current air treatment equipment. Once this is complete, the hygrometer will be permanently mounted in the system to provide continuous monitoring of the compressed air dew point.

Dave Shambach, a Smithsonian Institution (SI)-contracted architect, is continuing work on developing a design proposal to add a new standing seam metal roof and snow melt system.

Firetrol completed installation of the new fire alarm system and we await final check and approval by SI. Keltron, who makes a box that monitors the alarm system, delivered a second replacement for the first computer. However, it has exhibited some problems during boot-up. Inspection and final acceptance of the system has been delayed to late July.

GLHN, a local engineering firm, has completed a set of drawings documenting the summit HVAC systems. They are also preparing a construction bid package for some modifications to the aging HVAC systems that can be made relatively quickly. These include replacing pumps, adding balancing valves and air-bleeders, and installing air conditioning for the control room.

During this quarter our SI-funded contractors temporarily connected the second chiller (the "Carrier on the Rock"), confirming the unit is operational. This unit was acquired in 2003, placed on its present location on a rock east of the IRF at the summit, and mothballed. This fall we plan to put this unit into service as the primary chiller for the observatory, and retain the current Carrier chiller unit as a back-up. The Carrier chiller is the release point for about 80% of the waste heat generated by the observatory during nighttime operation. Switching over to the Carrier on the rock will move this waste heat about 25 meters further from the telescope and should result in a slight improvement in the seeing at the site.

A long range study is underway to investigate the use of an ice house. In this scenario we would run the chillers during the daytime to freeze a mass of water/glycol mixture tailored to freeze at a sufficiently low temperature. We would circulate glycol through the ice to provide cooling for nighttime operations. This would eliminate most of the heat generated during operations.

Access to the mountain was intermittent through May and June due to work on buried power lines. This included re-cabling a portion near kilometer 8. The ridge generator supplied power during power interruptions, but on Sunday, June 1, we lost all power when an automatic fuel switch at the generator failed to switch to a new fuel tank. FLWO crewmen worked late into the night and power was restored a little past midnight. In all, 4.5 hours was lost to this power outage.

Computers and Information Technology

One of the ongoing challenges within the MMTO is managing the growth of computing and IT requirements. A decade ago a single computer, *hacksaw*, served as the telescope operator's (TO) workstation in the MMT control room and largely met the telescope observing and operations needs at the summit. This computer was also the sole summit-based web and database server, hosting only a few web pages and small databases. Observer's needs were met by a second workstation, *alewife*, also located in the control room.

Approximately eight years ago, in July 2006, a headless server located in the MMT 3E server room became the new *hacksaw*. A client-server computer architecture was adopted with *hacksaw* being the server, and two new Linux computers, *yggdrasil* and *hoseclamp*, becoming the client computers used by the TO. *Alewife* remained the observer's computer.

More recently, *hacksaw* has evolved into a virtual machine (VM), running on one of two available VM host computers (*vmhost1* and *vmhost2*), both located in the 3E server room. MMT software and data, primarily located in the "/mmt" directory, have been migrated to one of two possible network-attached storage (NAS) devices (nas1 and nas2). Y*ggdrasil* and *hoseclamp* have been replaced by more current hardware and revised names: *chisel* and *pipewrench*. *Alewife* has been replaced by an iMac, *pixel*, and a second iMac backup, *gilead*.

All of these computers, *hacksaw, chisel, pipewrench*, and others, currently share the same NAS box for the crucial /mmt directory. This shared resource, the network file system (NFS)-mounted /mmt directory on a NAS box, has made excessive activity on one computer (particularly extensive disk I/O within the /mmt directory) affect all other computers using that shared resource. For example, MySQL queries by MMT staff, which access the database files stored within the /mmt directory, can noticeably slow down the user interfaces for the telescope operator on *chisel* or *pipewrench*. Similarly, transferring files via "rsync" or listing files within the /mmt directory can affect all other computers. This interdependency has been the topic of many conversations within the software group and with staff experiencing computer sluggishness.

The FY15 MMTO Program Plan being written this quarter contains several high-level tasks involving the separation of software within the /mmt directory into two or more categories, and the migration of substantial bodies of software to one or more new computers (e.g., a new MySQL database computer). This is a natural evolution as the software/IT framework at the MMTO continues to grow in scope and size. Much of the work during this reporting quarter relates to separating these two types of software, both of which are large bodies of software at the MMTO, and isolating control systems software from operations software.

MMTO software has been informally divided in two categories, referred to variously as critical and non-critical, observing and non-observing, or control system and operations software. Interactions with staff from other observatories at the bi-annual SPIE Software and Cyberinfrastructure for Astronomy conferences suggest that the control system and operations software categories are the most commonly used worldwide. In detail, it can be difficult to determine in which category a specific program or body of code belongs, and the decision can vary between individuals.

In general, control system software is software that is directly interacting with and controlling hardware, primarily the telescope mount, primary mirror, and secondary mirror. It can also include software used to control instrumentation or other equipment that is required with observing. Control system software typically contains some form of controller or servo logic for hardware control and may be implemented with a real-time operating system.

Operations software is everything else. In SPIE terminology, this software forms the "cyberinfrastructure." This category of software includes telemetry, data logging/ analysis/visualization (i.e., data pipelines), virtualization, and all other software involved in observatory operations. Directly related to this operations software at the MMTO are the extensive databases of images and telemetry data involved with logging. These data storage requirements are currently in the hundreds of gigabyte range for the MMTO and are growing daily.

In the control systems/operations software paradigm, it is critical that control systems software be isolated from operations software. Any operations software activity, such as data reduction and visualization, should not affect control software in any way. Control software must be sufficiently isolated from operations software so that the two sets of software do not interfere with each other. Interference could result from shared resources such as computers, network-attached storage (NAS) devices, network file systems (NFS), and network communication hardware. As an example of possible interference between the two types of software, plotting of telemetry data by a staff member should be completely transparent to the telescope operator and observers.

Computers and Storage

A new computer, *ops.mmto.org*, was purchased to aid in the separation of control system and operations software. Potential operations software roles for this computer include image processing (skycam and webcam images), telemetry logging and visualization (the miniserver MySQL background logs), and data analysis and visualization (web-based charting tools). Details of the role of this computer and how it integrates into the computing and IT environment will mature with time. At the present time, this computer does not NFS-mount the /mmt/ directory from *nas2*, and user access is restricted to local administrative accounts. It is being configured as a typical "LAMP" (Linux Apache MySQL PHP/Perl/Python) server.

Computer and Storage Administration

Linux

The newly released CentOS 7, an open-source production-level Linux distribution by Red Hat, has been installed on *ops*. Red Hat uses the Fedora Linux distribution, commonly used at the MMT, as its development-level Linux distribution, RHEL ("Red Hat Enterprise Linux") as its production-level commercial Linux distribution, and CentOS ("Community Enterprise Operation System") as its open-source production-level Linux distribution. Traditionally, CentOS has been the closest to RHEL of the various open-source RHEL derivatives. Red Hat will support RHEL/CentOS 7 for over ten years, with the End of Production Phase 3 on June 30, 2024. The Linux kernel will remain at version 3.10 throughout the life of RHEL/CentOS 7, helping to ensure stability in a production environment.

The *f9wfs* computer had "prelink" problems that were corrected. Prelink is a Linux program that is intended to speed up a system by reducing the time a program needs to begin. This machine is several years old and is due to be retired.

Planning and testing of a new method for the compaction of *hacksaw's* virtual disk image (vdi) file was done. This will be implemented over summer shutdown. This method should make future vdi file compaction unnecessary and will speed up backups, reducing *hacksaw's* scheduled down time.

The four Linux-based NAS boxes, *nas1, nas2, nas3* and *ao-nas*, were updated to the current DSM 5.0-4493 release. The current versions of Redis and node.js were installed on *nas1, nas2,* and *nas3*.

Mac

Mac OS-based systems: *pixel, gilead, telstat, telstat2,* and *pluto*, were also kept current with OS updates. These computers run Mac OS 10.9 ("Mavericks"). XCode and MacPorts software was also updated on these computers.

The Mac mini, *telstat2*, has been used during aluminization testing at Sunnyside. This computer and two 40" TV/HDMI monitors have been moved to the mountain in preparation for summer shutdown and installation in the control room. The computer will be reconfigured for the mountain subnet. The existing "telstat" Mac mini and its two 40" TV/HDMI monitors will provide telescope status (telstat) information to the TO, while *telstat2* and its two monitors will provide status information tailored for the observer.

The location of these and other computers in the control room will depend on the control room redesign during summer shutdown. New control room desks were specified and ordered from Goodmans in Tucson, and delivery is expected in early August. Also researched and ordered for the control room were new recessed lighting, sound absorbing panels, and TV mounts for the telstat displays.

Network

Hardware

Two dedicated network address translation (NAT) gateway boxes were purchased and configured to handle the traffic between MMT's private and public networks. One was installed at the mountain and one at the campus office. These gateways were installed to handle gigabit-rate NAT translation between the private and public subnets, which has caused excessive loads on the *hacksaw* and *mmt* VMs during high data transfer rates. These NAT gateways relieve the mountain and campus servers, *hacksaw* and *mmt*, respectively, of that burden.

The four Lantronix Spiders used at the MMT provide secure KVM (keyboard, video, mouse) server management over an IP network. These devices allow remote access to the consoles of our four virtual machine hosts (*vmhost1* through *vmhost4*). Outdated firmware and settings had caused these Spiders to stop working. The firmware for the Spiders was updated to the current version. Keyboard/mouse configuration settings were changed to "Force USB Full Speed Mode" to allow proper keyboard and mouse behavior over the Spider KVM connections.

Administration

Work continued on an ongoing issue of "broken pipes" on long data transfers or network connections. The MMT worked with other organizations within Whipple Observatory (*e.g.*, VERITAS and the Ridge telescopes), Steward Observatory (*e.g.*, Large Binocular Telescope), the Smithsonian Institution/Washington (*i.e.*, Jun Wu), and the University of Arizona University Information Technology Services (UITS) on this issue. Current status on the issue is documented in both a UITS 24/7 trouble ticket Incident INC000000265136, "Long rsync's have broken pipe errors," and in emails between numerous individuals. Network connections were terminating with "broken pipe" errors after variable time periods, between a few minutes to a few hours. This issue has been observed between the mountain and campus subnets as well as from MMT subnets to off-campus locations. The issue was preventing the periodic transfer of new "/mmt" data from the mountain *nas2* to the campus *nas3* computers.

During this work, it was found that traffic between the mountain and campus subnets was traversing the Steward Observatory firewall twice, potentially contributing to these data transfer issues. Modifications were made by UITS so that both subnets are behind the same firewall and that no firewall restrictions exist between them.

MMT staff coordinated with UITS staff, particularly K. Newell and R. Watkins, to address these problems. An MMT computer was temporarily set up on the MMT's campus subnet (199.104.150/24) for UITS's use in testing network connections. Modifications by UITS to firewall and network settings have improved network reliability. The number of broken pipe connections has decreased substantially. Work continues on testing and monitoring these network connections for reliability.

Network reconfiguration was done on a number of wired machines and network devices on the mountain, including cell phone relays, to remove the use of DHCP (Dynamic Host Configuration Protocol) for these devices. With DHCP, computers request Internet Protocol (IP) addresses and networking parameters automatically from a DHCP server. Static IP addresses were assigned to these machines. Requests for DHCP leases have commonly exceeded the number available, particularly at the summit. This has typically occurred when a large number of users are present at the summit. The Cisco network switches are now performing DHCP for wireless devices at the summit.

Internal firewalls on our Fedora-based Linux machines were upgraded. These firewalls supplement network firewalls maintained by UITS and Steward Observatory between the MMT subnets and the rest of the world.

A software tool was developed to query the Cisco network switches to easily determine to which virtual local area network (VLAN) each switch port is assigned. This tool is used to aid in migration and monitoring of network devices, including VOIP phones, between local VLANs at the MMT.

Hardware/Software Interfaces

The current wavefront sensor software is an amalgamation of software written in Perl, Python, C, C++, TCL, Ruby, and even FORTRAN. This wide assortment of software types makes modifications to the code troublesome (e.g., constants, like the pixel scale of the various WFS cameras, are defined in multiple locations). In order to make the system more easily expandable (to new cameras such as the f/9 SBIG, future instruments such as Binospec and MMIRS, and new modes such as off-axis WFS, we have started to convert the entire system to a Python code base (with some select calculations farmed out to very efficient C code). This is an on-going project, but a substantial amount of work has been completed in prototyping existing TCL and Ruby GUIs into Glade3.

MMT Control and Operations GUIs

Work continues on the telescope status (telstat) displays for the control room. The control room will have a total of four HD TV/HDMI monitors to display telstat information. New content for these monitors has been under development for the four displays. The goal of the redesign is to take advantage of the new display size, simplify the layout, and create a uniform style. This will be accomplished using shared CSS classes, styled Highcharts, and shared server-side code running on the new *ops* server. It is planned that the four telstat monitors will be in place in the control room after summer shutdown.

New work related to the telstat displays includes the development of a node.js server, *param-server*, which allows event-driven programming for the telstat displays. This node.js server is a "Redis/socket.io" data service for any of the 2000+ data parameters available from the MMT miniservers. It also allows complete access to the MySQL miniserver background logs through an application programming interface (API). All interactions with this data service, the Redis and MySQL/MariaDB servers, and web browser clients are asynchronous.

This server acts as a relay between Redis events from miniservers and socket.io events to web browser clients. Although the *param-server* node server can run on any computer, it is most reliable when it is running on the same computer as a Redis server (see Figure 5). As described elsewhere, Redis servers are currently running on *nas1, nas3,* and *ops* computers. The *param-server* data service is running as a system service on the *ops* and *nas3* computers. The data server is launched whenever the computer boots and remains running while the host computer is in operation. The data server handles each client connection separately so that, if no client is connected, there is no system load on the *param-server* data service. This Redis slave/*param-server* data service is very lightweight, and is able to run even on a microcomputer such as an $89 Beaglebone Black embedded Linux-based computer. A Beaglebone has been provisionally configured in this manner and is running on a public-accessible port 80, allowing public access to this data service from web pages served from other MMT web servers.

Among other software improvements for the *param-server* data service, the manner in which data are sparsely sampled for web browser charts has been greatly improved in efficiency, up to 100 times more efficient than in previous implementations. This new approach makes use of the "id" field in the miniservers MySQL background log tables in much the same way as the build-in "rowid" index is used in SQLite and the "rownum" index in PostgreSQL. MySQL does not natively support this type of automatic row indexing. For the InnoDB versions of the MySQL/MariaDB databases (see below for separate discussion on MySQL databases), the "id" field is now re-indexed so that rows are indexed sequentially based upon ascending timestamps. This allows very rapid access to sparely spaced data without use of computationally intensive MySQL count() or mod() functions previously used for time-series plotting. This approach is implemented in node.js for the *param-server* and has also been implemented in PHP for the xy8flot background plots on *hacksaw, ops,* and *nas3*.

The cell graphical user interfaces (GUIs) were moved outside of a browser context. These Java-based GUIs are now running as stand-alone Java applications rather than web-based Java applets.

A bug in the seeing server that prevented it from sending its data was fixed.

Telemetry, Logging, and Database Management

Several new technologies are included in the aluminization software framework being developed for the upcoming primary mirror re-aluminization. One of these new technologies is Redis, a key-value store and data structure server. Redis also supports master-slave replication, which was implemented on various MMT computers during this reporting period.

Performance of different database systems (*e.g.*, MySQL and SQLite), database structures and indexing approaches (*e.g.*, "spreadsheet" vs "eventlog" formats), and MySQL storage engines (*e.g.*, MyISAM and InnoDB) were also evaluated. A summary of ongoing changes to database systems at the MMT is included below. These changes included several computers, many of which are configured for master-slave replication.

Database replication can be used on many database systems, including Redis and MySQL, usually with a master/slave relationship between the original and the copies. The master logs the updates, which then ripple through to the slaves. The slave outputs a message stating that it has received the update successfully, thus allowing the sending (and potentially re-sending until successfully applied) of subsequent updates. New data are written only to the master. Database clients read data only from the slaves. This prevents large, complex data queries on slave databases from affecting the master database. The complex queries are handled, as needed, with a higher priority than replication. When the query is complete, replication continues automatically from where it paused.

A Redis database master has been running on *hacksaw* for several weeks (see Figure 5). New data are pushed to this master from each of the miniservers, using both "set" and "publish" Redis protocols. During this reporting period, Redis slaves were set up on two NAS boxes (*nas1* and *nas3*) and on the new operations computer, *ops*. Redis slave setup is extremely simple, involving modification of only one line in the Redis configuration file to specify the Redis master host and port details. The purpose of using Redis slaves is to isolate the Redis master, running on *hacksaw*, from excessive user loads.

Data for over 2000 parameters are available on the Redis master-slave system. Data are stored as JSON (JavaScript Object Notation) strings that include the following data for each parameter: key, value, datetime, and timestamp. The key value within the JSON string is somewhat redundant since the name of the Redis channel/parameter is the same as the key value stored in the string. Having the key within the JSON string makes processing of data much easier for clients. The timestamp parameter in the JSON string contains the UNIX timestamp in milliseconds. This is a common time format for time-series plotting tools such as the JavaScript Highcharts and Flot chart libraries. The datetime value is somewhat redundant with the timestamp, but makes the JSON string much more human-readable.

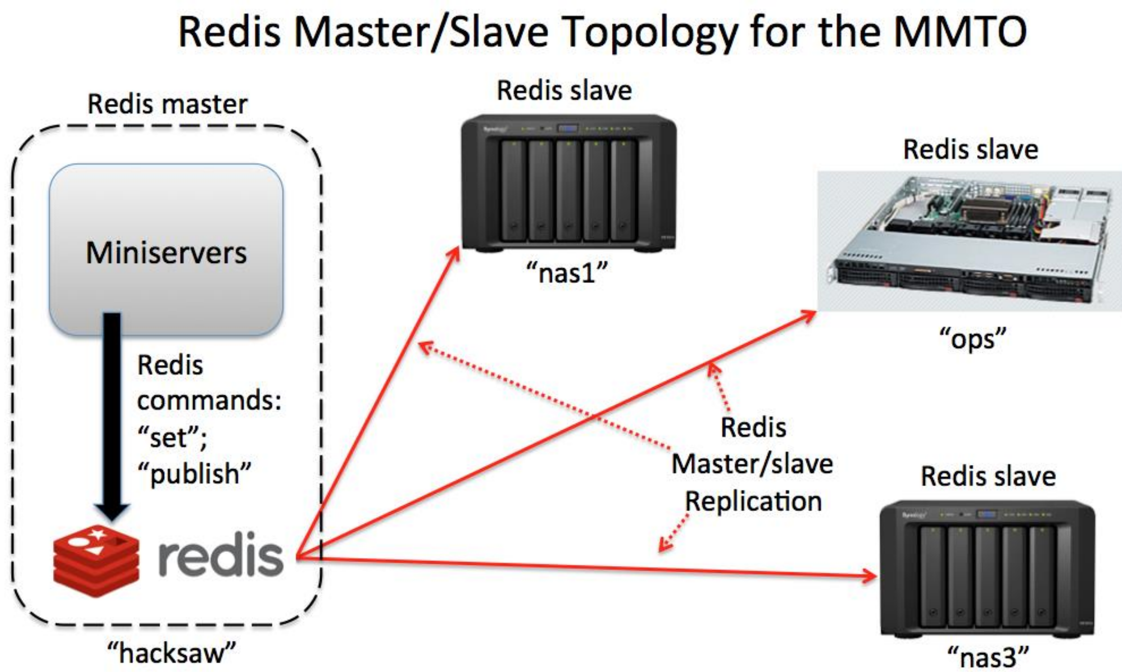

Figure 5. Redis master-slave topology at the MMTO. The Redis master runs on *hacksaw*. Redis slaves are currently running on *nas1, nas3,* and *ops*. Data parameters are sent from the miniserver to the Redis master via "set" and "publish" messages. These messages are only sent when the parameter value changes. This allows event-based programming throughout the telemetry system, including the web browser client. Data are automatically sent from the Redis master to all Redis slaves. Data queries are only made to Redis slaves, isolating the Redis master from user interactions.

Much work was done on database evaluation and implementation. One of the biggest impacts on operations has been database-intensive SQL queries on the *hacksaw/nas2* system. The MySQL files on *hacksaw* are NFS-mounted from *nas2*. For SQL queries that access large portions of database tables or indexes, SQL queries have been slow, up to several minutes to complete. One of these SQL queries can completely immobilize a computer that is NFS-mounted from *nas2*, including *hacksaw*, *chisel*, and *pipewrench*.

Two major changes to the database system were considered: 1) migration of the current MySQL database system to a simpler, potentially more robust non-MySQL database, specifically SQLite; and 2) implementation of the miniserver-related database system with a different table structure or schema, specifically, the "eventlog"-like table structure used for aluminization rather than the current "spreadsheet"-like table structure for the miniserver background logs.

SQLite is a software library rather than a database system or server. It is implemented as a self-contained SQL database engine. A database is implemented as a single file that contains the data, indexes, views, triggers, and any other information related to that database. SQLite is the most widely used SQL database engine in the world. SQLite is an application database that runs in-process rather than as a separate network-based server, such as MySQL or PostgreSQL, two widely used database servers.

SQLite implements atomic commit and rollback using a rollback journal or write-ahead log (WAL) file. There are several advantages to WAL over the rollback journal:

1. WAL is significantly faster in most scenarios.
2. WAL provides more concurrency, as readers do not block writers and a writer does not block readers. Reading and writing can proceed concurrently.
3. Disk I/O operations tends to be more sequential using WAL.
4. WAL uses significantly fewer file synchronization "fsync()" operations and is thus less vulnerable to problems on systems where the fsync() system call is broken.

Another advantage to SQLite over MySQL, at least for database use at the MMT, is the built-in "rowid" parameter for SQLite databases as previously mentioned. This parameter internally keeps track of the row id within a database table. In MySQL, a separate "id" field must be defined for similar functionality. During a VACUUM SQLite operation, the database is defragmented and the rowid is updated automatically. A similar update in MySQL can be done through a series of SQL commands – not quite as easily. These row ids typically do not change for miniserver data and are used to greatly speed up sparse sampling of data. A program for performing the auto-incrementing of an "id" field in MySQL was written as part of this work.

There are disadvantages to SQLite in general, and to WAL for SQLite in particular. WAL, by default, uses shared memory between the database file and its associated WAL file. This requires that the database and WAL file be on the same computer. NFS mounting of database files is not allowed. This would require that the miniservers run on the same computer as the SQLite database. Since SQLite WAL databases cannot use NFS-mounted files, the miniservers would need to be moved from *hacksaw* to another computer, the same computer where the SQLite databases are located.

Within the MMT miniserver framework, each miniserver is the only application that writes data to the database table associated with that miniserver. Using WAL option for SQLite, there is no limit to the number of readers that can access a database while the miniserver is writing to that database. This differs from the MySQL MyISAM storage engine, which is being used for the mmtlogs database on *hacksaw*. MyISAM locks at the table level. Only one client can access a database table at a time.

After many different tests using a variety of database structures and sizes, SQLite was found to be at least as fast as MySQL. It is extremely easy to configure and to use. It is supported by all of the major programming languages. However, SQLite is typically used for databases smaller than the mmtlogs background logs database. The upper limit for SQLite is huge, currently listed at 140 terabytes (TB) in the SQLite documentation. The mmtlogs database is currently around 150 gigabytes (GB) (1000 GB = 1 TB). During SQLite testing, the mmtlog database was divided into smaller databases, either with one database per one mmtlogs table or one database per one parameter.

Another disadvantage of SQLite is the generally lower quality of database administration tools compared to MySQL, such as "MySQL Workbench" and "phpMyAdmin." Although the same information can be obtained from command line prompts, these MySQL tools are invaluable for routine and detailed monitoring of database status and health. Finally, the MMT staff is familiar with MySQL. MySQL use is extremely widespread, and finding new database administration support staff would be easier than for less commonly used databases.

There have been several recent changes in the MySQL software community. Many Linux distributions, including Fedora, RedHat, Centos, and Synology's Disk Station Manager (DSM), have abandoned Oracle MySQL and have moved to another variation of MySQL, MariaDB. This is apparently related to the recent acquisition of MySQL by Oracle and the fear that it would not be available freely and open-source. Oracle MySQL and MariaDB are very similar, but typically do require some modification to database configuration files during the migration. As compared to SQLite, configuring and "tuning" a MySQL or MariaDB server for the data within the database and the hardware on which the server is running is much more complex, with hundreds of possible database configuration parameters.

There are many storage engines ("backends") available for MySQL databases, such as MyISAM, InnoDB, BLACKHOLE, CSV, MEMORY, ARCHIVE, and MERGE. SQL queries to the MySQL database are largely independent of the storage engine – the same SQL "insert" command would work for all storage engines. These different storage engines differ in the file structure and location for the database content on the database server. For example, MyISAM has separate data and index files for each database. MyISAM has been the default database storage engine for MySQL for years. In contrast, the InnoDB storage engine combines much of database data and indexing into one large file, typically called "ibdata1". This single file will be the size of all the databases on the server, commonly 10's to 100's of GB or more.

With the release of MySQL version 5.5, the InnoDB storage engine has become the default MySQL storage engine. MariaDB's version of the InnoDB storage engine is called the "XtraDB" storage engine, which the MariaDB Foundation considers to be an enhanced, highly compatible version of the InnoDB storage engine.

There are several major advantages of the InnoDB/XtraDB storage engine over the MyISAM storage engine.

- InnoDB/XtraDB is a transaction-safe ("ACID", i.e., "Atomicity, Consistency, Isolation, Durability" compliant) storage engine for MySQL that has commit, rollback, and crash-recovery capabilities to protect user data. This offers a higher level of data integrity than the MyISAM storage engine. We have had to disable re-indexing of MyISAM databases on *hacksaw* because of the minutes, to potentially hours, to re-index corrupted MyISAM tables during crash-recovery.
- InnoDB/XtraDB-based databases also support SQL transactions. These SQL transactions allow multiple database options and SQL operations to be handled as a single unit of work. For example, hundreds to thousands of rows of data can be inserted into a database in a single transaction. Transactions have been used with great success in the MMT re-aluminization database where data logging rates have been up to 100Hz.

The MyISAM storage engine offers more tools for database recovery. MyISAM databases can be recovered with the "CHECK TABLE", "REPAIR TABLE", and similar SQL operations. With the InnoDB storage engine, the only data recovery for a corrupted database is to dump the database and then restore it with indexes being rebuilt during the restore.

As part of evaluating a simpler, higher performance database structure, the current data in the mmtlogs miniservers tables were migrated to eventlog-like tables, similar to tables used within the re-aluminization database. The current mmtlogs database tables contain a large amount of redundant data. All data, even values that remain the same, are stored for each sampling timestamp. Each row in the database table, with a spreadsheet-like format, contains an "id" field, a "timestamp" field, and up to 150 or more parameter values associated with that timestamp. In the eventlog-log format, each row contains an id, timestamp, key, and value. Alternatively, a separate table can be created for each key with the table name matching the key name, reducing each row to an id, timestamp, and value. The eventlog table format is much more flexible than the spreadsheet table format. Parameters can be added or removed as needed. Also, a table entry is made only when the value changes, not at a fixed interval as in the current miniserver background logging.

After many tests, it was found that the eventlog-format logging required slightly more (around 10%) disk storage space than the current spreadsheet-format logging for the MMT miniserver data. This result was not expected. Performance for the two approaches is similar. The larger disk storage requirements for the eventlog-format database result mainly from the much larger index files associated with each database parameter. Each parameter value has two indexes, one as a unique "id" and one as a timestamp. The data from the hexapod and mount miniservers, for example, have well over 100 parameters with the same timestamp. In the spreadsheet-like table format, these 100+ parameters have only one "id" and one "timestamp" index. The index files are small. In the eventlog-format tables, individual parameter values are indexed. The storage space for these extra indexes is slightly more than simply repeating the value in the spreadsheet-like format.

These results hold for the data being stored in the MMT miniserver background log databases. For data in which fewer parameters are associated with a single timestamp, the eventlog-format could easily require less data storage requirements than a spreadsheet-format approach. Database storage requirements are dependent on table structure and the nature of the data being stored.

It was decided to stay with the current spreadsheet-like format for the mmtlogs database, both because of slightly reduced disk storage requirements and because of the large amount of work that would be required for modifying existing code to access data stored in the eventlog-like format.

During this reporting period, the mmtlogs database was migrated from MyISAM to InnoDB, independently of the MySQL installation on *hacksaw* (see Figure 6). The mmtlogs database on *hacksaw* remains in MyISAM format for a variety of reasons. As MySQL/MariaDB is currently configured on *hacksaw*, data would move from the NFS-mounted MyISAM database table files to the single ibdata1 file within the *hacksaw* virtual machine. This would greatly increase the size of the virtual disk image, which we want to keep at a minimum. Symbolic links could be used for the MySQL default data directory "datadir", but databases on *hacksaw* are currently both MyISAM and InnoDB files and are located in the default /var/lib/mysql location and also in the /mmt NFS-mounted file system via numerous symbolic links. Since this MyISAM-based MySQL server on *hackaw* has been our only database server and is in production, there is substantial risk in trying to re-configure MySQL on *hacksaw*. The MySQL/MariaDB configuration has remained unchanged in *hacksaw* with the mmtlogs database in a MyISAM format.

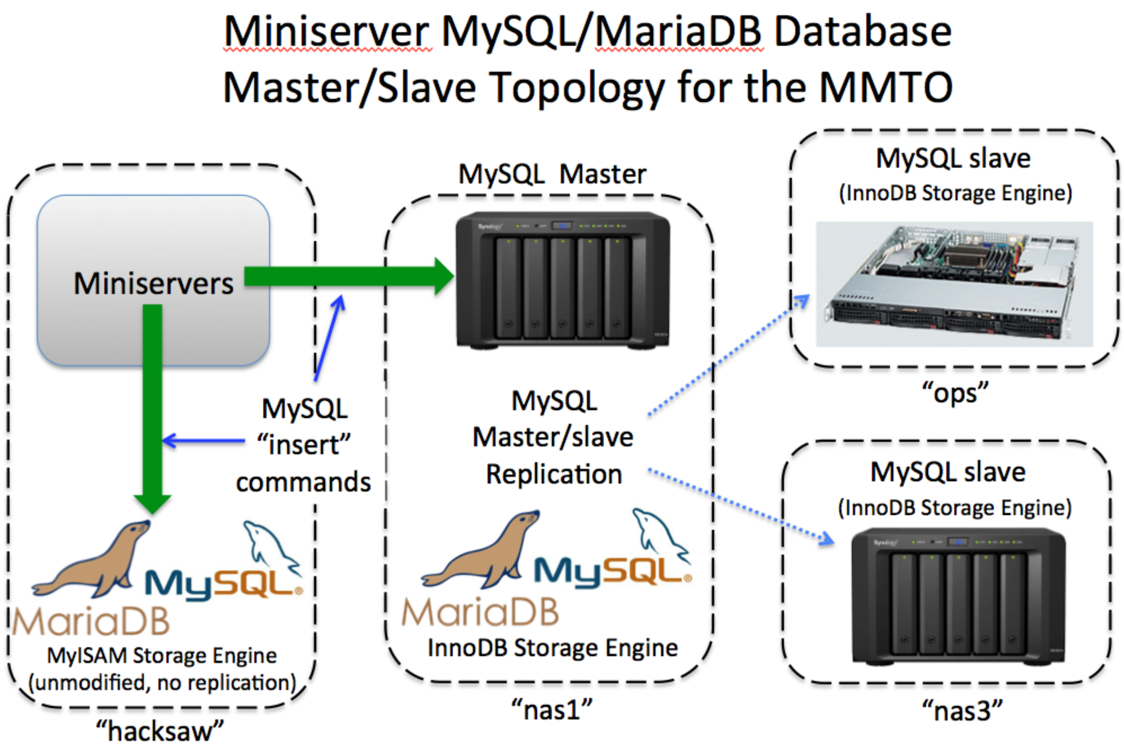

Figure 6. MySQL/MariaDB database master/slave topology at the MMTO. Prior to this recent work on the database administration, the only copy of the MMT "mmtlogs" database was on *hacksaw*, using NFS-mounting of data files and indexes in MyISAM storage format on *nas2*. This configuration remains. MySQL "insert" commands are sent to a new master MySQL/MariaDB database, located on *nas1*. This master is currently replicated onto two slaves: one on *nas3* and one on the new *ops* computer. It is planned that most database queries will be to the *ops* computer, preventing CPU load level spikes on *hacksaw/nas2*. The *ops* slave will be moved to the summit when it is more completely configured. The *nas3* slave will remain at the campus office and can serve as a backup MySQL slave. It can also provide database information if network connection between campus and the summit is unavailable.

Operator's Paddle

The operator's paddle (Figure 7) was disconnecting frequently, which required the telescope operator to cycle the power. Some changes were made to the Ardiuno code to help improve stability. A USB cable was also connected to the telstat Mac Mini computer to enable logging from the Arduino. This allows for remote troubleshooting. A new watchdog script runs on telstat that pings the paddle Arduino every minute and, if it fails to respond, will reset it using the serial connection. We have found that the lockup problem is due to the limited memory on the Arduino while it handles both paddle control and TCP/IP on the chip. D. Porter is evaluating a future replacement for the Arduino, most likely an Arduino Yun (which has a separate Linux CPU for TCP/IP) or a Beaglebone. Since the watchdog script and the code change, however, there have been no reported complaints of the paddle being down.

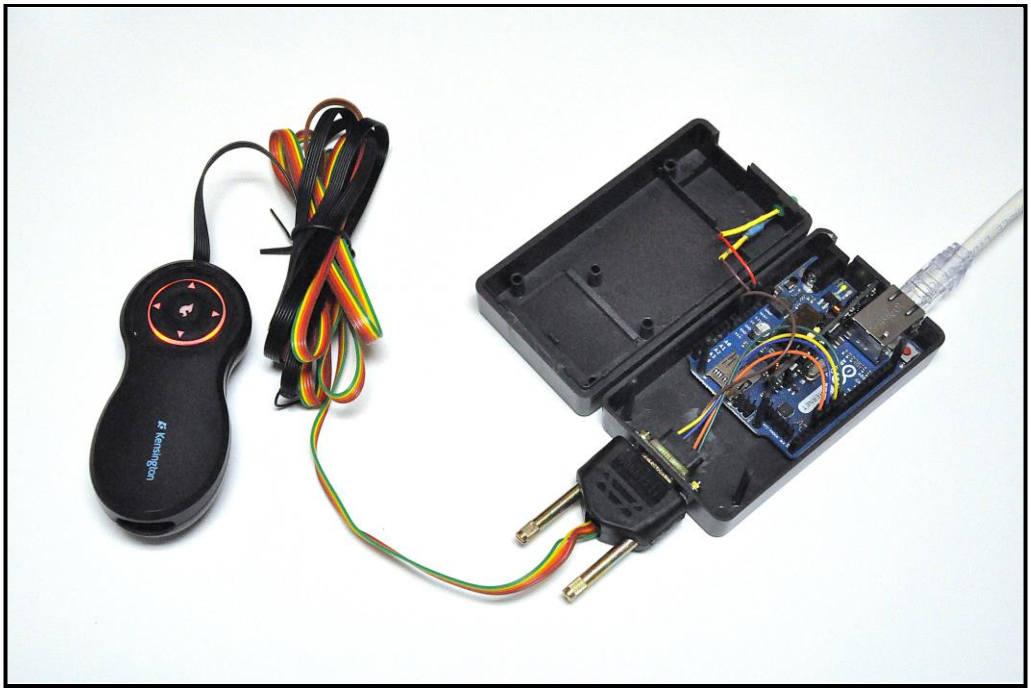

Figure 7. The operator paddle hardware, including the Arduino computer on the right.

Miniservers/Annunciator

Modifications were made to the "redis_tcs" mmtservice. This server reads "derived" parameters from the dataserver2/annunciator, and posts changed values to the master Redis server running on *hacksaw*. Dataserver2-derived parameters are parameters that may come from one of several sources or may be the same parameter expressed in different units. For example, the device that reads the average wind speed and direction can vary from weather stations located on either the east or west side of the MMT building. In this case, the weather sensor with the highest wind speed is used as the data source. The logic for determining these derived parameters is within the dataserver2/annunciator code.

QR Code/Barcode database

D. Porter and T. Gerl continued work on a Quick Response Code (QR code) or barcode database for equipment and spares at the summit.

Weather and Environmental Monitoring

Seeing

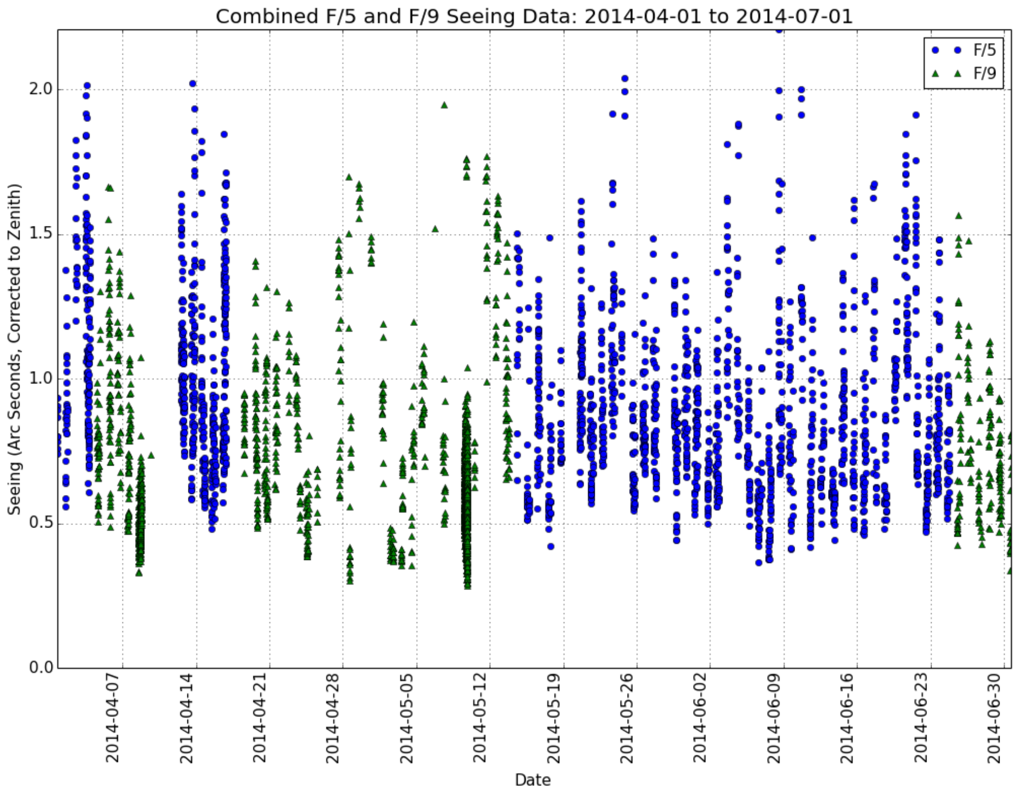

Figure 8 shows the seeing data from both the f/5 (blue) and f/9 (green) WFS datasets as a function of time during the April-June quarter. Seeing measurements between the two configurations are historically identical. Through this quarter, the f/9 seeing was remarkably tight, leading to slightly differing median seeing values.

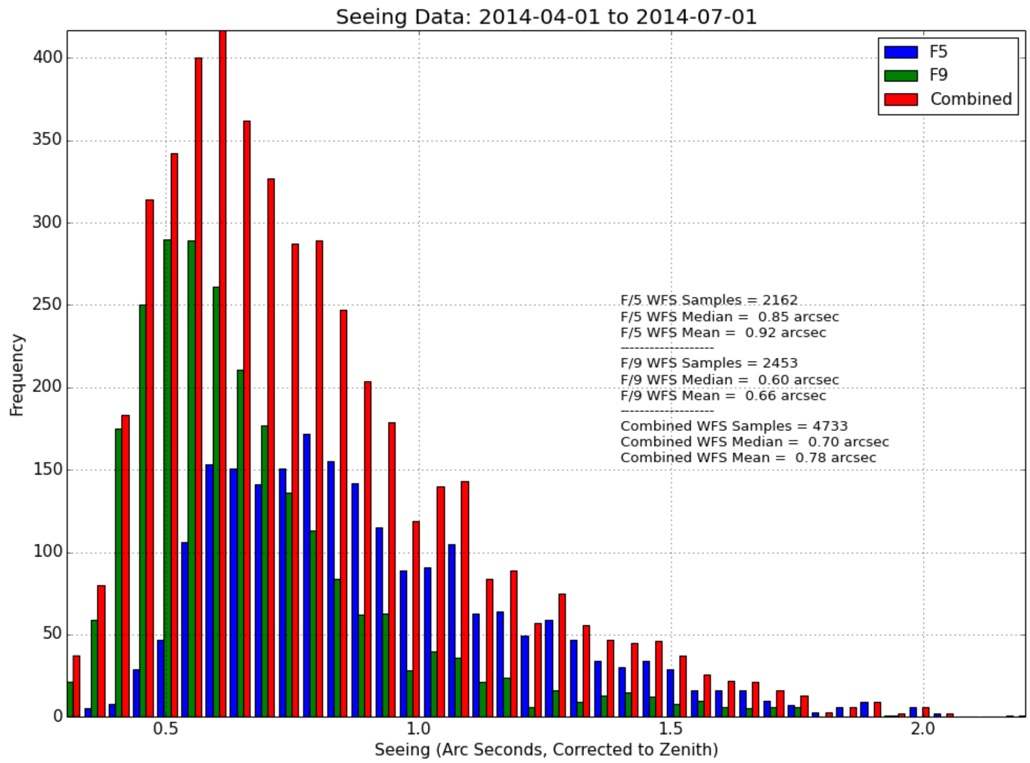

Figure 9 shows the histograms of seeing values for the April-June quarter for f/5 (blue), f/9 (green), and the full data set (red). The f/5 seeing (0.85”) was similar to the historical median MMT seeing, but the f/9 seeing was exceptional with a median seeing value of 0.6”. Taking the full dataset, the median seeing for this quarter was 0.7” compared to a historical median of ~0.8”, illustrating the excellent seeing.

User Support

Remote Observing

The MMTO supported a total of ten nights of remote observing this quarter. Three nights were for CfA, 6.5 nights were for UA/ASU observers, and a half night of Director’s time allocated to J. Farihi (IoA, Cambridge).

Data Quality Assessment

Work continues on a software pipeline that will run daily after a night of f/9 observing to check for evolution in several properties present in the data. This code will alert staff if there is a change in the bias level in the detector or the relative strength of arc lines, or a significant change in the count rates in the flat fields. We are still in the process of fine-tuning these alerts to ensure that real changes are noted while minimizing false positive alerts.

Data Archive

All of the data taken using the observing machine, *pixel*, continue to be archived and backed up on a NAS box on campus. Three users have requested access to the archive to recover a file that was missed when copying files after their run.

Documentation

Nothing to report.

Public Relations and Outreach

Visitors and Tours

4/4/14 – G. Williams gave a tour of the MMTO to two astronomers from the Korea Astronomy and Space Science Institute (KASI) who are also involved in UA-Korean astronomy collaborations. J.S. Kim, astronomer at Steward Observatory, accompanied them.

4/14/14 – Prof. Dae-Sik Moon and Dr. Suresh Sivanandam (former Steward graduate student), both of the University of Toronto, were given a tour of the MMTO by G. Williams.

4/21/14 - 4/22/14 – A 24-hour time-lapse video was shot by Sean Parker, a Tucson photographer hired by the SAO Science Education Dept. The video was to show the sky and landscape as seen from Mt. Hopkins and the MMTO and will be used for an exhibition and education project with the Boston Children's Museum.

4/29/14 – G. Williams gave a tour of the MMTO to seven astronomers from the Tokyo Astronomical Observatory.

5/3/14 – G. Williams gave a tour of the MMTO to the Huachuca (AZ) Astronomy Club.

5/23/14 – A high school student and his mother are shown at the MMTO during their visit to FLWO. He won an award from the Tucson Amateur Astronomy Association Science Fair for his project on black holes. Part of the award included a guided visit to FLWO. He is in the 11th grade and is interested in astrophysics, particularly cosmology. E. Falco hosted their visit (his shadow can be seen on the right taking this picture).

MMTO in the Media

CfA announced the discovery of a white dwarf binary system that is rotating very quickly and is one of the best gravitational wave sources yet to be identified. The data were taken with the MMT Blue Channel Spectrograph and at Gemini Observatory. http://www.cfa.harvard.edu/news/fe201419

A runaway star cluster from the M87 galaxy was discovered using the MMT and Hectospec. http://www.cfa.harvard.edu/news/2014-09. The story later appeared on websites such as Universe News, Space Daily, Fox News, and Tech Times.

Site Protection

4/29/14 – J. Hinz attended the Cochise County Code Task Force meeting where county lighting codes were discussed.

Appendix I - Publications

MMT Related Scientific Publications

(An online publication list can be found in the MMTO ADS library at http://www.mmto.org/node/244 *)*

14-24 A Multiwavelength Study of Embedded Clusters in W5-east, NGC 7538, S235, S252 and S254-S258
L. Chavarria, L. Allen, C. Brunt, et al.
MNRAS, **439**, 3719

14-25 CSI 2264: Characterizing Accretion-burst Dominated Light Curves for Young Stars in NGC 2264
J. Stauffer, A.M. Cody, A. Baglin, et al.
AJ, **147**, 83

14-26 The 400d Galaxy Cluster Survey Weak Lensing Programme. III. Evidence for Consistent WL and X-ray Masses at $z \approx 0.5$
H. Israel, T.H. Reiprich, T. Erben, et al.
A&A, **564**, 129

14-27 High-resolution Imaging of *Kepler* Planet Host Candidates. A Comprehensive Comparison of Different Techniques
J. Lillo-Box, D. Barrado, and H. Bouy
A&A, **566**, 103

14-28 Dusty WDs in the WISE All Sky Survey SDSS
S.D. Barber, M. Kilic, W.R. Brown, A. Gianninas
ApJ, **786**, 77

14-29 Reducing Systematic Error in Weak Lensing Cluster Surveys
Y. Utsumi, S. Miyazaki, M.J. Geller, et al.
ApJ, **786**, 93

14-30 An Anomalous Extinction Law in the Cep OB3b Young Cluster: Evidence for Dust Processing during Gas Dispersal
T.S. Allen, J.J. Prchlik, S.T. Megeath, et al.
ApJ, **786**, 113

14-31 Comparison of Galaxy Clusters Selected by Weak-lensing, Optical Spectroscopy, and X-Rays in the Deep Lens Survey F2 Field
S. Starikova, C. Jones, W.R. Forman, et al.
ApJ, **786**, 125

14-32 A Globular Cluster toward M87 with a Radial Velocity < -1000 km s^{-1}: The First Hypervelocity Cluster
N. Caldwell, J. Strader, A.J. Romanowsky, et al.
ApJ Lett, **787**, L11

14-33 MMT Hypervelocity Star Survey. III. The Complete Survey
W.R. Brown, M.J. Geller, S.J. Kenyon
ApJ, **787**, 89

14-34 Optical Spectra of 73 Stripped-envelope Core-collapse Supernovae
M. Modjaz, S. Blondin, R.P. Kirshner, et al.
AJ, **147**, 99

14-35 Optical Spectroscopic Observations of γ-Ray Blazar Candidates. I. Preliminary Results
A. Paggi, D. Milisavljevic, N. Masetti, et al.
AJ, **147**, 112

14-36 Hydrogen-poor Superluminous Supernovae and Long-duration Gamma-Ray Bursts have Similar Host Galaxies
R. Lunnan, R. Chornock, E. Berger, et al.
ApJ, **787**, 138

14-37 Improved Determination of the Atmospheric Parameters of the Pulsating sdB Star Feige 48
M. Latour, G. Fontaine, E.M. Green, et al.
ApJ, **788**, 65

14-38 Software framework for the upcoming MMT Observatory primary mirror re-aluminization
J. D. Gibson, D. Clark, and D. Porter
Proc. SPIE, **9152**, 91521K-14

MMT Technical Memoranda / Reports

None

Non-MMT Related Staff Publications

Spitzer Ultra Faint Survey Program (SURFS UP). I. An Overview
Bradac, M., Ryan, R., Casterano, S., et al., (J. Hinz)
ApJ, **785**, 108

Measuring the Stellar Masses of z~7 Galaxies with the Spitzer UltRaFaint Survey Program (SURFS UP)
Ryan, R. E., Jr., Gonzalez, A. H., Lemaux, B. C., et al., (J. Hinz)
ApJ, **786**, L4

The Baryonic Tully-Fisher Relationship for S4G Galaxies and the "Condensed" Baryon Fraction of Galaxies
Zaritsky, D., Courtois, H., Munoz-Mateos, J.-C., et al., (J. Hinz)
AJ, **147**, 134

Reconstructing the Stellar Mass Distributions of Galaxies Using S4G IRAC 3.6 and 4.5um Images. II. The Conversion from Light to Mass
Meidt, S. E., Schinnerer, E., van de Ven, G., et al., (J. Hinz)
ApJ, **788**, 144

Appendix II - Service Request (SR) and Response Summary: April – June, 2014

The MMT Service Request (SR) system is an online tool to track ongoing issues that arise primarily during telescope operations, although the system can be used throughout the day and night by the entire staff. Once an SR has been created, one or more responses are created by staff members to address and eventually close the SR. These SRs and associated responses are logged into a relational database for later reference.

Figure 10 presents the distribution of SR responses by priority during the period of April through June 2014. As seen in the figure, the highest (47%) of responses were of "Important" priority. "Critical," "Near-Critical," and "Low" priorities were each 16% with "Information Only" at 6%. "Critical" SRs address issues that are preventing telescope operation, while "Near-Critical" SRs relate to concerns that pose an imminent threat to continued telescope operation. There were a total of 32 SRs during this three-month period.

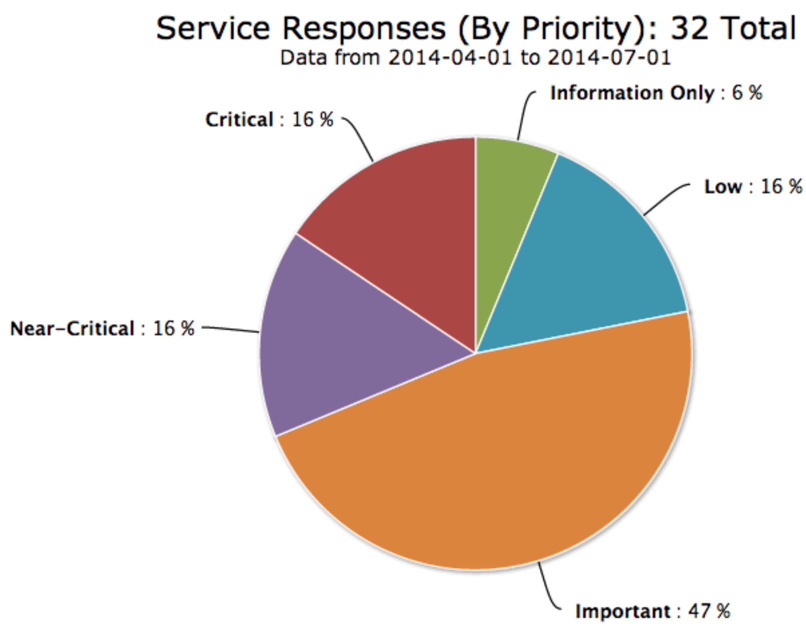

Figure 10. Service Request responses by priority during April through June 2014. The majority (47%) of the responses are related to SRs of "Important" priority, while "Critical," "Near-Critical," and "Low" priority responses are 16% each of all responses.

Figure 11 presents the same 32 SR responses grouped by category. These categories are further divided into subcategories for more detailed tracking of issues. The majority of the responses from April through June were related to the "Building" and "Thermal Systems" categories. There were four responses each for "Computers/Network," "Software," and "Telescope."

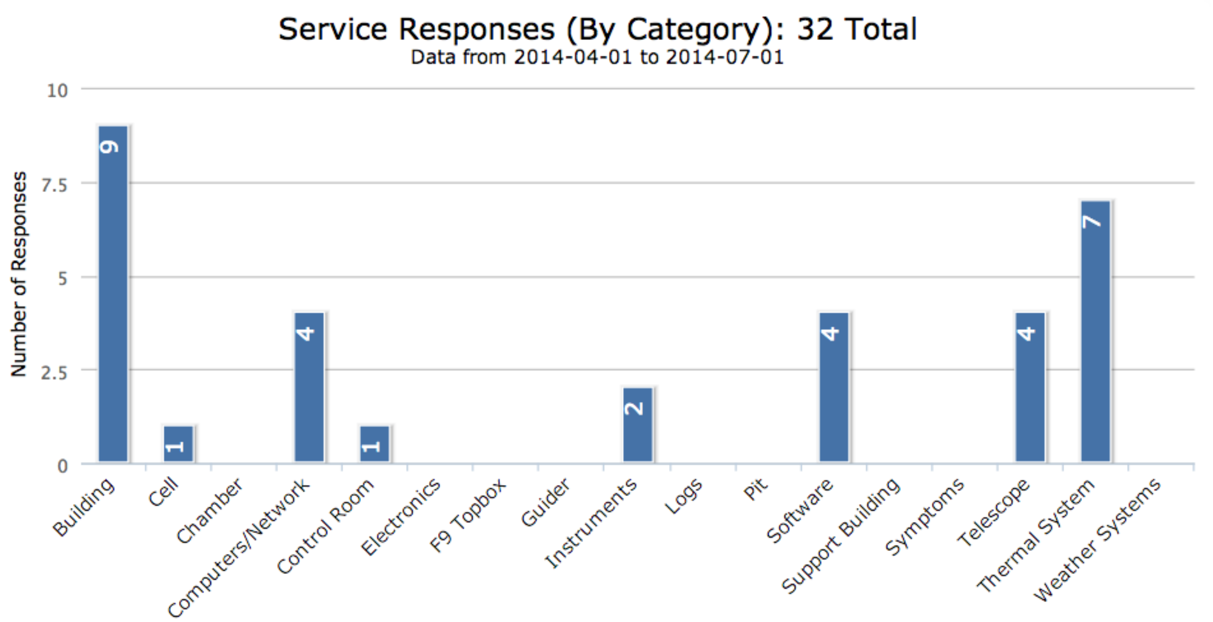

Figure 11. Service Request responses by category during April through June 2014. The majority of responses were within the "Building" and "Thermal Systems" categories.

Appendix III - Observing Statistics

The MMTO maintains a database containing relevant information pertaining to the operation of the telescope, facility instruments, and the weather. Details are given in the June 1985 monthly summary. The data attached to the back of this report are taken from that database.

Use of MMT Scientific Observing Time

April 2014

Instrument	Nights Scheduled	Hours Scheduled	Lost to Weather	*Lost to Instrument	**Lost to Telescope	***Lost to Gen'l Facility	****Lost to Environment	Total Lost
MMT SG	13.00	117.90	28.05	0.75	0.50	0.00	0.00	29.30
PI Instr	13.00	121.40	28.75	0.00	0.75	0.00	0.00	29.50
Engr	4.00	37.30	16.15	0.00	0.00	0.00	0.00	16.15
Sec Change	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total	**30.00**	**276.60**	**72.95**	**0.75**	**1.25**	**0.00**	**0.00**	**74.95**

Time Summary

Percentage of time scheduled for observing	86.5
Percentage of time scheduled for engineering	13.5
Percentage of time scheduled for sec/instr change	0.0
Percentage of time lost to weather	26.4
Percentage of time lost to instrument	0.3
Percentage of time lost to telescope	0.5
Percentage of time lost to general facility	0.0
Percentage of time lost to environment (non-weather)	0.0
Percentage of time lost	27.1

* Breakdown of hours lost to instrument
0.75 Spectrograph GUI issue

** Breakdown of hours lost to telescope
0.25 Oscillation
0.50 WFS coding issue
0.50 Hexapod limit issue

May 2014

Instrument	Nights Scheduled	Hours Scheduled	Lost to Weather	*Lost to Instrument	** Lost to Telescope	***Lost to Gen'l Facility	****Lost to Environment	Total Lost
MMT SG	12.00	102.10	16.15	0.00	0.00	0.00	0.00	16.15
PI Instr	17.00	137.20	29.00	0.00	0.00	0.00	0.00	29.00
Engr	2.00	16.80	0.00	0.00	0.00	0.00	0.00	0.00
Sec Change	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total	**31.00**	**256.10**	**45.15**	**0.00**	**0.00**	**0.00**	**0.00**	**45.15**

Time Summary

Percentage of time scheduled for observing	93.4
Percentage of time scheduled for engineering	6.6
Percentage of time scheduled for sec/instr change	0.0
Percentage of time lost to weather	17.6
Percentage of time lost to instrument	0.0
Percentage of time lost to telescope	0.0
Percentage of time lost to general facility	0.0
Percentage of time lost to environment (non-weather)	0.0
Percentage of time lost	17.6

Year to Date May 2014

Instrument	Nights Scheduled	Hours Scheduled	Lost to Weather	Lost to Instrument	Lost to Telescope	Lost to Gen'l Facility	Lost to Environment	Total Lost
MMT SG	63.00	653.80	234.30	1.50	0.66	12.50	0.00	248.96
PI Instr	77.00	763.20	161.40	1.75	26.40	0.00	0.00	189.55
Engr	11.00	111.40	16.15	0.00	0.00	0.00	0.00	16.15
Sec Change	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total	**151.00**	**1528.40**	**411.85**	**3.25**	**27.06**	**12.50**	**0.00**	**454.66**

Time Summary

Percentage of time scheduled for observing	92.7
Percentage of time scheduled for engineering	7.3
Percentage of time scheduled for sec/instr change	0.0
Percentage of time lost to weather	26.9
Percentage of time lost to instrument	0.2
Percentage of time lost to telescope	1.8
Percentage of time lost to general facility	0.8
Percentage of time lost to environment (non-weather)	0.0
Percentage of time lost	29.7

Use of MMT Scientific Observing Time

June 2014

Instrument	Nights Scheduled	Hours Scheduled	Lost to Weather	*Lost to Instrument	**Lost to Telescope	***Lost to Gen'l Facility	****Lost to Environment	Total Lost
MMT SG	5.00	38.50	2.50	0.00	0.50	0.00	0.00	3.00
PI Instr	25.00	194.00	3.50	0.00	0.25	4.50	0.00	8.25
Engr	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Sec Change	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total	**30.00**	**232.50**	**6.00**	**0.00**	**0.75**	**4.50**	**0.00**	**11.25**

Time Summary

Percentage of time scheduled for observing	100.0
Percentage of time scheduled for engineering	0.0
Percentage of time scheduled for secondary change	0.0
Percentage of time lost to weather	2.6
Percentage of time lost to instrument	0.0
Percentage of time lost to telescope	0.3
Percentage of time lost to general facility	1.9
Percentage of time lost to environment	0.0
Percentage of time lost	4.8

** Breakdown of hours lost to telescope

0.25 Hard point error

0.50 Guider errors

*** Breakdown of hours lost to facility

4.50 Generator failure

Year to Date June 2014

Instrument	Nights Scheduled	Hours Scheduled	Lost to Weather	Lost to Instrument	Lost to Telescope	Lost to Gen'l Facility	Lost to Environment	Total Lost
MMT SG	68.00	692.30	236.80	1.50	1.16	12.50	0.00	251.96
PI Instr	102.00	957.20	164.90	1.75	26.65	4.50	0.00	197.80
Engr	11.00	111.40	16.15	0.00	0.00	0.00	0.00	16.15
Sec Change	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total	**181.00**	**1760.90**	**417.85**	**3.25**	**27.81**	**17.00**	**0.00**	**465.91**

Time Summary

Percentage of time scheduled for observing	93.7
Percentage of time scheduled for engineering	6.3
Percentage of time scheduled for secondary change	0.0
Percentage of time lost to weather	23.7
Percentage of time lost to instrument	0.2
Percentage of time lost to telescope	1.6
Percentage of time lost to general facility	1.0
Percentage of time lost to environment	0.0
Percentage of time lost	26.5