

Research Article

Converting from Classical to Robotic Astronomy: The Lowell Observatory 0.8-m Telescope

Marc W. Buie

Southwest Research Institute, 1050 Walnut Street, Suite 300, Boulder, CO 80302, USA

Correspondence should be addressed to Marc W. Buie, buie@boulder.swri.edu

Received 3 June 2009; Accepted 2 December 2009

Academic Editor: Alberto J. Castro-Tirado

Copyright © 2010 Marc W. Buie. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

The Lowell Observatory 0.8-m (31-in) Telescope was augmented by the addition of a robotic mode of operation in early 2001. This system executes any predetermined sequence of observing instructions without supervision. Described herein is a general description of the system, lessons learned from the conversion, and a few general algorithms for focusing, collecting twilight flat field images, and scheduling standard star observations.

1. Introduction

The power and value of autonomous observing has long been recognized in the amateur astronomy community while adoption has been slow at professional observatories. In January 2001 I converted the Lowell Observatory 0.8-m telescope from a classically-operated manual telescope to a robotic facility. At that time the telescope was poorly utilized and deemed unimportant enough to permit my experimenting with the system. The perception was that such a conversion and subsequent robotic operation was either dangerous to the hardware or disruptive to the operation of the telescope—a fear that was not realized. Since 2001 roughly 1/3 of the time was scheduled as classical observing and the rest operated in a robotic mode.

The ground rules for the conversion were (1) classical observations must still be possible without any changes in operation or procedures and (2) instrument changes had to continue to be supported. The project was further constrained by the lack of development funds and minimal technical staff support. This constraint led to a very frugal development path with minimal changes to the existing system. The new mode became operational and scientifically useful in March 2001 and immediately caused this telescope to be the most heavily used of any of the Lowell Observatory facilities. During this period of operation the system had to be modified to work with two different cameras as well as a total changeout of the dome control system.

In this paper the design objectives and constraints involved in moving this telescope from classical hands-on observing to robotic hands-off observing are summarized. The observing queue system and two camera systems will be discussed along with changes needed in telescope control. I also present some of the non-obvious algorithmic components of the system as lessons learned that may be useful to other projects. Also discussed are safety issues common to systems of this type. At the end I will present a brief synopsis of the types of observations carried out by this system since its implementation.

2. System Overview

This system was created to permit autonomous collection of images at the Lowell Observatory 0.8-m telescope on Anderson Mesa. The primary goal was to provide a means for collecting data without requiring an observer to be present. This approach makes data collection much more cost-effective in circumstances where a large number of nights are needed to yield a single useful result. For example, collecting a lightcurve of Pluto (period ~6.4 days) within one apparition requires a large number of nights to create a final composite lightcurve. Any type of survey or monitoring observation is also very well suited to this system.

Figure 1 is a flowchart showing the principal hardware subsystems, computers, software and the communication

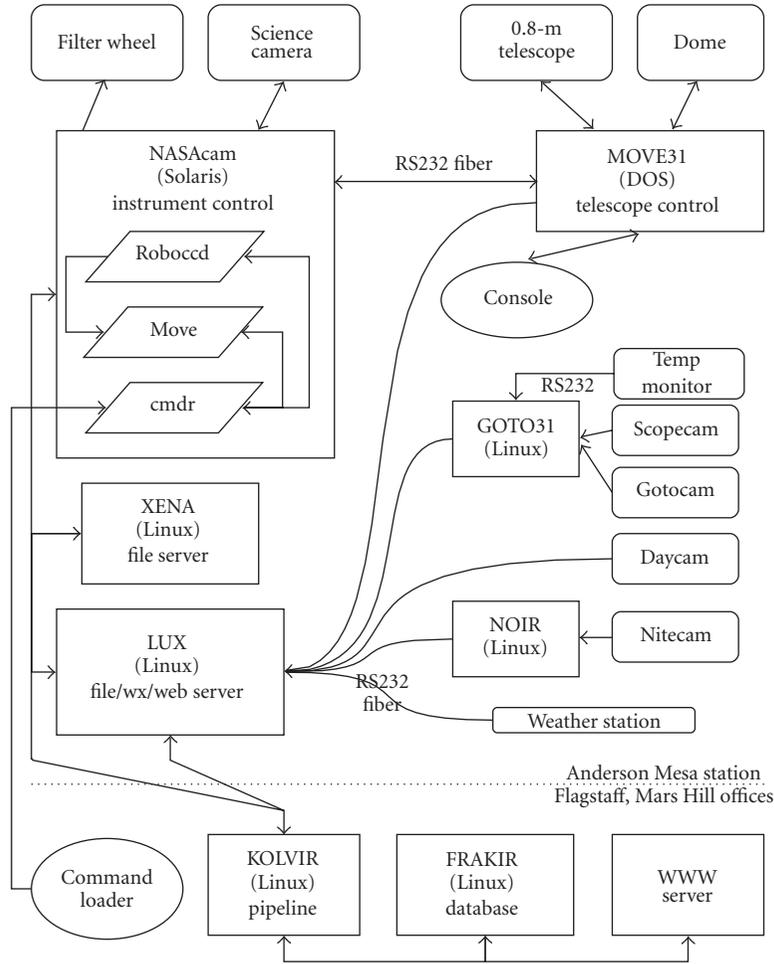


FIGURE 1: System overview; see text for details.

paths between them. The bulk of the figure describes systems in place at the Anderson Mesa Station where the telescope is situated. At the very bottom of the diagram a few machines are shown that are located at the main observatory office 20 km away. In this diagram, all of the rectangles depict specific computers (9 total). The rounded rectangles stand for specific hardware. The parallelograms show processes that are always running while the system is active. Ovals indicate where user input can enter into the system. All computers communicate at each site via a 100 Mbs ethernet LAN, typically using NFS for communications but occasionally using rsync, scp, or direct TCP/IP sockets. Communication between sites is provided by a dedicated T1 link but this link is not required for the system to run and collect data.



FIGURE 2: 31" dome and control building.

3. Hardware

This section describes the hardware used by the robotic system. A photograph of the facility exterior is shown in Figure 2. The white building on the left houses the telescope and telescope control computer. The low brown building houses a traditional control room and a computer room.

3.1. Telescope. The telescope is a 0.8-m closed-tube reflecting telescope with a single instrument port at the cassegrain focus. The output beam is $f/15$ with a focal plane scale of 17.5 arcsec/mm. The scale is poorly matched to modern CCD detectors requiring the use of a focal reducer to achieve a



FIGURE 3: 31" telescope and PCCD camera.

more useful field of view. The telescope tube is constructed of aluminum and as a result has a very strong shift in focus position with changing temperature. Figure 3 shows the telescope with the first generation CCD system.

The facility already had a sophisticated observatory control system that was the result of decades of local software and hardware development. The control program, named MOVE31, is written in FORTRAN with a small amount of assembly code and operates on an MSDOSv6.2 computer. The telescope has three digital stepper motors, one each for the Right Ascension and Declination axes and one for moving the secondary mirror for focus control. Position knowledge comes from relative encoders. Initially only the focus mechanism had a home switch for calibrating its absolute location. More recently a pair of switches was added to provide a home location from which to restore pointing in case of error. However, all of the startup and observing strategies grew out of the environment where loss of pointing meant the loss of a night and subsequent on-site effort to recover. The english yoke mount imposes a north declination limit of roughly 60° beyond which the telescope cannot point. There is a slight protrusion of the north pier that the telescope can track into. With classical observing one can work around this protrusion and extend the declination coverage. For robotic observing the northern limit was made more conservative so that the protrusion can never be reached. However, the limit is implemented in the scheduling software, not in the telescope control program.

The telescope has had a variety of positional limit switches over the years. All of the switches work well at normal tracking rates. The protections are not as effective while the telescope is slewing. As a result, the robotic system was forced to take on a larger pointing limit responsibility than it should. Over the years there was only one critical failure where the telescope was driven into the north pier, breaking the drive coupler that was designed to fail in such cases, protecting the drive motor and gears.

3.2. Dome. The telescope is protected from weather by an Ash dome. It uses high-power a/c motors operated via computer controlled relays for rotation and for opening and

closing the two-part dome slit. The dome slit has a long upper panel that is lifted up and over the top of the dome when open. It also has a short lower panel that drops down below the bottom of the dome. Opening or closing the dome takes roughly 3 minutes. The telescope mirror becomes vignetted below 18° altitude.

Power to the dome slit motors is provided by a bus-bar contact that is engaged over a 5° range in dome azimuth. To open and close the dome one must successfully stop on the correct position before the open/close operation will work. The original dome encoding system was handled with a friction wheel drive relative encoder and the power contact was wired to act as a fiducial switch. This system required elaborate software to maximize the likelihood of closing the dome in case of problems. In later years the dome encoding system was replaced by a bar code scanner providing an absolute position every degree, greatly simplifying the calibration process and making for considerably more reliable operation with much simpler software. Over the years there have been a few instances where a failure occurred and the dome was not properly closed at the end of the night. Procedurally this has led to a policy that someone always verify that the dome is closed after sunrise.

Dome control during the night is handled completely by the MOVE31 computer once automatic dome control is enabled. A minor software modification was required for robotic observing such that notification of the completion of a slew had to wait for the telescope *and dome* to both reach their final position. With classical observing an observer would know from the noise of the dome motion that the system was not ready.

3.3. Science Camera. The robotic system was originally developed with a thermoelectrically cooled Photometrics CCD system (PCCD) and a 3 : 1 focal reducer. The detector was 576×384 and had a pixel scale of 1.2 arcsec/pixel for a field of view of 11.5×7.7 arcmin. The focal reducer had a slight pupil ghost centered on the detector so the prime location for placing a single source was offset by a quarter of the FOV from the actual center of the device. This offset thus required better than about 3 arcmin pointing from a slew to ensure successful target acquisition. Full-frame readout time and saving the data to disk for this detector was 8 seconds. The shutter was a commercial two-leaf system that was not particularly fast. Photometrically accurate exposures had to be longer than 3 seconds but the system could take useful exposures as short as 0.002 seconds. The shorter exposures were not useful for science data but came in handy for other operational support modes. The thermoelectric cooling restricted exposure times to 5 minutes or less and the system required dark-frame calibrations.

The camera system also included an integrated 10-position filter wheel that provided fast cycling between filters and reliable knowledge of the state of the filter wheel. The wheel could only move in one direction and stepping by one position took less than one second. The worst case was for moving backward by one position which require moving forward by 9 positions and took about 3 seconds. Any time

a motion crossed the fiducial the controller would stop at the fiducial, recalibrate and then then move on to the final destination. Thus, observing with any two-filter pattern (e.g., alternating between *B* and *V*) would lead to a verification of filter position knowledge and a recalibration once per pattern. The typical observing pattern for this system put a lot of cycles on the filter wheel. Twice during the robotic usage of this camera the system wore out a critical part but the recalibration strategy and error logs gave several months notice before total failure.

The performance of the original system was limited by its plate scale, FOV, high-readnoise, low QE, and ultimately wore out and was removed from service in July 2005. A new camera was constructed for the telescope that was based on a custom focal reducer matching the CCD to the plate scale of the telescope. The initial CCD was a Loral detector that was thinned by M. Lesser of Steward Observatory. This batch of devices is prone to delamination and combined with a slow vacuum leak in the dewar this detector never collected science grade data even though it was in service for 9 months. By January 2007 a replacement CCD was put in service. This camera is known as NASAcam, so named for the source of the funds that made the camera possible.

NASAcam is built around an EEV 2048 × 2048 device. It has a image scale of 0.45 arcsec/pixel for a 15 arcmin field of view. The nominal gain is $2.15e^-/DN$ and a read-noise of $6.2e^-$. Lower read-noise was possible from this system but this setting was a compromise between readout time and readout noise. The full-frame read and store time is roughly 9 seconds. The detector is cooled by a Cryotiger closed-cycle refrigerator that maintains a regulated temperature of $-113^{\circ}C$, sufficient to completely eliminate dark-current. The system is intended to run full time, providing a long-term stable environment for the detector. This system has two 10-position filter wheels. The filter wheel system is a little slower than that for PCCD. The shortest move takes just under 3 seconds while the longest move requires about 9 seconds.

The requirements for adapting this system to a new camera are relatively modest. The most critical is the availability of a Linux or Solaris driver. The only other requirement for efficient operation is the ability to read out a portion of the image with the expectation that the readout time is much faster for small subframes. System performance is also improved if the image last read is available to the application at some memory location.

3.4. Support Cameras. An essential element of this system are cameras to provide on-site visual feedback. Four cameras based on inexpensive surveillance camera hardware are used to monitor daytime sky conditions (DAYcam), nighttime sky conditions (NITEcam), the sky at the current telescope location (GOTOcam), and a view of the telescope and dome (SCOPEcam).

DAYcam is the simplest of all the cameras. The camera is a D-Link DCS-1000 and is a self-contained color video camera and web server. The stock lens was replaced with a 3.8 mm fixed-aperture fixed-focus lens that provided a field

of view of about 90° . This camera is housed inside a north-west facing window to protect it from the weather. Every 2 minutes a picture is sent via FTP to a central network location. These images are managed and culled by simple Perl scripts that file them by day and delete images taken when the sun is below -5° altitude when the sky is too dark for this camera. The pointing of the camera is set so that the sun and moon can never pass through the field of view. These color images are required for times when you cannot discern between high thin cirrus and a photometric sky without seeing the color of the sky. The data rate from this camera is about 4 Mb/day.

NITEcam provides information about clouds at night, even when there is no moonlight. To accomplish this, I used a Watec 903 K b/w video camera that has a 1/3" CMOS CCD detector with a minimum lux rating of 0.0002 lux. In comparison, DAYcam only works down to 2.5 lux. The camera is mated to a Computar HG3808AFCS-HSP video-mode auto-iris lens, focal length 3.8 mm, with a maximum aperture of $f/0.8$. The lens settings are set to full average brightness control. At night, the natural signal from the camera with the moon down does not yield a very useful image though it's not bad once you have illumination from the Moon. To grab the video, I use a Hauppauge WinTV-GO card mounted in an old surplus Pentium 200 Mhz computer running RedHat 7.3. The camera is located in the same window bay as DAYcam and also avoids looking at the sun and the moon. The setup of the Watec camera is automatic electronic shutter control on, hi-AGC on, and AGC on (AGC = automatic gain control). During the day an image one frame per minute is collected and saved. The night-time sensitivity is enhanced by coadding 200 individual video frames and then subtracting a predetermined dark/bias image thus increasing the sensitivity by a factor of 10. This sensitivity is sufficient to detect unilluminated clouds in silhouette against the dark night sky (no moon). The software also has some automatic gain logic that keeps the image in range on brighter nights. This camera provides full-time images of the sky and generates about 19 Mb of data per day.

SCOPEcam is a stock b/w CCTV camera that provides a classical observer with a remote real-time view of the telescope and dome while observing in the control room. To support robotic observing, the signal was split and routed to a computer with another WinTV video capture card. The camera is not sensitive enough to provide useful images at night with natural lighting so the MOVE31 computer turns on a 5W red light in the dome whenever the telescope is slewing (or by explicit command). However, during the day it is immediately obvious if the dome is open even just a centimeter. The most important role this camera plays is in verifying that the dome is closed at the end of the night. This camera is also useful for debugging failures during robotic observing. To minimize the amount of data collected from this camera, images are saved only when something either moves (e.g., telescope slews) or if the lighting changes dramatically (e.g., a light is turned on). The data flow from this camera is extremely variable but averages to roughly 9 Mb per day.

The final camera, GOTOcam, has the same hardware and software as NITEcam except the camera lens. This image features a pixel scale of roughly 3 arcmin and a field of view of 14×10 degrees and is boresighted with the telescope. This camera makes it possible to recover pointing by positioning a bright star at a location that is known to be coincident with the science detector. These images also can be used to decide if data are affected by clouds. In practice, this data stream is rarely used and is marginally worth implementing. Images are collected only when the dome is open so the data rate is variable. On a clear night the camera generates roughly 40 Mb of data.

3.5. Environment Data. On-site weather conditions are collected with a Davis Instruments weather station. Special modification were made to increase the length of the sensor leads and to add some protection against lightning. The unit is connected to the logging computer via an RS232 fiber link. This station provides absolute barometric pressure, outside air temperature and relative humidity, dewpoint temperature, and wind speed and direction. Real-time information is available via socket communication and archival information is stored in text files. The weather data is primarily used to support predicting the weather conditions for the upcoming night and if it will be safe to open the telescope dome but it is not required in order to run the system. The local weather forecasts and satellite imagery is also used to support the nightly weather decision. In the early days this data was also used to estimate the focus position for the telescope at the beginning of the night.

With the development of NASAcam came an opportunity to improve the knowledge of the immediate telescope environment. The single most critical piece of missing information was the temperature of the telescope tube. To get accurate temperatures we used an OMEGA D5111 module that supports four temperature sensors (AD590) and has an RS232 port. One sensor was glued to the primary mirror (but done so that it could be removed when the mirror gets realuminized). One sensor is attached to the exterior of the telescope tube. Finally, there is one sensor dangling in the air at about the same height as the average position of the primary mirror. The tube temperature works well enough to predict the telescope focus position to within about $30 \mu\text{m}$ (focus operation has a typical step size of $6 \mu\text{m}$).

The telescope temperature data also provides information related to image quality. As expected, the mirror is often warmer than the ambient air. Surprisingly, the telescope tube was found to be significantly colder than the air and is just as detrimental to image quality as a hot mirror. The painted aluminum telescope tube radiatively couples to the night sky and over-cools during the night. This problem was easily fixed by wrapping the telescope in a loose layer of aluminized mylar. This simple solution broke the radiative link and afterward the tube was seen to track the air temperature very well. The seeing definitely improved as a result of this simple fix. At this point, the image quality is now limited by the quality of telescope tracking since there is no auto-guider.

4. Software

The basic design of the system is a collection of independent processes that share as little information as possible. Also deeply rooted in the system design is that all operations are unsupervised. That means there is no graphical display of any data, there are no real-time monitoring windows peeking into system processes, and queue manipulation is limited to adding commands or removing everything. Lastly, the software is designed to be running all the time. A specific effort was made to avoid forking and multi-threaded code for an easier path to reliable software. This coding constraint also led to the decision to use message queues for inter-process communication rather than TCP/IP sockets. The message queues provide a very simple and fast communication method but do require that all of the tightly coupled processes be running on the same computer. Communication via TCP/IP, where needed, is accomplished by adding a separate trivial process that accepts a socket connection and passes them along to a message queue.

Central to the very first versions of the system was the concept of event and anomaly logging in all of the programs. Every operation that can generate a system error (e.g., opening a file) will write a time-stamped system error message to a log file. Other informative messages about logical operations or decisions are included in the log as needed. This system has had its share of anomalies and software errors over the years. It is usually not practical to debug a problem by replicating the actions. These log files have been an essential and 100% effective tool for fixing the software without the need for error replication.

4.1. cmdr—Observing Queue. The top level controlling process is *cmdr* and it maintains a large FIFO of commands and passes them along, one at a time, to the appropriate agent for execution. If the FIFO becomes empty the system becomes idle. The *cmdr* daemon gets its commands over a inter-process communication (IPC) channel from the interactive program send. In many cases, this process has no knowledge of the consequences of any command and there is no concept of a time limit on any command. The communication protocol is to send a command via IPC channel to the appropriate agent, listen for optional information messages from the agent, and log the termination and time of execution when the completion message is returned. Information messages and command completions are logged with a timestamp as they are received. Commands consists of a single line of text. Single word lines are queue control commands. Everything else is a single character that is associated with either an external agent or a type of internal command followed by the command intended for the agent. There are four such categories defined: t: telescope commands, c: camera commands, f: file to be executed, and s: synchronization or timing commands.

The “file” command is really a special category since *cmdr* never actually sees such a command. Commands are give to send before being sent to *cmdr*. If send sees a file command it will process that file and send the commands

found in the file. Therefore, loading a full night's observing is usually a matter of loading the one file that contains all the commands. Any command that is itself a file command is expanded until nothing but non-file commands are sent. Thus a script for a night can appear to have relatively few commands, most of them file commands, that will be expanded into a very long list of actual commands that will occupy the fifo. All memory of the original organization of the files and commands is lost in the fifo other than the sequential order of execution. For this reason, no tools were developed to inspect the queue since it would be very hard to relate to the original script.

4.2. move—Telescope Communication and Control. The telescope and its dome are operated by a stand-alone computer (MOVE31 which refers to the computer and the control program) that maintains full knowledge and control of the telescope and the dome. This system supports a set of commands that can be sent via an RS232 serial port. All commands return an explanatory code upon completion. Some commands return instantly and some take as long as 3-4 minutes to complete.

The move process encapsulates all of the knowledge and idiosyncrasies of the MOVE31 system thus providing a clean commanding interface. This process takes care of preparing and sending the command and then waiting for the response. It also is aware of how long commands should take and imposes an appropriate length timeout period in case of errors. This process logs all activities, successful and unsuccessful, but any time there is a error the incident is also reported via email. Telescope failure codes are rare and always treated as a serious problem. A great deal of effort was expended to ensure that all errors are meaningful and worthy of response and intervention. Additional informational messages are also sent when the daemon is started or stopped and whenever the dome is opened or closed. This process also maintains an approximation of the telescope and dome state. It does so by querying the telescope every two seconds for its status and then parsing and saving that information. If another process wants to know the telescope state, the answer is returned instantly from the most recent query. This periodic query is not logged unless there is an error (such as a garbled ASCII string response or complete lack of response).

The move daemon listens to its own IPC channel for commands. There are three defined command sources, (1) primary command channel (usually from cmdr), (2) secondary command/query channel (roboccd), and (3) backdoor command/query channel (movecmd). The secondary channel is used to get the telescope status information to be saved along with each data frame. This channel is also used for pointing updates, small telescope offsets, and focus commands that are generated from the camera. The backdoor channel is used exclusively for testing or recovering from failures.

4.3. roboccd—CCD Camera Operation and Control. The CCD camera and filter wheels are controlled with the roboccd daemon. Image data is saved via NFS to a disk

served by XENA. Exposure durations are controlled by a hardware timer in the CCD electronics and the start time is taken from the system clock which is maintained by ntpd. Status information on the current state of ntp is recorded with each image. The roboccd daemon listens on its own IPC channel for commands from cmdr. Telescope focus is adjusted automatically for each filter based on differential focus offsets known for each filter. This program performs a few standard calculations on every image: (1) the mean and standard deviation of the background, (2) location and signal level of the maximum in the image, (3) instrumental magnitude, FWHM, and centroid location of the maximum. Quite a few decisions can be made that affect data collection with this limited knowledge. This information is recorded to the data headers and it is also sent to the log files on LUX. Thus is it possible to eavesdrop on the data collection process and get a reasonable idea of how things are progressing.

4.4. Synchronization. Timing and synchronization operations are handled internally by cmdr and are unusual to see in a classical observing system. Normally, all of these functions are performed implicitly by the observer as the night progresses. The following section lists the available synchronization commands

1st HH:MM:SS: Pause until the local sidereal time passes HH:MM:SS.

pause HH:MM:SS: Pause for the indicated length of time.

time HH:MM:SS: Pause until the UT time passes HH:MM:SS.

sun-above D.d: Wait while the Sun is above D.d degrees altitude. Used for sunset timing.

sun-below D.d: Wait while the Sun is below D.d degrees altitude. Used for sunrise timing.

pos-above D.d HH:MM:SS DD:MM:SS: Wait while the position (α, δ) is above D.d degrees altitude.

pos-below D.d HH:MM:SS DD:MM:SS: Wait while the position (α, δ) is below D.d degrees altitude.

The 1st and time commands cannot pause for longer than 12 hours. If the time is within 12 hours in the past, the condition is considered to have been met. Of these commands, the 1st and sun-below/above commands are the most heavily used. These commands permit building surprisingly flexible observing scripts that can be used night after night and year after year. They also make it possible to write good scripts that do not need to know exactly how long everything takes. In general, the best scripts will have a small amount of deadtime in the night to allow for timing variations in the actual schedule. A well honed full-night script will usually be scheduled within a few minutes of the full duration of the usable night.

4.5. Predefined File Commands. There are a collection of predefined command files that encapsulate operations that are used every night. By using these files, centralized changes

and improvements can be made that will automatically be incorporated into future nightly observing plans. The `alopen` script takes care of waking up the telescope, updating the clock on the telescope computer, opening the dome, turning on dome and telescope tracking, setting the file name for the data to be taken that night, and enabling the dead-man timer. Some operations benefit from reading out a small fraction of the science array. The file `acqsubar` sets the size of the sub-array for target acquisition to a 3 arcmin square in the center of the CCD. A second file is `focsubar` and is used to set the size of the sub-array for focusing at 80 arcsec square. In addition to speeding up these operations, reading a sub-array greatly reduces the chance for confusion from nearby sources or cosmic ray strikes. At the start of a new robotic run the telescope coordinates must be verified. A set of 12 `bigacqXX` files are defined (XX is the right ascension for the target star) that perform a full-field acquisition on a very bright star and updates the pointing. Accurate focusing is handled by a 24 scripts, `focusXX-25`, where XX is the right ascension of the focus star. All stars are at a declination of $+25^\circ$. Each script has a star $V=7-8$ that has nothing brighter within 3 arcmin. This brightness gives a good signal-to-noise ratio image in one second. These scripts include all the operations to slew to the star, take the data, and adjust the focus accordingly.

Most observing programs on this system strive to collect absolute photometry requiring the use of all-sky photometric standard stars. The list of Landolt standards [1–3] includes areas that provide two or more standard stars at the same time and spaced at roughly one-hour intervals on the sky. There are separate command files for these fields with different combinations of filters such as *BV* or *VRI*. These files contain the target pointing needed to place the stars on the field and also contain the relevant exposure times needed for optimal signal-to-noise ratio images.

At the end of the night, the system must be properly stowed. This is also the best time to take calibration frames for this telescope. There are commands for shutting down the telescope either with or without taking calibration data. Each of these also includes an instruction that marks data collection complete for that night.

5. Essential Algorithms

This project involved some algorithmic development and testing that were not obvious prior to starting. This section describes the most important of these lessons learned with the hope they will be useful in other systems.

5.1. Flat Field Collection. An essential calibration operation is the collection of flat field images for which I generally use twilight sky images. The challenge of collecting twilight flats is that the illumination level is constantly changing, affecting the exposure time needed. The system has the most recent mean background signal level making it easy to adjust before the next image. When taking a set of flats for any filter there are seven control parameters for the command and the default value is shown in parentheses: number of frames

desired (20), the highest signal level permitted for a usable flat ($maxgood = 43000$), the highest signal level permitted for the optimal flat ($maxbest = 28000$), lowest signal level permitted for the optimal flat ($minbest = 18000$), lowest signal level permitted for a usable flat ($mingood = 10000$), maximum exposure time allowed ($maxexp = 20$ seconds), and minimum exposure time allowed ($minexp = 2$ seconds). These default parameters are the result of optimizing against the dynamic range of the detector, readout time, speed of the shutter blades, and sensitivity of the detector and filter combination. The dawn flat steps are as follows.

(1) Move telescope to the “Chromey” spot [4] and turn on tracking.

(2) Take a single bias frame with a 200×200 sub-frame in the center of the CCD. This signal level is the zero-illumination reference value.

(3) Take a 0.002 second exposure with a 200×200 sub-frame. If the signal level is too high no flats are possible and the operation terminates.

(4) Take a one second metering exposure with a 200×200 sub-frame until the signal would be above $minbest$ with a $maxexp$ exposure time. If still too dark, wait for a minute before trying this step again.

(5) Calculate exposure time. First compute the time need to have the signal level be $minbest$ but no longer than $maxexp$ and no shorter than $minexp$. This new time will predict a new signal level. If the predicted level is less than $mingood$ or greater than $maxgood$ quit taking images and log an error that the desired number of frames was not reached.

(6) Take an image and save it. Exit when the desired number of frames have been taken.

(7) Offset the telescope 20 arcsec to the east and go back to step 5.

The control settings allow optimization for the largest number of frames in the shortest time with a useful signal level. The number of frames is more important than the number of photons collected. The current system can get flats on 2 filters per twilight with 20 frames per filter. Dusk flats are vastly inferior to dawn flats for this instrument and telescope and are very rarely used. Generally collecting flats for a given filter once a week is sufficient.

5.2. Focus Determination. An automated system must be able to find the best telescope focus without user intervention. A completely general automatic focus algorithm that assumes nothing at the start proved impossible. As implemented, the telescope must be within $100 \mu\text{m}$ of proper focus to ensure success. In the early days it was common to interactively determine an approximate focus at the start of a new run. After sufficient operational history was reached this step was dropped in favor of a temperature based initial predictor. The focus shift between the end of the previous night and the start of a new night is larger than $100 \mu\text{m}$. Therefore, at the start of each night the initial focus value is set based on the current tube temperature and the known trend of focus versus temperature.

The automatic focusing procedure is based on a series of images, changing focus for each, with a default step size

of $10\ \mu\text{m}$, starting 9 steps below the current focus value and ending 9 steps above. Using the default 1-second exposure time this sequence takes roughly 80 seconds. The default step size is normally used only for the first focus of the night. After the initial focus, the step size is reduced to $6\ \mu\text{m}$ to get a more accurate focus. However, the sequence always includes $-90\ \mu\text{m}$ at the start and $+90\ \mu\text{m}$ at the end. These flanking focus settings ensure that the sweep always sees an image that must be out of focus. For each image a figure of merit (FOM) is calculated that is the value of the peak pixel signal divided by a small-aperture integrated signal. The best focus is the one with the largest FOM. However, the best FOM must be at least twice the worst FOM. If not, the sequence is flagged as indeterminate and the focus is left at its original position. This safeguard protects against taking focus data when there is no real data (e.g., clouds, bad pointing). Every focus run is recorded in a special focus log file that is designed to be machine readable and record all of the information available to the data collection program (except for the actual images). In this way, refinements to the focus algorithm can be tested against a very large dataset of past focus runs to ensure that any algorithmic change truly leads to an improvement in operation.

Table 1 shows an example focus sequence. The table shows the focus position, the FWHM of the brightest pixel (in pixels), the peak signal (above background), total flux in a 5-pixel radius aperture, and the FOM. The initial setting was $5520\ \mu\text{m}$, taken with a V filter, at $\alpha = 19 : 57 : 52.8$, $\delta = +24 : 29 : 18$, hour angle = E00 : 14 : 44, airmass = 1.02, tube temperature = 19.7C. Best focus was found to be at $5562\ \mu\text{m}$ (shown in italics). This method has proven to be very fast and exceptionally robust. When conditions are bad it will refrain from making a change to focus. When conditions are excellent it gives the best focus subject to the chosen step size. Under conditions of poor seeing it will still give a good answer. Other common techniques such as long-exposure images (>30 seconds), functional fitting to FWHM versus focus, and minimum in FWHM were all attempted but had conditions under which they failed and reached a decision that made things worse.

5.3. Target Acquisition. Data quality and ease of reduction is considerably enhanced by accurate target acquisition. This system provides three different acquisition strategies depending on the pointing accuracy needed.

Blind acquisition refers to the simplest method of target acquisition. In this case, the telescope is commanded to the desired coordinates and data collection begins without any attempt to verify the pointing. The raw pointing accuracy of the telescope after a new pointing model determination is less than 10 arcsec (1σ) across the entire sky. Fortunately, it is easy to do better and this mode is generally used only for things like setting up for twilight flats.

A *self-referencing acquisition* is used if the target position happens to be the position of a point source with no other brighter objects within 3 arcmin. Generally this is only used on sources that are in the range of magnitude 6–9. In this case the acquisition consists of the following steps.

TABLE 1: Focus sweep.

Focus (μm)	FWHM (pixels)	Peak (DN)	Flux (DN)	FOM
5430	8.9	3619	$5.11\text{e} + 05$	0.0071
5466	8.7	5537	$7.13\text{e} + 05$	0.0078
5472	8.2	4564	$6.98\text{e} + 05$	0.0065
5478	8.5	4437	$7.25\text{e} + 05$	0.0061
5484	7.4	6452	$7.36\text{e} + 05$	0.0088
5490	7.6	6429	$7.29\text{e} + 05$	0.0088
5496	7.7	7707	$7.48\text{e} + 05$	0.0103
5502	6.3	10073	$7.62\text{e} + 05$	0.0132
5508	7.4	7206	$7.50\text{e} + 05$	0.0096
5514	7.3	6341	$7.46\text{e} + 05$	0.0085
5520	7.0	7535	$7.50\text{e} + 05$	0.0100
5526	5.4	13066	$7.53\text{e} + 05$	0.0174
5532	5.1	13832	$7.58\text{e} + 05$	0.0183
5538	5.6	11557	$7.50\text{e} + 05$	0.0154
5544	4.8	16368	$7.50\text{e} + 05$	0.0218
5550	5.4	12246	$7.44\text{e} + 05$	0.0165
5556	4.9	15461	$7.46\text{e} + 05$	0.0207
<i>5562</i>	<i>4.4</i>	<i>19354</i>	<i>7.65\text{e} + 05</i>	<i>0.0253</i>
5568	4.7	17501	$7.55\text{e} + 05$	0.0232
5574	5.7	11007	$7.48\text{e} + 05$	0.0147
5610	5.9	11352	$7.45\text{e} + 05$	0.0152

- (1) Slew the telescope to the nominal coordinates.
- (2) Take a short exposure (typically 1 second) with the 3-arcmin acquisition field setting.
- (3) If the brightest source in the image is too weak or has a FWHM less than 1.5 pixels the acquisition is aborted.
- (4) Given a good image, the telescope performs the small offset needed to put the source in the center of the CCD.
- (5) After the offset is completed the telescope pointing is updated to match the known coordinates of the source.

A *local-reference acquisition* is used when the desired pointing does not coincide with an actual target or if the target is too faint. In this case, acquisition becomes a two step process. First, a self-referencing acquisition is done on the nearest 8-9 magnitude star. This leaves the telescope nearby with very accurate coordinates. From there the final step is a blind acquisition of the requested position. If the local reference star is within a few degrees this method will get to within 1-2 arcsec of the desired location. The extra time for the local reference is generally quite small (<10 seconds).

5.4. Scheduling Standard Stars. Thinking about how to write software to build a night's observing session led to more than a few interesting revelations. One of the core lessons was how much more universal a timeline of observations is when it is laid out relative to the local sidereal time (LST). Using LST to schedule observations of standard stars is particularly effective since the airmass is known and constant for a given LST. Common practices developed with single-channel photoelectric photometry recommend observations of standards periodically spaced through the night that

TABLE 2: Standard star observing blocks.

LST	Horz	Low Field	High Field	Extra Field
00:47	west	SA112	SA114 (1.40)	
01:46	west	SA113	PG0231 (1.17)	SA114 (1.72)
02:18	west	PG2213	PG0231 (1.15)	SA114 (2.03)
02:48	west	SA114	PG0231 (1.15)	
03:23	east	Rubin149	PG0231 (1.17)	
04:14	east	M67	SA95 (1.22)	
05:00	west	SA92	SA95 (1.27)	
05:10	east	PG0918	Rubin149 (1.48)	
05:55	east	SA101	Rubin149 (1.33)	PG0918 (1.86)
06:47	east	PG1047	PG0918 (1.50)	
06:53	west	PG0231	Rubin149 (1.24)	SA95 (1.71)
07:57	west	SA95	M67 (1.11)	
09:54	east	PG1323	PG1047 (1.26)	
11:26	west	Rubin149	PG1047 (1.23)	
12:04	east	PG1633	PG1323 (1.49)	
13:29	west	M67	PG1323 (1.38)	SA101 (2.02)
13:33	west	PG0918	PG1323 (1.38)	
14:00	west	SA101	PG1323 (1.40)	
14:38	east	SA110	PG1633 (1.25)	
14:53	west	PG1047	PG1323 (1.50)	
16:39	east	SA112	PG1633 (1.10)	SA110 (1.41)
16:58	west	PG1323	PG1633 (1.11)	
17:36	east	SA113	SA110 (1.27)	SA112 (1.77)
18:14	east	PG2213	SA110 (1.22)	
18:35	east	SA114	SA110 (1.21)	
20:49	east	SA92	SA112 (1.22)	
21:08	west	PG1633	SA112 (1.22)	
22:14	east	PG0231	SA114 (1.21)	
22:48	west	SA110	SA114 (1.20)	SA112 (1.42)
23:49	east	SA95	PG0231 (1.50)	

are taken at a wide range in airmass. Standards for CCD photometry are much more effective if fields containing two or more standard stars are used. The Landolt standard fields mentioned in Section 4.5 are a set of 15 fields equally spaced along the equator. The choices were also influenced by requiring high-quality standards with a range of star color in the field.

Given this list of fields it is a simple matter to compute for each field the two LSTs when it is at 2.5 airmasses. At each LST, compute the airmass and hour angle of all the other fields. From this list, keep the field closest to the meridian that is on the same side of the meridian as the high airmass field at that time. These two fields are a low/high airmass field pair. Some of the pairs have a third well-placed field that is intermediate in airmass. The extra field is optional but if included will improve the quality of the calibrations. Table 2 provides a list of calibration opportunities ordered by LST for 35° north or south latitude. The column labeled “Horz” is short for horizon and indicates if the first field is rising in the east or setting in the west. The “Low Field” column is the shortened name of the field always at an airmass of 2.5. The

“High Field” is the field nearest the meridian. An extra field is also included if appropriate. The value in parentheses is the field’s airmass at this LST.

These fields and the temporal framework greatly simplify the process of sequencing a night either interactively or with software. For a given night you first eliminate all opportunities that do not fall within the observing window. Next, remove all sets involving fields that are too close to the moon ($<30^\circ$ for this system). As the science program is built for the night add these calibrations to the timeline at roughly a two hour spacing taking special care to put a calibration as close as possible to the start and end of the night.

5.5. Pointing Model. All telescopes, whether used robotically or classically, are much more effective when they can accurately point to a desired location on the sky. Most, if not all, telescopes use some method to map from raw telescope coordinates to sky coordinates and is used to take out systematic errors introduced by the mechanical system. The Lowell Observatory telescopes use the Wyoming analytic model [5] for this mapping. Robotic systems are particularly effective at collecting data needed to derive such maps and involves taking a set of images over the entire accessible sky. An astrometric solution for the image reveals the actual pointing while the header records where the telescope thought it was pointing. The current pointing model is removed from the header coordinates to provide the original raw telescope coordinates. These pairs of values can then be used to refine the pointing model. A normal pointing run is a grid of about 120 points uniformly spread over the sky and takes about 3 hours to collect. A regular pattern in either hour angle and declination or altitude and azimuth should be avoided since they often generate excess data at either the pole or the zenith. The improved data quality made possible by robotic means made it possible to find systematic limitations in the Wyoming model when applied to the Lowell telescopes. This telescope turns out to need three additional terms in the pointing model where the new independent variables are (1) hour angle, (2) hour angle times declination, and (3) hour angle squared. The addition of these three terms reduced the pointing error from a few arc-minutes to a few arc-seconds.

6. Safety

An unattended system has a very different risk profile to one where there are always on-site personnel. The differences between the two modes can often be subtle and yet surprisingly important. A significant challenge in developing this autonomous mode was in recognizing these differences and reducing or eliminating risk through software, hardware, and good safety practices. First and foremost is the need to maintain a safe environment for technical personnel that are on site from time to time for repairs or routine maintenance. The second priority is to safeguard the hardware.

6.1. Personnel Safety. The only risk posed by this system to personnel is through the moving components (telescope and

dome) that are influenced by computer control. Minimizing these risks therefore involves controlling when and how the computer can move the telescope. It is set by policy that the nighttime is the province of the system and going into the dome should be done in consultation with the cognizant observing supervisor. During the day, the system is supposed to be idle and if not, the activities must be cleared through the technical staff. A policy statement has provided sufficient protection since the staff are rarely onsite unless there is a problem in which case everyone is involved anyway. Most of the accommodations for safety are surprisingly transparent to the staff. For instance, the system was built to tolerate power failures during the day, a common step during maintenance. It is unlikely that the facility is a completely safe environment but the goal was to make sure the new mode of operation did not add any risk that was not already present.

6.2. Hardware Safety. The nature of working with hardware is that it does not always work as expected. Considerable attention was paid to protecting the telescope and dome from harm. Despite best efforts in software no system can be 100% reliable and understood by its software and notification is the last line of defense in case of error. Any error generated by the telescope is treated as a serious problem. In the early days there was a daemon that collected error notification requests and pass them along to a pager. It resent the page every 10 minutes until confirmation ensuring that someone saw the message. After a couple of years the spectrum of system errors changed enough that any error was unlikely to have a meaningful remedy during the night. With that realization the pager was retired and all notifications were relegated to email.

The telescope has a safety feature known as the dead-man timer. This is a mode, activated at night for robotic observing only, that causes the MOVE31 computer to keep track of command activity. Each time a command is received the time is reset to 0. If the timer reaches two hours the telescope will automatically stow, shut down tracking, and close the dome. This is largely a protection against a catastrophic failure of cmdr, the computer it runs on, or the RS232 communication link.

The most important design philosophy for hardware safety is that the system must not assume anything. Most, if not all, errors can be traced back to an assumption being made in software. For instance, the telescope control software (MOVE31) used to assume that if it was not commanding the dome it could not be moving. When faced with a hardware problem such as a stuck power relay to the dome rotation motor this assumption becomes part of the problem. Clearly, there was little harm in having the system occasionally poll the dome position while stationary and if it were seen to move it would throw an error. In this case there is nothing the software can do directly but sending an email messages vastly increases the chances that someone can intervene and deal with the problem.

The other important lesson learned is that errors, once seen, do not go away and must be understood and fixed, never deferred for later. In most cases the log files permit

reconstructing all elements of the problem and allow tracing through the software to find the problem. In cases where this was not possible it invariably led to finding some error condition that was not properly logged. By adding new logging code such problems can eventually be solved.

The final safety issue for the hardware relates to the weather. In the case of Anderson Mesa, the problem is not too difficult since the weather conditions rarely change on a timescale short compared to a night. A weather forecast in the afternoon can generally be relied on to indicate if there will be precipitation or fog during the night. If there is a chance of either, the telescope is not opened at all. This site does not suffer from problems caused by high winds during a workable night. All of these weather issues mean that there is no need for real-time decision making based on current weather conditions. A go/no-go decision just before sunset is sufficient to eliminate weather related risks.

7. Routine Operation

A typical night is handled by a single script that is intended to fill the entire night. Hand generated scripts are only feasible if the same observing pattern is used night after night. Machine generated scripts are much less prone to error and can change from night to night.

The night always starts with a command to wait for the sun to set. Once that happens, the dome is opened and all the housekeeping steps needed to bring the telescope and camera to an operational state are executed. If the night is the first robotic night a full-frame pointing verification is performed very early during twilight. Next, wait for the sun to get 9° below the horizon and then do a first focus update. Every focus command is preceded by a synchronization command that keeps the command from executing too early. At this point the system is ready to begin observing. Science observations are interspersed with the necessary focus (once an hour) and standard star observations (every two hours).

At the end of the night twilight flats and bias frames are collected, if desired. The flat field command begins after the sun reaches an altitude of -9° . Once flats are collected the dome is closed, the telescope is parked, and then bias frames are collected. With PCCD, darks were also taken after the dome was closed and before the bias frames were taken. If no flats are taken the system usually still waits until the sun reaches -9° altitude and then shuts down. This timing marker makes it easy to spot wasted time in the schedule. At the very end a marker file is written with the data to indicate that data collection has ceased for the night. Once that marker propagates to the pipeline processing computers the night can be processed.

8. Example Data

8.1. Pluto Lightcurve Monitoring. The project that led to the development of robotic observing is lightcurve monitoring of Pluto. Figure 4 shows one example lightcurve from 2000 [6] in the last year before the robotic system was functional. The project requires a complete lightcurve each year to

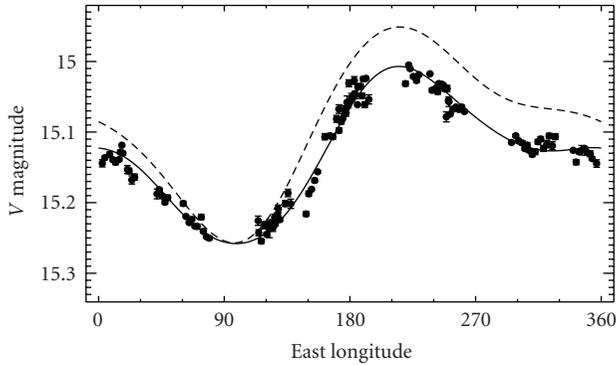


FIGURE 4: Lightcurve of Pluto in March-July 2000. The rotation period of Pluto is roughly 6.4 days and requires many nights to obtain a complete lightcurve. The points with error bars are the result of 35 nights of observation and nearly 1000 individual images. The solid curve is a 4-term Fourier series fit to the lightcurve. For comparison the lightcurve from 1986 is shown (dashed line).

document the system brightness during the current period of rapid change in the illumination geometry. The amount of effort required to collect data like this became prohibitively expensive for a long-term project.

8.2. Deep Ecliptic Survey Calibration. The Deep Ecliptic Survey [7–9] was an observational project to discover 1000 new trans- and ultraneptunian objects. The project used wide field imaging cameras on the Kitt Peak Mayall 4-m telescope and the Cerro Tololo Blanco 4-m telescope to survey the sky within 5° of the ecliptic. The robotic system has been working on calibrating those images by obtaining absolute photometry over most of this ecliptic zone. The survey has over 1300 primary fields and 1200 secondary fields that need supporting photometry requiring roughly 250 photometric nights to complete. This work is nearing completion and already has a photometric database of 515,000 observations from the first generation (PCCD) camera and 3.2 million observations from the current (NASAcam) camera. This observing project is supported by a dedicated single-purpose scheduling program that keeps track of which fields have already been done and builds a complete night of observations in a few seconds.

8.3. Astrometry. The current camera is an excellent instrument for making astrometric observations. The optical prescription was optimized to minimize image quality (PSF) variations over the field of view. There is a barely measurable quadratic distortion term required for fitting an astrometric control network to the images. In all cases, the astrometry from this system is limited by the quality of the supporting astrometric catalogs. To date, this system has been used to collect astrometry in support of occultation predictions, asteroid orbit improvement, and even astrometry of the background stellar field that the New Horizons spacecraft will see as it approaches Pluto in 2015. Often these astrometric observations can be slipped into the observing sequence with minimal disruption and with little advanced warning.

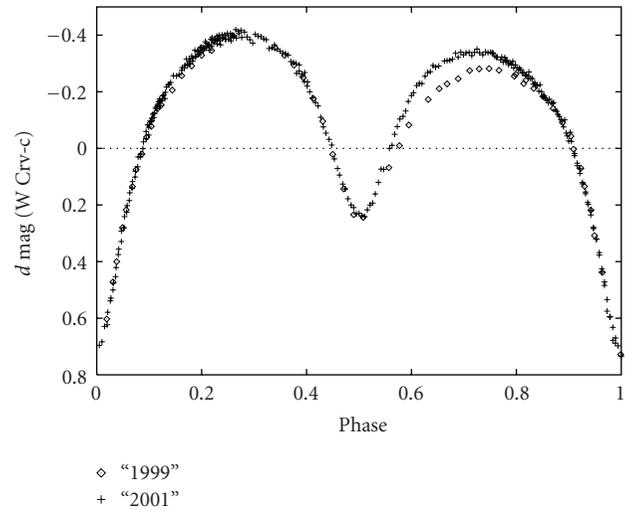


FIGURE 5: Lightcurve of W Corvi in 1999 and 2001 with the 0.8-m telescope. The 1999 data (diamonds) were taken with a classically operated CCD system in V . The 2001 data (+) were taken with the PCCD camera and the robotic system in $BVRI$ (only V data are shown). Figure used with permission from Odell [10].

8.4. Asteroid Lightcurves. A few test observations were taken to investigate the utility of this system for collecting asteroid lightcurves and near-earth asteroids in particular. The system supports observations of moving targets and is especially good at nearby targets where the topocentric corrections during the night are large.

8.5. Variable Stars. The simplest observations of all are photometric observations of variable stars where the photometry can be done relative to field stars. These can be done on non-photometric nights for which there is more time available than is typically used. Figure 5 shows a lightcurve of W Corvi collected with this system [10]. The fast filter wheel makes it very easy to collect multi-color observations of systems like this as well. As with all of the example projects, reducing the cost of data collection makes it much easier to collect long time base observations where the scientific breakthrough comes from a lifetime of monitoring rather than a single night of data.

8.6. Education. This system was used for a test project with the University of Hawaii to involve students in secondary education. This project made time available for students, typically in support of science fair projects, where the student picks a targets and an observing pattern. The data are collected on their behalf, calibrated, and sent to the teachers and students involved. The cost of supporting these educational projects is trivially small but has a profound impact on the students that are involved. We have worked together to build the process by which students can directly request the observations that feed into an automated scheduler but there has been insufficient funding thus far to complete this project.

9. Summary

Here are some of the benefits of this autonomous observing system:

- (i) Greatly reduced costs of data collection
- (ii) Fatigue free “observer”
- (iii) Much, much less prone to data collection errors
- (iv) Perfect data headers (simplifies reductions)
- (v) Zero-cost dawn twilights that improved data quality
- (vi) Extensible to much more elaborate observational patterns
- (vii) Eliminates all decision-related dead-time during the night.

In many ways this mode is the lifeblood for small telescopes everywhere. Robotic operation can maintain the scientific relevance of smaller apertures in cases where observing operations costs dominate.

Acknowledgments

This work was supported in part by NASA grants, NAG5-13380, NAG5-12229, NNG06-GI23G, NAG5-12835, NNX09-AB43G. Thanks to the technical staff at Lowell Observatory for their dedication to the Observatory telescopes. Special thanks to Larry Wasserman for his help above and beyond the call of duty on special modifications to the telescope control software that made this project a productive reality. David Tucker and REU student Brian Keeney helped with the early prototyping work on this project. Peter Collins was invaluable during the conversion process to NASAcam and also helped with crucial debugging and system oversight and monitoring.

References

- [1] A. U. Landolt, “UBV photoelectric sequences in the celestial equatorial selected areas,” *Astronomical Journal*, vol. 78, pp. 959–981, 1973.
- [2] A. U. Landolt, “UBVRI photometric standard stars around the celestial equator,” *Astronomical Journal*, vol. 88, pp. 439–460, 1983.
- [3] A. U. Landolt, “UBV RI photometric standard stars in the magnitude range $11.5 < V < 16.0$ around the celestial equator,” *Astronomical Journal*, vol. 104, no. 1, pp. 340–371, 1992.
- [4] F. R. Chromey and D. A. Hasselbacher, “The flat sky: calibration and background uniformity in wide-field astronomical images,” *Publications of the Astronomical Society of the Pacific*, vol. 108, no. 728, pp. 944–949, 1996.
- [5] E. J. Spillar, D. Dumbrell, G. L. Grasdalen, and R. R. Howell, “The Wyoming infrared observatory telescope software system,” *Publications of the Astronomical Society of the Pacific*, vol. 105, pp. 616–624, 1993.
- [6] M. W. Buie and W. M. Grundy, “Continued evolution in the lightcurve of Pluto,” *Bulleting of the American Astronomical Society*, vol. 32, p. 1083, 2000.
- [7] R. L. Millis, M. W. Buie, L. H. Wasserman, J. L. Elliot, S. D. Kern, and R. M. Wagner, “The Deep Ecliptic Survey: a search for Kuiper belt objects and Centaurs—I: description of methods and initial results,” *Astronomical Journal*, vol. 123, no. 4, pp. 2083–2109, 2002.
- [8] M. W. Buie, R. L. Millis, L. H. Wasserman, et al., “Procedures, resources and selected results of the Deep Ecliptic Survey,” *Earth, Moon and Planets*, vol. 92, no. 1–4, pp. 113–124, 2003.
- [9] J. L. Elliot, S. D. Kern, K. B. Clancy, et al., “The Deep Ecliptic Survey: a search for Kuiper belt objects and Centaurs—II: dynamical classification, the Kuiper belt plane, and the core population,” *Astronomical Journal*, vol. 129, no. 2, pp. 1117–1162, 2005.
- [10] A. P. Odell, *International Bulletin Variable Stars* no. 5514, 2001.