The Astrophysical Journal, 605:846–853, 2004 April 20

A CCD TIME-SERIES PHOTOMETER

R. E. Nather and Anjum S. Mukadam
Department of Astronomy, University of Texas, Austin, TX 78712; nather@astro.as.utexas.edu, anjum@astro.as.utexas.edu
Received 2003 November 4; accepted 2004 January 9

ABSTRACT

We describe a high-speed time-series CCD photometer for the prime focus of the 82 inch (2.1 m) telescope at McDonald Observatory and summarize the observational results we have obtained since it was placed into regular use in 2002 February. We compare this instrument with the three-channel time-series photometers we have previously used for the asteroseismological study of pulsating white dwarf stars, which used photomultiplier tubes (PMTs) as the detectors. We find the CCD instrument is about 9 times more sensitive than the PMT instruments used on the same telescope for the same exposure time. We can therefore find and measure variable white dwarf stars some 2.4 mag fainter than before, significantly increasing the number of such objects available for study.

Subject headings: instrumentation: photometers — stars: imaging — stars: oscillations — techniques: image processing — techniques: photometric — white dwarfs

1. INTRODUCTION

We have designed and placed into operation a CCD camera system optimized for high-speed time-series measurements of oscillating white dwarf stars. Our experience indicates that CCD instruments designed for more general use will have characteristics that are unacceptable for rapid time-series measurements or will seriously compromise the data quality when used for this purpose. Until this instrument was available, we relied on time-series photometers using photomultiplier tube (PMT) detectors for these measurements (Nather & Warner 1971; Kleinman, Nather, & Phillips 1996). We are now in a position to compare and contrast these two approaches to the same measurements and to demonstrate where CCDs are superior for this work and where they are not.

One basic goal in the design of this instrument is to take advantage of the improved quantum efficiency that CCD detectors offer, so that we can obtain usable data on faint stars that the PMT instruments can not measure and data of better quality on those (i.e., less-faint stars) they can. We also hoped to discover additional variable white dwarf stars to increase the limited number of objects available for asteroseismological study. A more long-term goal was to provide an instrument whose timing accuracy was high enough to allow a search for small deviations from the smooth secular change in pulsation frequency due to white dwarf cooling (e.g., Kepler et al. 2000), opening the possibility that such deviations, if found to be periodic, could demonstrate the presence of planet-sized objects in orbit around the white dwarf star (Mukadam, Winget, & Kepler 2001; Winget et al. 2003; Mullally et al. 2003).

We have achieved our short-term goals in the 15 months the instrument has been in operation and have established a list of target objects for the longer term search for primordial planets. Our search for white dwarf variables has already doubled the number available for study (Mukadam et al. 2003a), primarily because of increased instrumental sensitivity. Our practical limit using the PMT photometers on the 82 inch (2.1 m) telescope at McDonald Observatory was $\sim$17.0 mag; we now obtain light curves of comparable quality at 19.4 mag, a gain of about a factor of 9 in overall sensitivity. To obtain a similar gain with the PMT photometers, they would have to be attached to a telescope 6 m in aperture.

2. HARDWARE

2.1. *The CCD Camera*

Our CCD time-series photometer, Argos, is based on a commercial CCD camera made by Roper Scientific, the Princeton Micromax 512 BFT NTE-CCD.[1] Its specifications are shown in Table 1.

The CCD chip is back-illuminated to improve its blue sensitivity (most white dwarf stars are blue) and can transfer its 512×512 pixel image to the on-chip buffer in 310 μs. We have built a mount to support the camera at the prime focus of the 82 inch telescope (f/3.9), obtaining an image scale to match the 13 μm $\times$ 13 μm pixel size: 3.05 pixels arcsec^{-1}. The field of view for this image scale is $2'.8$ on a side, large enough that we have not had any trouble finding our targets and suitable comparison stars. Since the target can be placed almost anywhere in the chip, the usable area to search for comparison stars is $\sim$25 arcmin^2.

The camera incorporates a thermoelectric cooling system that keeps the chip at -45 °C, where the dark count of 1–2 ADU s^{-1} is smaller than the counts coming from the moonless sky ($\sim$3–7 ADU s^{-1}). The readout noise of 8 electrons rms is negligible for all except the shortest exposure times, for which it is comparable to sky noise. The image readout time of 280 ms is comfortably shorter than our minimum exposure time of 1 s. The prime-focus mount design includes a manual two-position filter slide, which can double as a dark slide when we want to take dark or bias frames. The camera came with an internal shutter for this purpose, but we removed it when it proved to be unreliable. The system layout is shown in Figure 1.

The CCD camera connects to the ST-133 controller (“A/D” in Fig. 1) via a 10 foot (3 m) analog cable, both of which are mounted at the prime focus of the 82 inch telescope. The interface card housed in the camera control computer (see § 2.6) connects to the controller via a 75 foot (23 m) digital

[1] See http://www.roperscientific.com/micromax.html.

TABLE 1
SUMMARY OF CAMERA SPECIFICATIONS

Attribute	Specification
Pixel size	13×13 μm
Pixel array size	512×512, back illuminated
On-chip storage	512×512, frame transfer operation
Frame transfer time	310 μs
Readout rate	1 MHz, 16 bit A/D conversion
Readout time	0.28 s, full frame with no binning
Cooling	Thermoelectric + fan air exhaust
Chip temperature	-45 °C
Readout noise	8 electrons rms
Gain	2 electrons ADU^{-1}
Dark noise	1–2 ADU s^{-1} pixel^{-1}
Optical coating	Broadband antireflection
Quantum efficiency	30% at 3500 Å, 80% 4500–6500 Å, 40% at 9000 Å
Linearity	~1% below 40,000 ADU (saturation at 65,000 ADU)

high-speed communication cable. Frame transfer operation is started by a single synchronizing pulse from a timer card ("Divide by n" in Fig. 1), which serves to end one exposure and begin the next one; the pulse is sent to the camera via a 75 foot coaxial cable.

2.2. *Wavelength Response*

While the CCD chip is more sensitive than the bi-alkali PMT detectors, its wavelength response is different enough that combining CCD and PMT data on the same star can not be done directly. Our chief interest lies in blue pulsating DA white dwarfs, whose pulsation amplitudes are a function of wavelength (Robinson et al. 1995; Nitta et al. 1998, 2000). Including photons redward of the PMT cutoff (ca. 650 nm), which are less modulated by the pulsation process, reduces the measured amplitude (Kanaan et al. 2000). We thought we might improve the signal-to-noise ratio in our light curves by inserting a blue filter (1 mm Schott BG40), and indeed found the measured amplitudes to be higher by about 35%–42%. However, photon losses in the filter increased the measurement noise, and we found no significant improvement in the signal-to-noise ratio by using it.

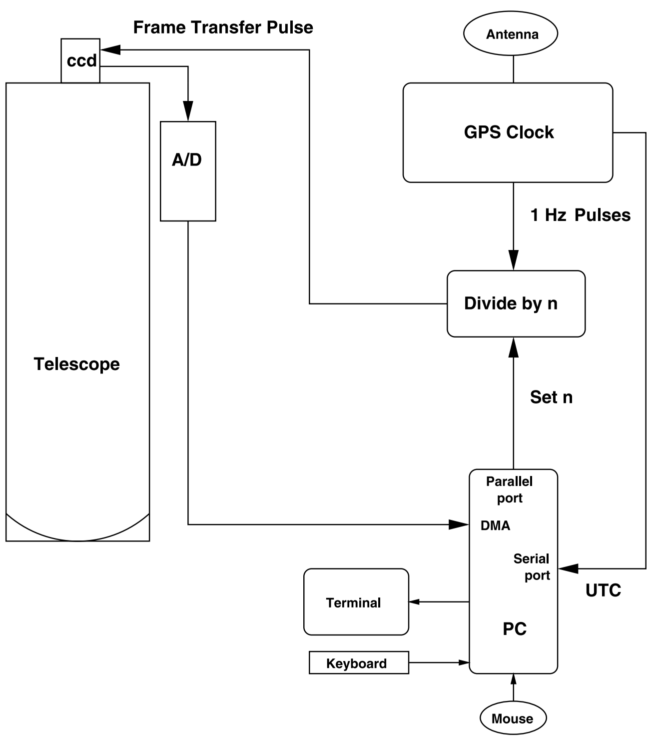

FIG. 1.—Major components of the Argos camera system and their interconnections

2.3. *The Prime Focus Mount*

The camera mounting plate can be moved by *x-y* adjustment screws over a distance of half an inch to center the CCD chip on the optical axis of the telescope. The tip/tilt of the camera can be controlled by a push-pull arrangement of screws to set the CCD chip perpendicular to the axis. The alignment procedure is adequate but awkward. When properly aligned, the corners of the chip are 2′ from the optical axis, where aberration from coma due to the parabolic primary mirror is calculated to expand a point image to about 1″ in diameter. We rarely experience subarcsecond seeing at McDonald Observatory, so we have not been able to verify this calculation.

2.4. *Baffling Argos*

Scattered light was initially a significant problem. The first three nights of the commissioning run proved beyond doubt that the shiny aluminum surfaces of the mount had to be darkened; we chose hard black anodizing for the purpose, as it does not corrode easily.

The single crude baffle in the original design was replaced with a five-stage baffle system consisting of two thin baffles very close to the camera and three other baffles in the body of the mount. The two camera baffles and the mount baffle closest to the camera have square-shaped apertures with rounded corners and are derived by projecting the light beam backward from the CCD chip. These openings are a few percent larger than the converging light beam from the primary. The other two mount baffles have circular apertures, which are 5%–7% bigger than the light beam. The edges of all the light baffles are at an angle of 45° with respect to the optic axis to reflect light away from the CCD camera.

The original flat-field images were very strange but have now been improved so that people no longer laugh at them. Our current images are flat to within a few percent; the variations in the flat field come from structural nonuniformities in the CCD chip itself. This pattern is stable and can be removed, giving us residual variations of less than 1%.

2.5. *The Timing System*

A time-series photometer must know precisely when an exposure is started, and precisely how long it takes. These are two different requirements, involving both a time epoch and an interval. The Argos timing system is based on a Global Positioning Signal (GPS) clock designed primarily for precision timekeeping, although position information is also available. It consists of an oven-controlled crystal oscillator disciplined by signals obtained from a GPS receiver.[2] The time epoch is claimed to have an error of about 50 ns, considerably more precise than we need, but comforting.

We have assembled a simple countdown register on a small circuit card that can accept the 1 Hz timing pulses from the GPS clock and can provide an output pulse to initiate frame transfer in the CCD camera. The exposure intervals are thus contiguous and determined directly from the clocking hardware. The timer card plugs into the parallel port on the camera control personal computer (PC), so the countdown value (i.e., the exposure time) can be set into it by software. Thereafter, it operates independently, initiating frame transfer operations at the established timing intervals. Exposure times can be set to any integral number of seconds from 1 to 30.

Immediately following a 1 Hz timing pulse, the GPS clock provides information from which the precise epoch of the pulse can be determined. The information arrives encoded in packet form at the serial port of the PC at 9600 baud, where it can be read and decoded by the software control program. The PC serial port is buffered, so that the epoch (the time and date) for each pulse can be determined, even if a second packet arrives before the first one is read by the program.

A second clock, somewhat less accurate, is available if the camera control PC is running the Network Time Protocol (NTP)[3] software, which obtains timing information over the Internet by periodically contacting time servers and adjusting the PC system clock accordingly. It allows for internet time delays as well as it can and averages the best readings it finds to keep the system clock in proper synchronization. The Argos control program relies primarily on the GPS clock for timing but can use the NTP-disciplined system clock if the GPS time signals are not available. For observer assurance, it compares the time ticks from the two clocks and displays their time difference in a status display window. After both clocks have been running for a few hours, the time difference is usually within a few milliseconds.

2.6. *The Camera Control Computer*

The PC that controls the camera has fairly modest requirements by modern standards: it must have a PCI bus to accept the camera control card, a parallel port (for the timer card), a serial port (for the GPS packet information), enough memory to run the Linux operating system comfortably (256 Mbyte is enough, but more is always better), and enough disk space to hold the images (527 kbyte each) as they arrive. Our current camera control PC runs a Pentium III at 1 GHz and is not pushed for time. The software prefers a display resolution of 1280×1024 so that the various windows do not overlap each other. A 17 inch (43 cm) liquid crystal display (LCD) works fine. The PC also needs an ethernet card to connect to the Internet, so the NTP software can discipline the system clock, and to receive pointing information from the computer that controls the 82 inch telescope. This connection is also used to transfer image data to our Argos data archive in Austin (slowly) and to allow a remote login to run the camera for testing purposes (even more slowly).

We usually operate a second PC as well, with access to the disk on the camera control computer, so that arriving image data can be examined by software not concerned with the data acquisition and recording process. We also make a more durable copy of the data on CD-ROM. Someday, when the DVD format wars are over, we may move to that medium to minimize the number of disks required.

3. SOFTWARE: THE USER'S VIEW

3.1. *Program Requirements*

The control program is called Quilt 11 (q11), the most recent in a series of programs designed to control time-series photometers. The name was originally chosen because the first of the series, Quilt 1, started as a patchwork of software routines. It was written in 1970 and has undergone 10 complete rewrites in different languages for different computers since that time. The basic operations have not really changed much.

Once the camera has been set to operate in frame transfer mode, images arrive via direct memory access (DMA) at the end of each exposure and appear magically in memory. The program must first associate each image with its epoch (start time and date) before it is written to disk. This is not quite as straightforward as it sounds: a new exposure starts when a frame transfer operation finishes so that the epoch is available right away, but obviously the image is not—it is still being exposed. It only shows up after the next timing pulse (and its epoch) arrives, and then only after the readout process has finished. The program must keep all this straight so the proper epoch is associated with the appropriate image.

The data images are recorded in the Flexible Image Transport System (FITS)[4] format, with the epoch and other operating parameters in the header. Each image has its own file, so the file name must be generated automatically and the names must be sequential so they can be kept in proper time order.

3.2. *Controlling q11*

Figure 2 shows what the computer "desktop" looks like with the program in operation. The window labeled "Commands" accepts simple typed keyboard commands (go, stop, abort, etc.) to control the camera, as well as commands to edit entries in the window labeled "Setmode," which provides the q11 program with information and parameters that it cannot determine for itself. The user may then mark the target and comparison stars on the image (small, labeled circles appear in response to mouse clicks) to enable on-line extraction and display of plotted light curves.

In our design, the images arrive whenever a frame transfer pulse is generated by the clock—that is, all the time. These images are shown on the display screen by a program called DS9, written by William Joy and his colleagues at the Smithsonian Astrophysical Observatory (SAO) and made

[2] See http://www.trimble.com/thunderbolt.html.
[3] See http://www.ntp.org.
[4] See http://fits.gsfc.nasa.gov/fits_intro.html.

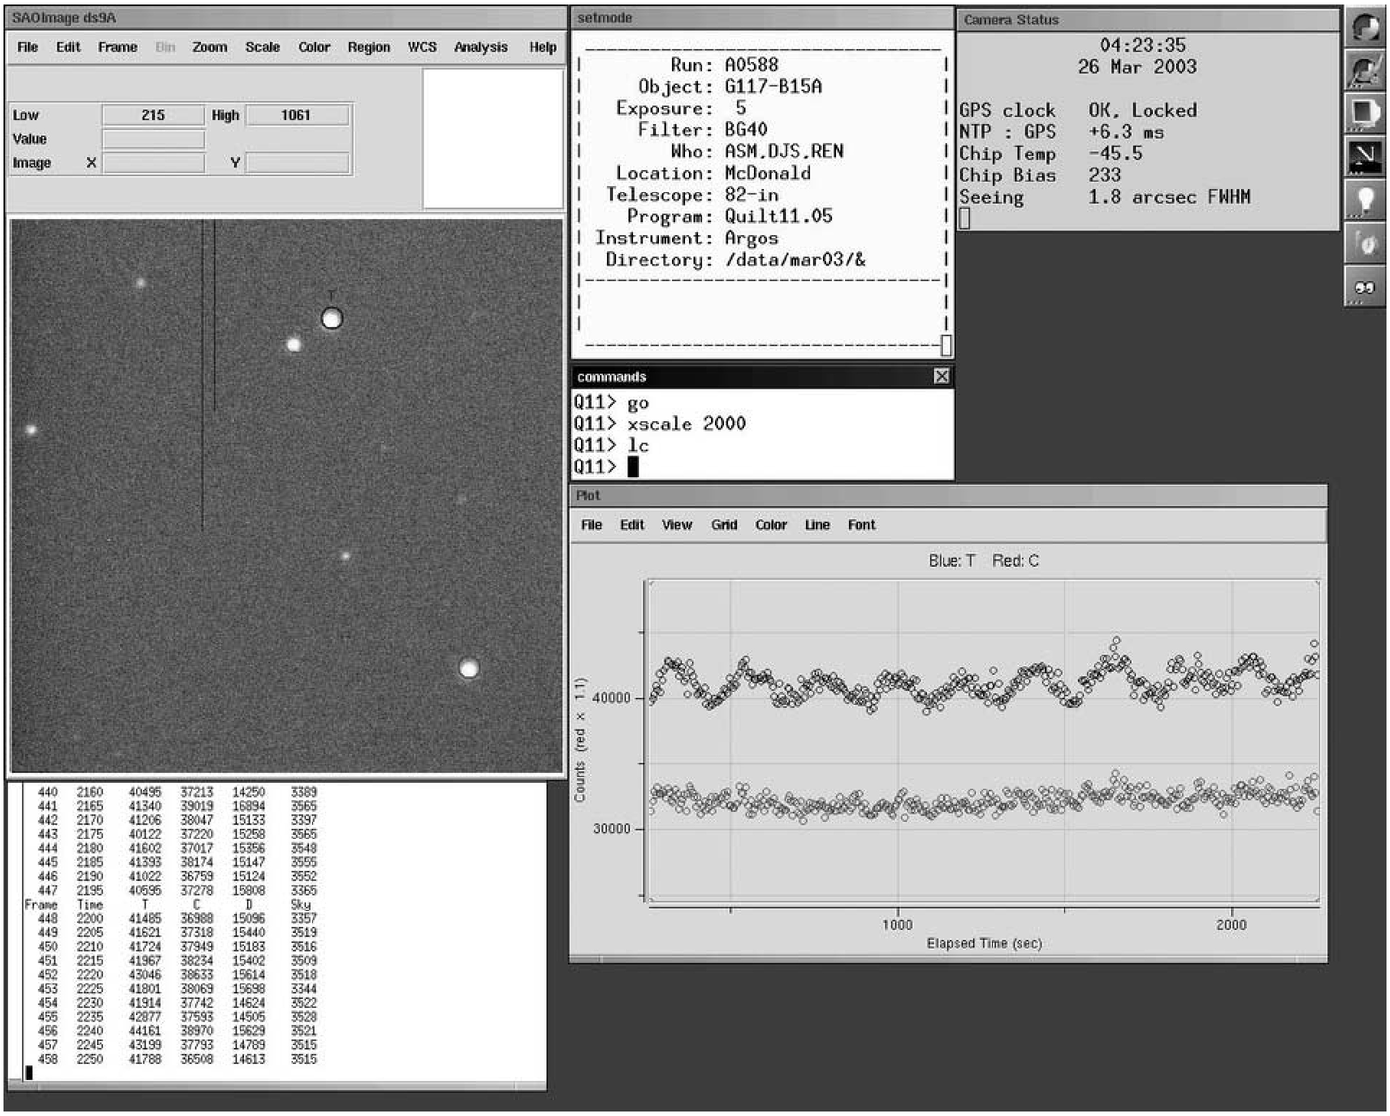

FIG. 2.—User's view of Quilt 11, the data acquisition program. Although shown here in black and white, the display is in color, with the upper light curve in blue and the lower one (the comparison star) in red.

available as Open Source Software. Plotted light curves are displayed by the DS9 plotting widget.

Once data recording starts (in response to the "go" command), the observer must act as an aide to the program to do things it cannot do for itself: keep the star images on the chip, keep the dome out of the light path, and be prepared to shut things down if it rains. In addition to plotting the light curve, the program also prints columns of numbers on the terminal used to start the program: the image number, the elapsed time, and the extracted brightness for each of the marked stars with sky removed. The last column shows the average sky value that was subtracted from the target star.

3.3. *Simulation*

The program can also be run in simulation mode, without a telescope, camera, or GPS timing system. Previously recorded data images are read from disk into the same buffer used by the camera, and the timing is simulated from the system clock ticks. Users can use this mode to review data previously recorded and to learn how the program works. The same code is used in both modes, with very few exceptions. If the program is started with a command-line path name to a data run, simulation mode is assumed; without one, it assumes that it must run the camera for real. Simulation mode goes through all the motions involved in a real run but does not write anything to disk.

3.4. *Software Design*

The q11 program is written in the C language and consists of five separate executable processes in simultaneous execution. It is designed to run on a PC under the Linux operating system and to tolerate the presence of other programs running at the same time on the same CPU. Both incoming time and image data are buffered to maintain proper real-time operation. Details of its architecture and the on-line extraction algorithms are presented in the Appendix.

3.5. *Display*

The data acquisition and online extraction processes run in a thread separate from the display process because of timing considerations: the display routines (DS9 and its plotting widget) are written in Tcl/Tk, an interpreted language, and are

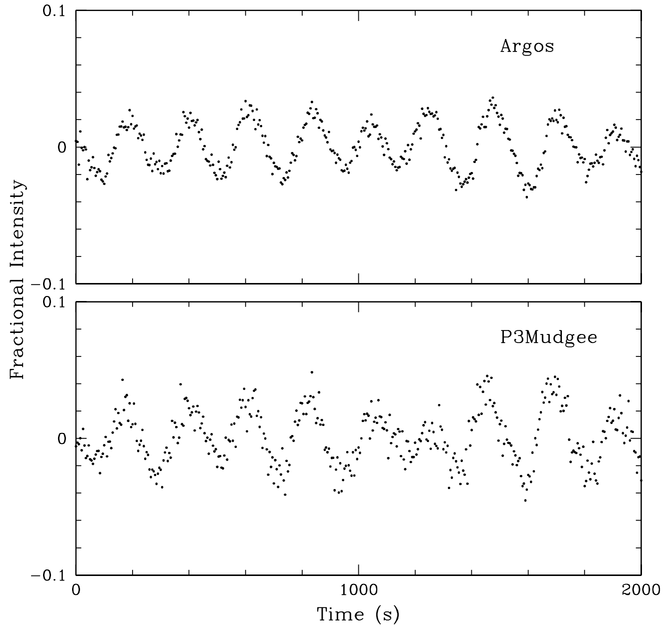

FIG. 3.—*Top*: Light curve of G117-B15A taken with Argos at an exposure time of 5 s under average seeing conditions. *Bottom*: Light curve of the same star observed with the three-channel PMT photometer, P3Mudgee, at the same telescope and with the same exposure time under excellent seeing conditions.

therefore slow. At the shortest exposure times, the acquisition process can easily keep up, but the display routines cannot. Arranged as separate threads of execution, they run in parallel, so acquisition gets its needed amount of CPU time even if the display falls behind. Should this happen, the user still sees the printed luminosity values appear right away, but the plotted values appear in clumps rather than one at a time, i.e., whenever the plot widget gets updated. No data points are lost. Display of some of the incoming images may be skipped, but the DS9 window always shows the most recent one when it is updated. The q11 program can run under any window manager, but users notice how much more slowly the display windows are updated using Gnome or KDE, compared to WindowMaker, which is much smaller and faster.

4. COMPARING: CCD VERSUS PMT

4.1. *Good Things*

4.1.1. *Digital Image Preservation*

The most notable improvement offered by the Argos instrument is the ability to record individual images for each integration for later inspection and analysis. This is very much like the historical transition that took place when photographic observations replaced visual ones. Extracting the measured brightness of the target, comparison star, and sky is done directly with the three-channel PMT photometer, the equivalent of 3 large pixels, through fixed apertures that isolate them from the rest of the stars in the field. What they see is what you get, with no going back.

The digital images arriving from the CCD are recorded as individual disk files and can be replayed without loss in the same time sequence as they were taken, with their time intervals the same as the original exposure times for simulation, or faster for analysis and reduction. Different extraction techniques can be applied for comparison, and the light curves with the highest signal-to-noise ratios can be chosen for further analysis. When a data point appears that does not fit well with those surrounding it, the corresponding image can be examined and often reveals the cause. ("Oh. I guess that's where I dropped my flashlight.") We find far more satellite tracks through the images than we ever expected.

4.1.2. *Greater sensitivity*

The higher quantum efficiency of the CCD coupled to the wider optical bandwidth gives higher photon counts from the same target stars, compared to the PMT photometers. Argos is designed such that light from the primary mirror forms an image of the star field directly on the CCD, so it has fewer optical surfaces than the PMT photometers. Figure 3 shows the light curves of the same white dwarf pulsator G117-B15A taken with the two different instruments 5 weeks apart on the 82 inch telescope with 5 s exposures. Observing conditions were excellent for the PMT run but only average for the CCD; even so, the CCD data are less noisy.

The ability to measure fainter stars is illustrated in Figure 4. The target star, a new DA variable WD 0815+4437 (Mukadam et al. 2003a), has a B magnitude comparable to 19.3, and clearly shows the pulsations; the Fourier transform (FT) shows two significant peaks (unresolved in this short run) and hints at more. Smoothing the light curve with a running average of three data points (to suppress noise at higher frequencies) shows the variations more clearly.

4.1.3. *Marginal Photometric Conditions*

The PMT photometers separate the target and comparison star fields optically, and do this before the image of the field is formed at the focal plane. This works but introduces a vignetting property that users must be aware of and avoid. This means the comparison star must always be at least 3′ away. Changing transparency from thin clouds does not always act on the target and comparison stars at the same time, so removing the effects of clouds by dividing the target light curve with that of the comparison star does not always work very well.

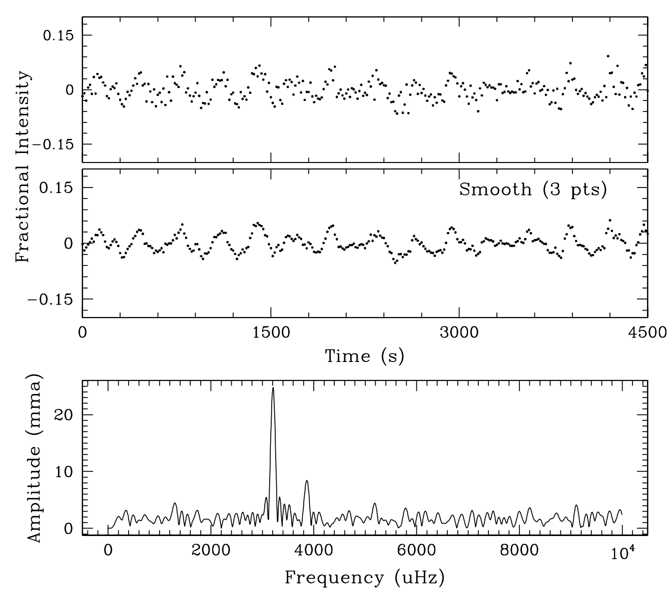

FIG. 4.—*Top*: Light curve of WD 0815+4437, a faint ($B \approx 19.3$) new DA variable star (Mukadam et al. 2003a). We acquired the data on 2002 February 1, with an exposure time of 15 s including the BG40 filter. *Middle*: Same light curve after a three-point smoothing. *Bottom*: FT of the light curve, where the scale of the y-axis is in units of millimodulation amplitude (1 mma = $0.1\% \Delta I/I$).

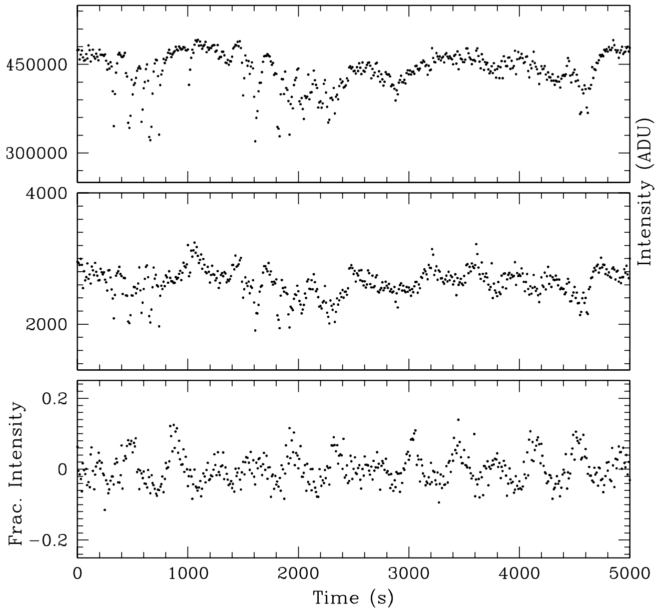

FIG. 5.—Cloudy weather conditions. *Top*: Summed light curve of two comparison stars in the field. *Middle*: Raw light curve of the (much fainter) target star, a new DA variable. *Bottom*: Reduced data on the star after it has been divided by the comparison light curve.

This technique of cloud removal works far better on the CCD images. In Figure 5 we show our data on the new DA variable WD 0949−0000 (Mukadam et al. 2003a, 2003b) taken through light clouds with 10 s exposures in 2003 April 2. The top panel shows the sum of two comparison stars, collectively brighter than the target by a factor of 150. The center panel shows the faint ($B \approx 18.8$) target star. The cloud-induced variations disappear when the target is divided by the comparison light curve.

4.1.4. *The Human Factor*

The PMT photometers require considerable skill and experience on the observer's part to get everything set up properly and to continually tend the observations during a run. Failing to align the two stars properly in their apertures, or to guide often enough but not too often takes some time to learn. Inexperienced observers can (and do) take data of poor quality until they have made most of the common mistakes and learned from them. The Argos photometer demands far less skill and experience; once the CCD instrument is set up properly, far less is demanded of an observer, and inexperienced observers can (and do) take good quality data on their first observing run with the instrument. Target star fields are much easier to locate and verify, and, unlike using the PMT photometers, faint targets are as easy to find as bright ones.

4.2. *Bad Things*

4.2.1. *Read Noise*

A photomultiplier detector amplifies individual photoelectrons until they can be readily detected as individual events; a CCD does not. In a CCD, each pixel-well collects unamplified photoelectrons until they are read out, amplified en masse, and sent to an analog to digital converter (A/D) over a 10 foot cable. The cable and its connectors are by far the weakest link in the instrument and can cause real grief if not properly maintained. In the Argos instrument the default amplifier gain is set so that two electrons yield one ADU, and it has a noise equivalent of 8 electrons. This means that 64 or more electrons must be collected before their stochastic noise and the amplifier noise are equal. We try to keep pixel counts below 40,000 to avoid the onset of nonlinear behavior; the PMT instruments have a somewhat wider dynamic range. These are not major problems but can hardly be considered assets.

4.2.2. *Short Exposure Times*

PMT detectors can count individual photon events, and these events can be placed into accurate time bins as short as desired. The CCD detector must accumulate many photon events before they can be measured, so the minimum time bins must be much longer. The PMT instruments work well in measuring lunar occultations, where 1–2 ms time bins are required, and for optical pulsar measurements, where time bins as short as 1 μs have been used (Sanwal, Robinson, & Stiening 1998). The smallest time bin available in the Argos instrument, 1 s, is short enough to measure pulsating white dwarf stars but not short enough for these other measurements.

4.2.3. *Proprietary Secrets*

Roper Scientific, the corporate owners of Princeton Instruments, the makers of the CCD camera in Argos, have a policy of not revealing technical details of either their hardware or their control software; users are given access to the camera's operation by a series of calls to a software library that they provide. The user must therefore treat the camera and its controlling software as a black box and can infer or measure details of its operation (e.g., timings) only by doing experiments. This makes troubleshooting and time-critical program development both slow and difficult. The Argos instrument and its software are unlikely to evolve more effective operation so long as this policy is in place.

5. SUMMARY: RESULTS TO DATE

The Argos instrument was placed into regular operation (with a less capable version of the software) in 2002 February. As of this writing (2003 April) it has achieved the following results:

1. Its increased sensitivity has allowed the identification of 35 variable white dwarf stars previously unknown (Mukadam et al. 2003a), doubling the total number of these objects available for study.
2. The improved data quality and the resulting improvement in measuring the time of arrival of pulses from the DAV white dwarf G117-B15A, has yielded the first measured $\dot{P}$ (S. O. Kepler 2003, private communication) for the rate of change of the 215 s principal period (previous measurements had only set limits). This is the first such measurement for a star this cool (and this old) and opens the way to calibrating by measurement, rather than by theory, the ages of the oldest stars in our Galaxy (e.g., Winget et al. 1987; Hansen et al. 2002).
3. The control software and a second CCD camera were used successfully at Siding Spring Observatory on the 1 m telescope, in support of a Whole Earth Telescope run in 2002 May. The f/8 Cassegrain focal position gave about the same plate scale and the same size images as the prime focus position on the 82 inch (f/3.9) telescope.

The success of the instrument, despite its limitations, can best be judged by its use: the PMT photometers have not been used on the 82 inch telescope since Argos was placed into regular operation.

We thank Frank Bash, Director of McDonald Observatory, for providing the funds to purchase the CCD camera, and the Texas Advanced Research Program for operating funds under grant ARP-0543. We thank Gary Hansen for designing the timer card, Gordon Wesley and David Boyd for the mount design, and Phillip McQueen for advice on the baffle design. We also thank Antonio Kanaan for the use of his CCD data reduction routines and Darragh O'Donoghue for showing us that off-axis coma would not be a problem. We thank Denis Sullivan for his help in the commissioning run for Argos.

APPENDIX

A1. THE SOFTWARE STRUCTURE

The ideal environment for a real-time program is to have a CPU dedicated to its task and to have complete control over all of its operations. In an earlier era (before operating systems became common) this was practical, but it meant the program had to include many of the functions we now expect an operating system to perform. Simple operating systems that supported only one user and one task (e.g., MS-DOS) could be used despite some loss of control. Multiuser, multitasking operating systems represent a more significant challenge to a real-time program but also offer system capabilities that can help do the job, within limits. The Quilt 11 program is designed to function in this less predictable environment and seems to be successful at it.

A single CPU can keep only one process at a time in execution, but it can be switched rapidly between several so that they appear to be running in parallel. The Linux kernel does this to support and schedule many processes, most of which are dormant until needed for the function they perform. The q11 program consists of five such processes, three of them dormant most of the time. Figure 6 shows this software architecture as a flow chart.

A2. STAR IMAGE IDENTIFICATION

Before online extraction can begin, the user must first tell the program which star image in a field is the target star and which other images it should use as comparison stars. When the "mark" command is detected, the current image is copied to a separate image buffer, and the display process is directed there. New images will still arrive but will not be shown until marking has been completed.

The process by which stars on a new image are identified with those initially marked is based on a simple assumption: that a star image may move around on the CCD chip, but its location with respect to other stars in the same field will not change (much). The identification algorithm therefore needs to remember a pattern of star images found in one image, so it can identify elements of the same pattern in a new one. Pattern recognition is a famously difficult programming problem, but we are fortunate here on two counts: first, our patterns are very simple and can be reduced to a small list of x, y locations of points in a plane; and second, we are really interested in pattern re-recognition, which is a far easier problem to solve.

To make this work, the program must first isolate the individual stars in an image (described in § A3) and make a position list (plist) for each one, recording its x, y location along with other useful values. The position is taken to be the pixel with the highest count. From this position list, the marking process can first identify which image the user indicates with a mouse click and can then make a reference list (rlist) of distance and position angle (dpa) to other stars in the field for later comparison.

The identification algorithm finds a known star in a new image by first deriving a set of dpas for it (from the plist of image locations) and then comparing them with the rlists saved from the marked stars. In effect, it is asking each star in a new field: Are

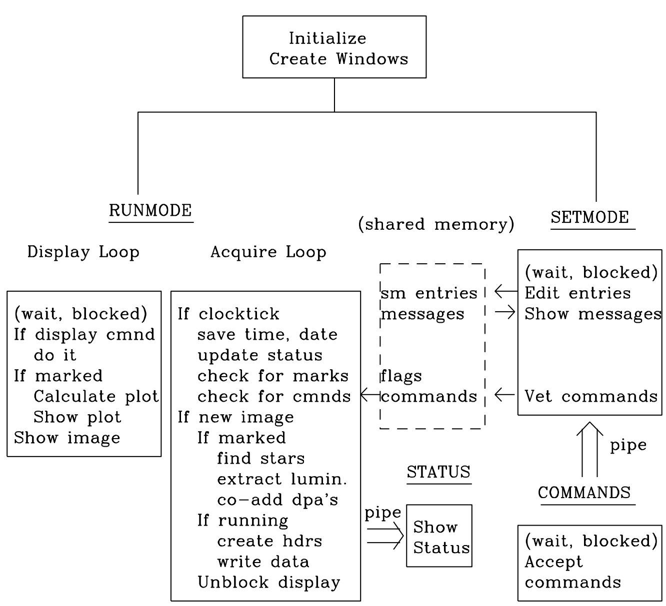

FIG. 6.—Quilt 11 architecture. Five processes (*solid boxes*) can execute in parallel, but four of them are blocked until needed. The Acquire loop polls the clock and the camera continuously, watching for the 1 Hz clock ticks or the arrival of a new image.

you on my list? Most often the answer is no; the dpas to other field stars do not fit one of the saved patterns, except by accident. An rlist can have up to 12 entries, and an accidental match to more than one or two of them just does not happen, even in a crowded field. Agreement with all of the rlist entries is common when a true match is found, but not required: the majority wins. If a known star drifts off the chip or is obscured by a cloud, the other stars are still identified properly. A running average of the 20 most recent reference patterns is kept to avoid the systematic effects of changing atmospheric refraction.

Even though the identification algorithm must examine every star in a new image, it does not spend much time doing so. The richest field we have encountered has about 230 stars detected in it, but the identification process needed less than 100 ms to do its job. It uses correspondingly less time if the field is less crowded.

A3. ONLINE LUMINOSITY EXTRACTION

The arrival of a new image triggers a flurry of activity that results in a new data point in the light curves of any marked stars. First, however, the star images must be isolated and located on the chip before they can be identified. The process of extracting the luminosity is melded in with the isolating procedure.

By analogy, we can think of the star images as luminosity mountains rising as peaks above the plane of the sky. If we flood the plane, and consider only the peaks that rise above flood level, then they are nicely isolated and can be treated individually.

We first approximate the sky level by finding the mean of all of the pixel values in the image (avl) and then setting a cut level (i.e., the flood level) enough above that mean to avoid finding false peaks due to sky noise:

$$\mathrm{cut} = \mathrm{avl} + 2.5\sqrt{\mathrm{avl}}. \qquad \text{(A1)}$$

We can now scan the image 1 pixel at a time and determine for each one if it is above the cut level, and therefore a part of some star image, or below it, and part of the sky level. In the process, we look for successive pixels above the cut level and contiguous with those on previous scan lines to define a luminosity "clump." We consider a clump as complete (and therefore isolated) when a scan in x finds no more pixels to include. The plist entry for this peak then contains the location of the maximum (in x and y), its total luminosity above the cut level, and the number of pixels summed.

This procedure works well for peaks that are not too broad but can sometimes yield small false peaks near a large one when the seeing is bad or the images are a bit out of focus. Akin to foothills in our analogy, these "skirt peaks" may confuse the identification procedure. To avoid this, when the scan is complete the plist is examined by a routine that identifies small peaks too close to big ones and merges them in. The resulting plist is then sorted by luminosity so the brightest ones appear first; identification can now proceed.

Once a peak has been identified with a user-marked star, its total luminosity (plus a small correction for the fraction of its luminosity below the cut level) becomes the next data point on its light curve, after a global sky value has been subtracted. This global sky value, found during the isolation scan, is the mean level of all the pixels below the cut level and is really composed of sky photons, dark count, and a bias value set by the electronics to ensure that only positive quantities are presented to the A/D converter.

The online extraction process was devised to provide the user with light curves in real time, the same as with the PMT instrument, to allow the data quality to be assessed. It was not intended to be a final data reduction procedure. However, cut-level extraction proves to work far better than originally expected and may evolve into a procedure that can rival the virtual aperture extraction technique (O'Donoghue et al. 2000); alternatively, that technique could be incorporated into the q11 program as a user-selected alternative.

REFERENCES

Hansen, B. M. S., et al. 2002, ApJ, 574, L155

Kanaan, A., O'Donoghue, D., Kleinman, S. J., Krzesinski, J., Koester, D., & Dreizler, S. 2000, Baltic Astron., 9, 387

Kepler, S. O., Mukadam, A., Winget, D. E., Nather, R. E., Metcalfe, T. S., Reed, M. D., Kawaler, S. D., & Bradley, P. A. 2000, ApJ, 534, L185

Kleinman, S. J., Nather, R. E., & Phillips, T. 1996, PASP, 108, 356

Mukadam, A. S., Winget, D. E., & Kepler, S. O. 2001, in ASP Conf. Ser. 226, 12th European Workshop on White Dwarfs, ed. J. L. Provencal, et al. (San Francisco: ASP), 337

Mukadam, A. et al. 2003a, ApJ, 594, 961

———. 2003b, in White Dwarfs, ed. D. de Martino et al. (ASI Ser. B, 105; Dordrecht: Kluwer), 227

Mullally, F., Mukadam, A., Winget, D. E., Nather, R. E., & Kepler, S. O. 2003, in White Dwarfs, ed. D. de Martino et al. (ASI Ser. B, 105; Dordrecht: Kluwer), 337

Nather, R. E., & Warner, B. 1971, MNRAS, 152, 209

Nitta, A., Kanaan, A., Kepler, S. O., Koester, D., Montgomery, M. H., & Winget, D. E. 2000, Baltic Astron., 9, 97

Nitta, A., Kepler, S. O., Winget, D. E., Koester, D., Krzesinski, J., Pajdosz, G., Jiang, X., & Zola, S. 1998, Baltic Astron., 7, 203

O'Donoghue, D., Kanaan, A., Kleinman, S. J., Krzesinski, J., & Pritchet, C. 2000, Baltic Astron., 9, 375

Robinson, E. L., et al. 1995, ApJ, 438, 908

Sanwal, D., Robinson, E. L., & Stiening, R. F. 1998, BAAS, 30, 1420

Winget, D. E., Hansen, C. J., Liebert, J., van Horn, H. M., Fontaine, G., Nather, R. E., Kepler, S. O., & Lamb, D. Q. 1987, ApJ, 315, L77

Winget, D. E., et al. 2003, in ASP Conf. Ser. 294, Scientific Frontiers in Research on Extrasolar Planets, ed. D. Deming & S. Seager (San Francisco: ASP), 59