Publications of the Astronomical Society of the Pacific
105: 1164–1174, 1993 October

The Berkeley Automatic Imaging Telescope

MICHAEL W. RICHMOND,[1] RICHARD R. TREFFERS, AND ALEXEI V. FILIPPENKO[2]

Department of Astronomy, University of California, Berkeley, California 94720
Electronic mail: richmond@astro.princeton.edu, treffers@bkyast.berkeley.edu, alex@bkyast.berkeley.edu

Received 1993 February 22; accepted 1993 May 30

ABSTRACT. We describe the design and control system of the Berkeley Automatic Imaging Telescope (BAIT), an observatory that performs all its tasks with minimal human intervention. We concentrate on the system's software, especially that which communicates with astronomers at remote sites and that which schedules targets during the night. BAIT has been in nearly continuous operation since 1992 January; we present an analysis of its performance during its first eight months.

1. INTRODUCTION

The nature of astronomical observing has undergone several great changes: the detector of choice has progressed from the human eye through photographic plates to electronic devices, telescopes have evolved from refractors to reflectors to multiple-segment or spin-cast reflectors, and the subject material has expanded from visible optical photons to include the entire electromagnetic spectrum. Inextricably intertwined with these has been the role of the observer. A minority of astronomers now endure long nights outside in the cold air, or hours of guiding while folded into the tiny space of a prime-focus cage. Instead, those who must still work through the night now enjoy the luxury of a heated and well-lit room. Many astronomers simply arrange to have someone else perform their observations for them, especially those who use facilities in space. Since instruments are often situated at remote sites, it is difficult and expensive for researchers to transport themselves to the observatory. With the advent of worldwide electronic communication networks, and the ubiquity of digital detectors, it is much simpler and cheaper to bring the data to the observer. However, most observatories still require *someone* on-site to follow the astronomer's instructions, operate the instruments, and distribute the resulting information.

Following in the footsteps of several pioneers, we have tried to take the process one step further by eliminating the human operator at the observatory site. Unlike *remote observing*, in which the astronomer controls the telescope and instruments in real time (i.e., during the night) from a distant location, *robotic observing* allows the astronomer to submit instructions for observations to be made at some time in the future, and then simply wait for the results of those observations to be sent back to her. She may work on other projects, sleep, eat, or go on vacation while the automatic system performs all the necessary tasks. In theory, at least, no human presence is required on-site except for maintenance.

We have devised a system in which all communication between the astronomer and the observatory takes place over the Internet computer network. Users send instructions in the form of simple ASCII text files via electronic mail (e-mail), and receive the results of their observations either by e-mail or via simple file-transfer programs, such as FTP. Astronomers on other networks may submit their instructions via gatewayed mail to Internet, but retrieval of their results is not yet so straightforward.

The Berkeley Automatic Imaging Telescope (BAIT) project was begun by one of us (A.V.F.) in 1989 with NSF funding. In its present form, BAIT consists of a 0.5-m reflecting telescope equipped with a CCD camera and an autoguider. Articles in several conference proceedings contain the scientific motivation (Filippenko and Richmond 1991; Filippenko 1992b) and major progress reports (Richmond and Filippenko 1991; Richmond et al. 1992; Treffers et al. 1992), while a brief update is given by Richmond (1991). We emphasize that BAIT is *not* designed exclusively for any specific project or dedicated search; rather, it is an instrument that may be—and is—used by a wide variety of groups, for projects ranging from optical studies of gravitational lensing to supernova followup photometry to undergraduate instruction.

2. OTHER AUTOMATIC TELESCOPE SYSTEMS

Several groups began to automate parts of the observing process in the 1960's, as computers became available to academic institutions. We have not made an exhaustive literature search, but find that groups at the University of Wisconsin (McNall et al. 1968; Code 1992) and at Kitt Peak (Maran 1967) succeeded in producing complete observatories that required minimal human oversight. Sterling Colgate and collaborators at New Mexico Tech spent several decades developing a very ambitious system that culminated in the Digitized Astronomy Supernova Search (Colgate et al. 1975; Colgate 1982; Pearce 1986). Considering the primitive state of computer technology in the 1960's and early 1970's, the accomplishments of such early projects are especially impressive. However, we were unfamiliar with these facilities when we began work on BAIT, and have come to realize the many similarities between their systems and our own only recently.

BAIT most closely resembles two more recent automatic telescope designs, and in fact we learned many valuable lessons by examining them in detail and talking with

[1]Present address: Dept. of Astrophysical Sciences, Princeton University, Princeton, NJ 08544.
[2]Presidential Young Investigator.

© 1993. Astronomical Society of the Pacific

© Astronomical Society of the Pacific • Provided by the NASA Astrophysics Data System

their creators. The Automatic Photoelectric Telescopes (APTs) designed and built by Genet and Boyd (Genet 1986a; Genet 1986b; Genet and Hayes 1989) have had great success in obtaining high-precision photometry of relatively bright ($m_V \leqslant 12$ mag) stars. A number of institutions have installed their own APTs or subscribe to the APT Service on Mt. Hopkins. However, since an APT uses a photoelectric photometer, rather than a CCD chip, it can measure only one object at a time, and its integration time must be sufficiently short that guiding is not required. In addition, the reduction and analysis procedures for photometer-derived data are completely different from those for images.

The Supernova Search telescope of the Lawrence Berkeley Laboratory (LBL), on the other hand, used a CCD to take images of galaxies in a completely automatic manner (Smith et al. 1988; Perlmutter et al. 1992). Since the LBL group based their effort at Leuschner Observatory (using a 0.75-m telescope which we have recently converted to BAIT-like operation), we have been able to copy many parts of their operation. There are three major differences between BAIT and the LBL instrument. First, without an autoguider, the LBL telescope was limited to exposures of 3 or 4 min at most, and lacked the extra pointing precision our guider provides. Second, the scheduling system it used was much less flexible than ours, since their project involved a repeated survey of nearby galaxies. Finally, we have attempted to create a means of communication between telescope and astronomer that was not needed for the LBL supernova search.

Several other astronomers have also recently developed automatic imaging telescopes for astronomical research (e.g., Honeycutt and Turner 1992; Schmidt 1992; see other contributions and references in Filippenko 1992a). As with the LBL system, however, these are not equipped with autoguiders; hence, tracking errors limit the integration times to a few minutes. Although faint objects can be measured by coadding many individual images, we felt that superior results could be obtained with a few relatively long exposures. Moreover, we wanted to construct a general-purpose instrument that could satisfy the requests of many different users, who might need different modes of operation, and who might be located thousands of miles from the observatory.

3. TELESCOPE AND DETECTOR HARDWARE

BAIT is currently built around a 0.5-m classical Cassegrain telescope, made by Tinsley Laboratories in Berkeley in 1954; our original plans were for a 0.75-m telescope with a modern, compact design (AutoScope; see Genet and Hayes 1989), but delays with the Ritchey–Chrétien optics led us to adopt the smaller, existing telescope as a testing platform. The telescope is fairly slow, with focal ratio $f/11$, which yields a scale of $33''$ mm^{-1} at the focal plane. There is a relatively small field of view which is free from coma; although our main detector lies entirely within the nearly stigmatic region, the off-axis autoguider does not, and consequently suffers from significantly asymmetric stellar images. Once the new optics are completed, they will yield a much larger coma-free area, which should help the guider to acquire faint stars.

The site of BAIT, Leuschner Observatory, is not a very good one: since it is $\sim$15 km from Berkeley, the sky is filled with light from the cities and suburbs of the San Francisco Bay area. The sky brightness is about 18 mag arcsec^{-2} in the V band. Leuschner Observatory is only about 300 m above sea level; during the summer months, fog often rolls in from the Pacific just after sunset, burning off only after sunrise. Even when there is no fog, the sky transparency can change significantly during a single night, making photometric calibration quite tricky. Finally, the seeing is mediocre at best [about $2''.5$ full width at half-maximum (FWHM)] and poor on average ($\sim 4''$ FWHM). The best feature of Leuschner is its proximity to the Berkeley campus, which allows us to test and maintain equipment with minimal expense of time or money for travel. We hope to move BAIT to Lick Observatory, a much better site, but the relocation costs are substantial and no funding source has yet been found.

A Thomson TH 7895 CCD chip is the main detector of BAIT. Its array of 516×516 pixels, each 19 μm on a side, gives a field of view of just over $320''$. The pixels, which subtend about $0''.63$, oversample stellar images at Leuschner, but should be well matched to the seeing at Lick (or any other reasonably good site). The camera head and electronics, provided by Photometrics Ltd., are cooled; not by liquid nitrogen, which requires daily refilling of the dewar, but by a simple thermoelectric unit with a closed-cycle liquid coolant that is pumped through the CCD head. Although this method cannot reach the same low temperatures as liquid-nitrogen-cooled dewars, getting down only to -45 °C rather than -90 °C, the CCD electronics are able to keep the dark current down to reasonable levels by running in Multi-Pin Phase (MPP) mode. As a result, the chip accumulates only about 0.4 electrons per second per pixel, well below the contribution of the sky through a broadband filter. The readout noise of about 14 electrons is similarly negligible, except in very short exposures. Our main concern with the chip is its lack of response in the blue spectral region. Despite a special ultraviolet-sensitive coating, the quantum efficiency of the chip falls off quickly at wavelengths shorter than 5000 Å.

The poor blue response of our detector has driven us to use a slightly different set of filters for most objects; rather than the traditional photographic trio of U, B, and V, we prefer to use the redder filters V, R, and I. Only rarely do we use B, and the U filter requires very long integration times on any object fainter than about tenth magnitude. As a result, our photometry of supernovae, in particular, is not easily compared to results in the literature, most of which are in the B and V passbands. As silicon detectors become the standard, however, our problems may be shared by other observatories, and gradually V, R, and I may become the standard bands for supernova work (see, for example, Ford et al. 1993). In addition to the broadband $UBVRI$ filters (based on the bandpasses of Bessell 1990), the 20-position filter wheel of BAIT has a number of other filters

© Astronomical Society of the Pacific • Provided by the NASA Astrophysics Data System

which may be used at any time: Strömgren u, v, b, and y, narrowband ($\Delta\lambda \approx 100$ Å) Hα and Hβ filters, and a pair of narrowband filters just off the Hα and Hβ lines for continuum subtraction. Several filter positions remain empty, so there is room to add additional filters for special projects.

A very important piece of hardware for an automatic telescope is the weather station. BAIT's station constantly monitors the temperature, relative humidity, and wind speed; it also includes a raindrop sensor and a "cloud sensor," which is simply an infrared detector with a 12–13 μm filter pointed at the sky. If the relative humidity or wind speed exceed critical values, or if a single drop of rain hits the sensor, direct circuits from the station to the dome shutter motor cause the shutter to close (bypassing all the control software, which might possibly be hung up or unable to respond immediately). In addition, there is a solar cell which prevents the dome from opening when the Sun is at an altitude above about 10°, and a "dead-man timer" which closes the dome shutter if no telescope activity occurs for 20 min. All of the weather information is logged by a special process throughout the night, from which one can make rough estimates of the night quality when analyzing images at a later date.

4. COMPUTER CONTROL SYSTEM

BAIT is controlled completely by computers, and usually runs unattended (although one can interrupt the system—either on site or over a network—and manually point the telescope, make an exposure, etc.). One workstation acts as the "master" computer, controlling a number of other computers which act as its "slaves." The master/slave arrangement is almost identical to that of the LBL system (Smith et al. 1988), and similar to the relationship between DIRECT and other computer processes used by the Digitized Astronomy Supernova Search (Pearce 1987). Most of the machinery (motors, gears, cameras, etc.) is controlled by PCs. They are cheap to replace or repair (an important consideration for machines sitting in the dome itself), are always in stock at local stores, and have a very large selection of control cards which allow us to buy much of the system off-the-shelf, rather than having to design and build it ourselves. Each is connected to the master computer by a local Ethernet-based network running TCP/IP. The slaves sit passively until they receive a command from the master computer; after completing the requested action (or trying and failing), they send back a short status message. All the communication is via ASCII strings, which are written into a log file for record keeping and trouble shooting. The programs run by the slave computers are therefore relatively simple and robust.

The master computer, on the other hand, must handle communications with all the slaves, and with observers as well. We are using a Sun SPARC station 2 running UNIX. Its software is more complicated and, therefore, prone to failure. Most of the problems that we encountered while first testing BAIT were caused by bugs in the communications software on the master computer. Making changes to the code itself is a perilous task; in order to allow the equipment to evolve without forcing rewrites of the software, we have placed much of the information about each device attached to BAIT into a single configuration file. One can edit the ASCII file (a small section is shown in Fig. 1) with any text editor. Changing the site parameters in this configuration file is all that is required if we move BAIT to another location. Should our main CCD chip need repair, we can simply replace the camera head with a spare one (which houses a Thomson 7883 chip), then comment out the "Th7895" section shown in Fig. 1 and remove the comment markers from the "Th7883" section. No recompilation is necessary, and the system can be running when changes are made.

```
#automatic telescope configuration file for the B.A.I.T
#observatory parameters
OBSERVAT='Leuschner'
OBSLAT  =  37.91833                    / leuschner latitude in decimal degrees
OBSLONGH=  8.1437972                   / leuschner longitude in hours
OBSELEV =  304.                        / leuschner elevation (m)
#
#ccd camera parameters
DETECTOR='Thomson TH7895'
CCDPIXSZ= '19.0 19.0'                  / pixel size along columns and rows (microns)
CCDSEC  = '[1:512,1:512]'              / number of columns and rows
CCDROT  = 0.25                         / CCD rotation angle (degrees)
CCDHAND = -1.0                         / +1 if (x,y)::(ra,dec) same handedness
CCDDARKC= .4                           / CCD dark current, electrons per sec per pixel
CCDWELL = 250000                       / full well capacity, in electrons
CCDREADN= 14                           / CCD read noise, in electrons
CCDGAIN = 25.4                         / adus per electron
#
#DETECTOR= 'Thomson Th7883'
#CCDPIX  = '23.0 23.0'                 / pixel size along columns and rows (microns)
#CCDSEC  = '[1:384,1:576]'             / number of columns and rows
#CCDREADN= 40.0                        / CCD read noise, in electrons
#CCDROT  = 180.0
#CCDHAND = -1.0
#CCDDARKC= 1.0                         / CCD dark current for Thomson
#CCDWELL = 250000                      / full well capacity, in electrons
#CCDGAIN = 13
#
```

FIG. 1—Sample section of the configuration file for BAIT. The lines beginning with "#" are "comments" not used by the program.

Note the similarity between the configuration file's format and that of a FITS file header. We hope that the near-universal usage of FITS will give other astronomers a head start if they attempt to understand the BAIT system.

With only a small amount of additional software, we have enabled the master computer to control simultaneously both BAIT and another telescope at Leuschner Observatory [the 0.75 m used until 1992 by the LBL supernova search group (Perlmutter et al. 1992)]. The same software is used for each telescope, but the processes running each telescope read different configuration files. The amount of effort needed to operate and maintain two instruments is therefore not very much larger than that required for one. We would eventually like to place some sort of spectrograph on the larger telescope, allowing simultaneous photometric and spectroscopic study of objects.

5. AUTOGUIDER

A fully automatic telescope must not only guide itself, but also acquire the stars which it will use to guide. The latter task is the harder one, especially when the guide stars are used to confirm the pointing of the telescope as well. Some tradeoffs to consider when choosing a guider are discussed by Richmond (1990). We have constructed an autoguider which searches for the proper guide star for a given object, acquires it, and makes small corrections to the tracking for arbitrary periods of time.

© Astronomical Society of the Pacific • Provided by the NASA Astrophysics Data System

TABLE 1
Statistics of the Galactic Stellar Distribution (from Allen 1973)

Visual mag $\leq m$	Stars in 4′ × 4′ Galactic eq.	Stars in 10′ × 10′ Galactic eq.	Stars in 4′ × 4′ Galactic pole	Stars in 10′ × 10′ Galactic pole
8	0.01	0.05	0.002	0.01
9	0.03	0.17	0.006	0.04
10	0.08	0.49	0.016	0.10
11	0.24	1.5	0.04	0.25
12	0.67	4.2	0.09	0.59
13	1.8	11	0.22	1.4
14	4.7	29	0.45	2.8
15	11.7	73	0.83	5.2

The guider's detector is a relatively inexpensive (~$1000), thermoelectrically cooled camera (the SpectraSource Lynxx-Plus CCD camera) based on the Texas Instruments TC211 chip. One of the reasons for its low cost is its small size: the CCD is only 165×192 pixels, each of size 13.75×16 μm^2. At the plate scale of our telescope, the camera has a field of view of only about two square arcminutes. We cannot simply add strong reducing optics to the camera, since a relatively small pixel size is required for good guiding. As a compromise, we have inserted moderately strong (~3×) reducing optics which yield pixels about 1.″4 in size and a total field of view about 4′ on a side. At a better site, we might have to make the pixels somewhat smaller to reduce guiding errors, but it will be easy to change the magnification by modifying a spacer on the transfer lens.

Even with such large pixels, the guider's field of view is too small; it must be large enough to include at least one star of sufficient brightness to permit guiding, no matter where the telescope may be pointing. Our current optics and site limit us to guide stars of roughly twelfth magnitude or brighter. As Table 1 shows, there will be many areas near the Galactic poles where not a single (sufficiently bright) star falls into the guider. It is clear, however, that if the guider's field of view were larger, about 10′×10′, then the chances of finding a suitable guide star would be very good near the equator and roughly 50% even at the poles. We therefore attached our guide camera to an x–y translation stage that moves it in two dimensions in the image plane, so that the effective field of view is limited by the vignetting of the telescope and the motion of the x–y stage rather than by the size of the CCD chip. The stage has a 50 mm range of motion in each axis, increasing the field of view to a square about 14′ on a side. Since the central portion of the field goes to the main camera, however, the effective field of the guider is only slightly larger than 100 square arcminutes, including the effects of vignetting near the field edges.

We use the *Hubble Space Telescope* Guide Star Catalog (GSC) (Lasker et al. 1990; Russell et al. 1990; Jenkner et al. 1990) to locate all possible guide stars around a target object and rank them in order of apparent brightness. Since our guider is based on a red-sensitive CCD, it sometimes measures a different "brightest star" for a field than the one indicated by the pseudo-V magnitudes listed in the GSC, but the effect is at most a few tenths of a magnitude and rarely causes a problem.

Using the GSC to pick guide stars, rather than trying to find them *ab initio* in the sky, provides several benefits. First, as each request for BAIT comes in over the network (see Sec. 6.2), the target position is checked for guide stars; if none should happen to fall within the guider's field of view, e-mail can be sent back to the observer immediately to inform her of the problem. Second, the relative scarcity of stars bright enough to be detected by the guider actually works to our advantage: after slewing to a new field, if BAIT does find a star at the predicted offset from the target position, it is fairly certain to be the right star—which in turn means that the telescope is pointed correctly at the target position. Conversely, if there is no guide star at the predicted position, BAIT knows that it must be slightly mispointed and can therefore attempt to reset its coordinate system. A more sophisticated approach to identifying stars in an acquisition guide image (using pattern matching of some kind) is usually unnecessary and, if only one star is present, would fail anyway. Finally, since the same guide star is used for all images of the same target, and that guide star is always centered on the same portion of the guider, all the images of a target are aligned with each other to within a few pixels.

We find that the action of the guider in "closing the loop" of the pointing system (providing constant feedback at each target acquisition) greatly simplifies the process of writing software to operate the telescope and to analyze the results. In this sense, the autoguider acts very much like the "starfinder" camera in the Kitt Peak Remotely Controlled Telescope (Maran 1967). In fact, we use the guide camera to refine the telescope's position even when exposures are so short that no actual guiding is necessary, by following steps 1 through 6 described below before every observation (but see footnote 4 in step 5). BAIT's autoguiding procedure works as follows:

(1) When a request comes in, the GSC is searched for a guide star with an appropriate offset from the main target's coordinates. The guide star's position and magnitude are written into the request and that position is used for all subsequent steps.

(2) When a request becomes active, BAIT calculates the offset of the guide star from the main target and moves the mechanical stage to that offset as the telescope slews across the sky.

(3) As soon as the telescope stops moving and the proper filter is in place, the guider takes a short (5 s) exposure and searches a small central portion of the chip (the "hot spot" region) for a star. In many cases, the telescope takes successive pictures of the same field without moving at all; this quick peek at the "hot spot" saves time if a guide star is already there. If a star is found, execution passes to Step 6.

(4) If no star falls within the "hot spot," a longer (20 s) exposure is taken and the entire guider image is searched for stars.

(5) If no stars are found, BAIT initiates a procedure to reset its coordinates (by pointing to the nearest isolated star from a list of bright ($m_V \leqslant 6$ mag) stars, taking a picture with the main camera, and looking for a bright point source), and then starts again at Step 3. If no stars are found again, the request may be marked with an error message and discarded, or it may be honored with an un-

© Astronomical Society of the Pacific • Provided by the NASA Astrophysics Data System

guided exposure.[3] In the latter case, an appropriate comment is put into the exposure's FITS header.

(6) If any stars are found, BAIT locates the brightest one and nudges the telescope to move the star into the center of the "hot spot" area.

(7) The guiding loop begins: the guide camera begins to take a series of short (1–20 s, depending on guide star brightness) exposures, reading out the "hot spot" region only. After subtracting the local sky, BAIT uses a simple centroid to determine the star's position. If it has moved from the center of the "hot spot," BAIT nudges the telescope to move it back in place.

(8) The guiding loop is continued for the length of the exposure being made through the main camera, and then ends. The dome position is checked occasionally during the loop and adjusted as needed.

Guiding is best on bright stars, for which the cycle from guide image to guide image takes only a few seconds. However, even the faintest stars the guider can currently use do not cause a large increase in the target's image FWHM. If we were to move to a site with better seeing, it might be necessary to change some of the guiding parameters (such as moving the telescope only a fraction of the offset from center, or decreasing the longest acceptable guide exposure time). Of course, better seeing, and improved optics, would improve the quality of the guide star images as well.

The current limit for guide stars, $m_V \leqslant 12$ mag, has not affected BAIT's operation very much. Only three or four of the hundred or so targets we have selected in the first eight months have failed to include at least one satisfactory guide star in the field nearby. For those targets, it is possible to take unguided images with short exposure times and coadd the results, or to make a single long-exposure image and hope for the best. (The telescope tracking varies substantially, but sometimes allows exposures of up to 5–6 min without significant trailing.) Of course, observers can check the GSC themselves to confirm that there is a good guide star candidate for their targets before creating a request.

6. COMMUNICATIONS BETWEEN BAIT AND THE ASTRONOMER

In order for astronomers to use a remote telescope, they must be able to send their instructions and receive the results of their observations. If a night assistant happens to be present at all times, a simple phone call will perhaps suffice. A completely unattended telescope, however, requires a more sophisticated (although not necessarily a more effective!) approach. BAIT is designed to be used over computer networks: astronomers send their requests by e-mail, receive a notice of their results by e-mail, and may retrieve their images via FTP.

[3]This depends on the value of the GUIDE keyword: requests with GUIDE ='guide' will be discarded, while requests with GUIDE='tryguide' or GUIDE='noguide' will be honored.

```
OBSERVAT= 'Leuschner'
INSTRUME= 'B.A.I.T.'
OBSERVER= 'Richmond'
MAILADDR= 'richmond@astro.princeton.edu'
SENDMAIL= T
OBJECT  = 'SN 1991T'
RA      = '12:31:36.9'
DEC     = '+2:56:28.0'
EPOCH   = 1950.0
EXPTIME = 600
FILTERS = 'R'
DAYSTART= '01/12/1991'
DAYEND  = '31/12/1991'
NUM-OBS = 10
END
```

FIG. 2—A simple request file. Ten observations of SN 1991T (in NGC 4527) are to be made during the month of 1991 December.

6.1 Request File Format

Requests for observations must be placed into a standard form before being e-mailed to the BAIT computer. Following a suggestion by Garrett Jernigan, we have chosen to use a format similar to that in the headers of FITS files (Wells et al. 1981; Greisen and Harten 1981; Grosbol et al. 1988; Harten et al. 1988) for "request files," those which observers submit to BAIT. Unlike true FITS, however, request files have no mandatory 80-character-per-line rule; instead, each line may contain only as many characters as are required, followed by a "newline" character to mark the end of the line. Thus, requests are simple ASCII text files which may be created or modified with any editor and sent easily by e-mail. Request files can have one or more *stanzas*, each delimited by its own END statement, which are executed sequentially and without interruption; each stanza specifies a single exposure. Thus, requests with multiple stanzas can ask for exposures of the same object through different filters (or through the same one), or a sequence of standard-object-standard images, or almost-simultaneous observations of several different targets. For convenience, information is carried over from one stanza to the next so that only keyword-value lines which differ from the previous stanzas must be listed.

There is a plethora of keywords for request files, but only a relative few that the average astronomer might wish to use. We will describe some of the more important ones briefly, and then explain how one can request several common modes of observations. A very simple request file is shown for reference in Fig. 2. The first two keywords, OBSERVAT and INSTRUME (note the eight-character limit to keywords, as in FITS), identify the site and telescope to which the request is being made; in the future, we expect to have several telescopes operating at Leuschner, and perhaps similar automatic systems at other institutions. The OBSERVER and MAILADDR keywords allow the observer to identify herself and given an e-mail address to which BAIT can send messages. The RA, DEC, and EPOCH keywords specify the position of the target, and OBJECT its name. The FILTERS value declares the filter through which the observation should be made, and EXPTIME its length (in seconds). The period over which observations should be made

© Astronomical Society of the Pacific • Provided by the NASA Astrophysics Data System

may be delimited by the dates—in form "day/month/year" following the international convention—given in the DAYSTART and DAYEND lines; if these are not supplied, the request is queued for observation as soon as it is received, and remains active until completed. Finally, the number of times the request ought to be executed can be specified in two ways: either by an explicit number of observations, NUM-OBS (default value one), or by an interval in days, INTERVAL (default value one).

The example above will yield several images of the target, one per night, taken at (roughly) equal intervals. Setting NUM-OBS=1 requests just a single image, of course. If, on the other hand, a number of images on the *same* night are desired, the user may either

(a) create a request file with a number of identical stanzas, in which case the observations will be made sequentially, or

(b) add a line NUMPERNI=x, in which case x copies of the request file will be made and scheduled at regular intervals throughout the night.

The latter option lends itself easily to abuse.

By default, requests are scheduled near the "best" time: when the target(s) is (are) as close to the meridian as possible (see Sec. 7.2). However, there are several ways in which the user can specify the time of observation more explicitly. Adding

UT-START='12:32:00'

to a request forces it to be scheduled for 12:32:00 UT exactly (or at least to within a minute or so). Alternatively, the starting time may be given as a Julian Date,

JDSTART=2448563.345,

although such a request must also include the correct (and identical) DAYSTART and DAYEND values, redundant as they are.

Forcing a request to be made at some given time greatly reduces the flexibility of BAIT's scheduling program. Moreover, if the request cannot be fit into a night at its exact time, due to bad weather or conflicting observations, it will not be executed at all. It is better, if possible, to give the system some leeway in scheduling the request, by giving a window in local sidereal time during which observations are permitted:

LSTSTART='10:23:00'

LSTEND='11:23:00'.

The larger the window, the higher the probability of getting an observation. Leaving the scheduling to the BAIT software, however, gives requests the greatest chance of being scheduled, and should be perfectly adequate for most purposes.

The request file allows astronomers to control not only the time of the observation, but also the procedures by which the data are processed. The default is for BAIT to remove the bias and dark current from an image, then divide, it by a median twilight sky flatfield frame; specifying

PROCEDUR='photo_proc'

yields this result. However, if the user desires to get the raw image itself, she may specify

PROCEDUR='photo'.

If the astronomer wants simply a list of the positions and brightnesses of all stars in the field (as determined by aperture photometry), she may add

PROCEDUR='photo_all',

in which case the resulting information (two ASCII text files, which are much smaller than a single image) is e-mailed directly back to her. For studies of variable stars in uncrowded fields, this procedure may be perfectly adequate. Since the software currently produces instrumental magnitudes, however, some additional information is needed to calibrate the data; the user may add exposures of standard stars to the request and get that information directly, or ask a member of the BAIT team for photometry of the standard fields that are scheduled every night.

We feel that the PROCEDUR keyword, which gives BAIT a powerful capability to adapt to the needs of its users, is a key innovation. By defining a new procedure (and convincing the telescope managers to install it), astronomers can change the way in which their data are processed or analyzed automatically. For example, we have written procedures to compare frames with previous ones of the same area and search for new objects, which we use to look for supernovae in images taken with the 0.75-m telescope at Leuschner (but not BAIT itself).

Observers with several projects running on BAIT simultaneously may wish to give some of them a higher priority than others. The PRIORITY keyword allows users to give their requests a relative priority; smaller numbers denote more important targets. A value of

PRIORITY=10

is used by default. Users who attempt to give all their requests a high priority will find themselves in some kind of trouble. The scheduling software (see Sec. 7.1) adjusts the priorities slightly (by increasing the priority of targets which have not been observed for some time, for example), but only by small amounts.

6.2 Network Communications

All communications between BAIT and the astronomer take place over the Internet network. The astronomer sends a request file to the user

bait@bait.berkeley.edu

via e-mail, either directly from another Internet node, or indirectly through a gateway from some other network. As soon as the e-mail reaches Berkeley, its format is checked for validity. If there is some problem, a message is sent back to the MAILADDR in the request, outlining the error and asking the user to fix the problem and try again. If the

© Astronomical Society of the Pacific • Provided by the NASA Astrophysics Data System

file has no errors, a message confirming its validity is e-mailed back to the user, along with a request ID.[4]

At the end of each night of observing, BAIT goes through the images that have been taken and sends a message to the user who requested each one. The message consists largely of the entire FITS header of the image (converted to normal ASCII form with a "newline" character at the end of each header line); all the information on the circumstances of the observation (object name, time, filter, length, etc.) is contained in the header. The header also includes the ID of the request which spawned the observation, such as

REQID='AugAef5',

and a second ID code which is unique to this particular image, e.g.,

DATID='SepD4fH'.

The DATID serves as the name of the FITS file which is stored on disk by the BAIT fileserver. All FITS files are given the ".fts" extension. Images files are stored in a compressed format, using the standard UNIX "compress" program, which appends ".Z" to file names. Thus, the complete file name of the example image would be SepD4fH.fts.Z. Instructions on retrieving the image via FTP are included in the message as well. Since FTP can only be used from one Internet node to another, users on other networks will have to borrow an Internet account in order to get their images.

On the other hand, users who specify

PROCEDUR='photo_all'

in their requests, asking for stellar coordinates and photometry rather than entire images, do not need direct Internet access. Since the reduced data files are both small and composed of normal ASCII characters, they are sent directly to the user by e-mail. The photometry files are stored on the fileserver as well, and may be transferred via FTP in the same manner as image files if necessary.

Data may be stored on the BAIT fileserver for only a relatively short period of time (roughly a week) before being replaced by more recent images. Users should therefore respond promptly when notified that their targets have been observed. Although every image taken by BAIT is stored in an archive on the site computer, locating and retrieving an image from the archive requires human intervention and is not recommended. Users are responsible for making their own copies of the data within a few days of observation; they may be less likely to request large amounts of telescope time if they have to keep up with the data flow themselves.

There are many details of the communication between BAIT and its users that remain unsettled. We would like to define facilities for listing, modifying, and deleting requests, to improve the manner in which errors are handled, and to allow users to predicate their requests on weather conditions. If possible, we plan to use commonly available networking facilities.

[4]The ID code is based on the date and time at which a request is received (REQID) or observed (DATID). The resulting names are awkward, but compact and unique. They serve as identifiers for all subsequent references to the request and its image(s).

7. SCHEDULING BAIT REQUESTS

One of the most important parts of an automatic telescope is the program that determines which objects are to be observed, and when. Designing such a program is difficult for many reasons, not the least of which is the lack of a unique "best" sequence in which to observe a given set of targets. Unless the targets are sparsely distributed and together require only a small fraction of the total length of the night, they cannot all be observed while crossing the meridian. The ideal scheduler must be able to handle all the requests for time on an oversubscribed instrument in a reasonable manner, and deal with the vagaries of weather as well. There have been a number of ideas promulgated in the literature [Johnston 1988; Genet and Hayes 1989 (Chapter 12); Barrett et al. 1991], but the BAIT software follows a novel approach.

We will describe its operation in two sections: *macroscheduling*, arranging observations over a period of many nights, and *microscheduling*, putting in order a set of targets over the course of a single night.

7.1 Macroscheduling

Observers can ask BAIT to check their objects once, or several times during some specified time period. As soon as it receives and verifies a request as valid, the BAIT software calculates an interval (in days, using fractions if necessary) by which observations of the request ought to be spaced. That interval is then written into the request file for future reference:

INTERVAL=3.1.

Whenever a request is honored, the data of the observation is also written into the request:

DAYLAST='12/06/1992'.

At noon of every day, a special process runs on the BAIT master computer; it goes through all the requests currently active, weeding out those which have expired and adding to the pool of candidate requests new ones which have arrived over the network and those whose starting date has arrived. It then calculates a probability for each request in the candidate pool, in the following manner:

–requests whose observation period is a single night (the upcoming night) are assigned PROBABIL=1.0;

–requests which have not been observed yet get PROBABIL=1.0;

–requests which have been observed, but require further images, are assigned a probability as follows:

let I=interval between observations and L=days since last observation

$$\text{if}(L \geqslant I) \quad \text{PROBABIL}=1.0$$

$$\text{if}(L < I) \quad \text{PROBABIL}=e^{-6.0*[(I-L)/I]}.$$

© Astronomical Society of the Pacific • Provided by the NASA Astrophysics Data System

We chose to use an exponential function, rather than a step function with a length of INTERVAL days, in order to compensate for missed observations. Bad weather or a heavy telescope load can delay a request for several days (longer delays may occur, of course, but these make any regular program of observation moot); observations made *at least* a fixed number of days since the last will therefore number fewer than desired, and the average spacing must be greater than that requested. By giving requests a small chance of occurring at smaller intervals, the total number of observations and the average period between them may be closer to that initially requested. The exponential function has a reasonably sharp cutoff, and a single parameter which may be tweaked to bring the average interval close to the desired one. Since both weather and telescope load vary over the course of a year, the exponential coefficient might be changed slightly from season to season; when the weather is good, for example, it could be increased. Different sites and different observing programs would require somewhat different coefficients, or perhaps even a different function for calculating the probability.

Once all candidate requests for the coming night have been assigned probabilities, the software generates a random number of each one and checks to see if it is less than or equal to the probability; if so, that request is added to the list of targets for the night. Otherwise, the request is removed from the pool of candidates for the current evening. When the telescope is very lightly loaded, this procedure seems unnecessary or even counterproductive, since it leaves the telescope idle for long stretches of time. A human operator, by contrast, would fill such gaps in the schedule with items the system has rejected. In practice, however, we find that the telescope is rarely given so few targets that it sits idle; instead, it is frequently oversubscribed, in which case the system must prune the candidate list in order to observe targets most efficiently. On the few occasions that a night's schedule does develop large gaps between objects, we add extra observations of standard stars (which are always useful) as filler.

7.2 Microscheduling

Deciding which targets to observe during a single night, and when, is a problem that usually does not have an easy solution; in fact, it is not clear that a single, "optimal" solution even exists. Whenever targets are bunched together, either on the sky or in time, one must make compromises in scheduling them. BAIT's approach is to try to observe each target at its "best" time, or as close to that time as possible.

The "best" time to observe an object may be defined in several ways. In the simplest case, it is the time when the object culminates on the meridian. Since one then looks through the smallest possible airmass, one (usually) gets the best image quality. However, there are complications. Objects which set soon after sunset, or rise just before dawn, cannot be observed when they cross the meridian because the Sun is still above the horizon. Some targets must be observed at a specific time or within a relatively narrow window—an occultation event, for example. A set of windows is required for those targets which need 10 or 20 observations throughout the night. Some requests might ask for images of several different objects (e.g., a variable star bracketed by standard stars) which lie in slightly different parts of the sky. If many targets are concentrated in a small region of the sky, as in the Virgo cluster or the Ophiuchus cloud, they may crowd each other as the region culminates.

Our solution is based on the idea of "windows of opportunity" for each target. Depending on its location, the times of sunrise and sunset, the maximum airmass at which observations are allowed, and any specific timing instructions from the user, one window (or occasionally two windows, when a circumpolar target exceeds the telescope's pointing limit just after sunset and comes back into view in the east just before sunrise) in time can be defined during which each object may be observed. Within that window, the "best" time is defined as the time at which the object's altitude is greatest. The problem then becomes picking a set of objects, and an order, in which each one is observed during its window and as close to the "best" time as possible.

The BAIT algorithm is relatively simple in theory, although nontrivial to implement. It requires a list of candidate requests, each of which has an observing window and a priority. First, the list is sorted by priority so that the most important objects appear at its head; objects with the same priority are placed in a random order, but all precede those with lower priorities. Next, a block of time is defined within which all observations must be made—initially, its termini are the times at which the Sun sets and rises beyond some altitude limit. The scheduler then considers each target in the list, starting with the most important one. It tries first to find a free block of time, large enough to accommodate the target's exposure length, at the "best" time for the target. If it succeeds, the object is scheduled for observation, and the block of time is reduced to smaller blocks before and/or after the observation. If there is no free block at the "best" time, the program searches outward from the "best" time, looking for any free block within the target's window of opportunity. If it finds such a free block, it schedules the target within the block, as before.

However, if this search also fails, then the program takes a different approach: it attempts to push previously scheduled targets to slightly earlier or later times in order to open up a contiguous space of free time sufficiently large to fit the current object. As before, the program tries to place the current object as close to its "best" time as possible. Scheduled targets may be shifted in time, but they may not be reordered, nor shifted so far that they move out of their window(s). A target with a low priority will therefore never cause one with a higher priority to be removed from the schedule. Although it is possible for a low-priority object to push a higher-priority object relatively far away from its original, optimal spot in the schedule, we find that this rarely occurs in practice. Unless the list of targets is severely concentrated in one region of the sky,

© Astronomical Society of the Pacific • Provided by the NASA Astrophysics Data System

the scheduling program initially chops the original, night-long block of free time into many small blocks of free time, separated by time allocated for observing targets. The large number of small blocks prevents any one target from being pushed far from its initial position, since small pushes in opposite directions tend to cancel each other.

After the scheduler has allocated objects for the entire night and exhausted the list of candidate targets, it simply waits for the night to begin, chooses the object with the earliest allocated time, and begins observing it. Upon completion of that observation, rather than keeping the same set of scheduled targets and doing them in order, the scheduler discards its list and starts all over again—with two small changes. First, it removes the object just observed from its pool of candidates; second, when creating the new "initial" block of free time, instead of using the start of night as the beginning of the block, it uses the current time. The program then follows exactly the same procedure described above for selecting objects from the remaining pool of candidates. Although this may seem like a waste of effort, it accomplishes several goals. First, if new requests have arrived since the last observation began, it is trivial to add them to the list of candidates and schedule them for the remaining portion of the night. Second, observations may not take the amount of time expected (the sum of exposure times plus some fixed overhead per slew and acquisition): if the telescope pointing is slightly maladjusted, acquiring a guide star may involve resetting the telescope coordinates by moving to a bright star. The resetting process may add 5 min to what could be a very short exposure, throwing the rest of the schedule behind. Bad weather, on the other hand, can cause an observation to last for a much shorter amount of time than expected: if the guide star cannot be acquired at all, then the observation is aborted. Selecting objects from scratch, just before each observation, avoids any timing problems that may occur. Third, starting over allows the system to go back and attempt again to observe objects that may have failed earlier in the night. When an object is selected, but fails to be acquired, it is replaced in the candidate pool and may again be selected at a later time, when the conditions may have improved. There is little time wasted running the scheduler repeatedly: the program takes only a few seconds to run on a Sun-2 workstation, even with a pool of about 50 candidates.

Weather information has been incorporated into the scheduler in a slightly indirect manner. Whenever the weather is bad, as directly indicated by the weather station or indirectly indicated by a string of successive failures to acquire guide stars, the scheduling program goes to sleep for an hour. At the end of that hour, it wakes up and tries again. By sleeping, the system avoids an excessive number of dome openings and closings, saving wear on many motors and gears. Since there is usually a period of increased cloudiness before a storm, giving up after a few successive failures tends to close the system down in advance of the rain that may come later.

The weather is important in a different way as well: some observations require the very best seeing or transparency, and should be scheduled only when conditions permit. BAIT does keep track of these parameters throughout the night by making a series of short exposures of standard stars whenever there is a small gap in the list of program objects. Software measures and averages the images' FWHM to determine the current seeing, and finds the difference between the brightness of the same stars on each frame as a measure of short-term variations in transparency. At Leuschner Observatory, at least, short-term variations in the brightness of a star are a good indication that the night is not photometric. Although the master computer has access to the most recently measured values of each parameter, we have not yet modified the scheduling program to check their values when it makes its list of targets for the current night. We hope to add this real-time capability soon.

TABLE 2
BAIT Performance, 1992 January–1992 August

Month	# images	# images/night	open hr	hr/night	avg hr/night	% time shutter open
Jan	674	22	33.4	1.7	12.1	0.14
Feb	300	11	17.9	1.6	11.3	0.14
Mar	597	19	41.4	2.8	10.1	0.27
Apr	1395	47	117.4	3.9	8.8	0.45
May	1534	49	93.3	3.2	7.6	0.43
Jun	1047	35	99.7	3.3	6.9	0.48
Jul	988	32	86.1	3.2	7.2	0.44
Aug	2293	74	134.3	4.1	8.4	0.49

Avoiding the Moon is a slightly easier task that has been implemented. The astronomer can specify how moonlight can affect her objects in two ways. First, she can add a line containing the MOONPHAS keyword to a request; observations of that request will then be scheduled only when the Moon is darker than the given phase (where 0.0 refers to new moon and 1.0 to full moon). By default, requests are given a value of MOONPHAS = 1.0. Second, she can request that no observation be made whenever the object falls within some angular distance from the Moon by adding the MOONDIST keyword to a request. By default, observing windows are calculated to avoid any period when a target lies within fifteen degrees of the Moon's position.

8. TELESCOPE PERFORMANCE

The first six months of BAIT operation, from 1991 December to 1992 May, were filled with problems caused by both software and hardware errors. However, the number of modifications gradually decreased and the telescope has been running virtually untouched since the beginning of the summer of 1992. We are now limited largely by the weather at Leuschner Observatory: few of the long winter nights are clear, and the hot summer season brings fog from the nearby ocean, which often covers the site as soon as the Sun sets. As Table 2 suggests, the best seasons appear to be spring and fall.

The mean exposure time for all images, including calibration frames, is about 240 s; excluding calibration frames, the mean rises to about 450 s. However, the great majority of exposure lengths are relatively short. About 75% of all exposures are less than 2-min long, or about 50% of all noncalibration exposures.

© Astronomical Society of the Pacific • Provided by the NASA Astrophysics Data System

TABLE 3
Observing Efficiency on Nights with Good Weather

Date	# images	obs. time (hr)	night length (hr)	% night open
3 August 1992	51	5.35	7.86	0.68
18 August 1992	65	5.55	8.47	0.65
21 August 1992	69	5.48	8.60	0.64

The overall observing efficiency (fraction of all possible night time during which the camera's shutter is open) has climbed to a respectable value of just a bit under 50%. It may fall somewhat in the winter, as the clouds return, but probably will remain higher than the 1992 January and February numbers, due to improvements in the system. For comparison, Benn and Martin (1987) found that the Isaac Newton Telescope (INT), a large (2.5 m) instrument operated on-site by an astronomer and night assistants, had an overall efficiency factor of about 36% over a 2-yr period. It is a little early to compare the two directly, because BAIT has not had a full winter season since the software has stopped evolving, but BAIT's efficiency will probably end up slightly higher than that of the INT. Benn and Martin state that the INT was used largely for spectroscopic work, and typically took longer exposures (half of them in the range 1000–1500 s, compared to half in the range 0–120 s for BAIT) of only a few objects per night. The overhead associated with slewing and acquiring new objects is therefore probably higher for BAIT, which makes its efficiency even more impressive.

The observing efficiencies mentioned above include time lost due to bad weather; Benn and Martin found that weather ruined about 16% of the total night time for the INT. When the weather is very good, a larger fraction of each night is spent collecting photons. We selected three very good nights during 1992 August and analyzed them separately; see Table 3. Even on nights when the weather was excellent, BAIT still spent at least 30% of the night on overhead (slewing, acquiring guide stars, reading out the CCD, etc.). The difference between BAIT's overall efficiency on clear nights alone and that on all nights is about $0.65-0.48=0.17$; thus, BAIT and the INT both lose about 16% of their time to the weather (albeit the figure for BAIT would probably be larger if a year-round average were taken).

BAIT is currently carrying out several scientific programs: monitoring some active galactic nuclei for many months, complementing space-based missions with visual observations from the ground, making light curves of active young stars, and providing data for undergraduate courses at U. C. Berkeley. One of its main goals is to make improved light curves of nearby supernovae. Papers describing observations of the peculiar Type Ia supernova 1991bg (Filippenko et al. 1992) and the Type IIb supernova 1993J (Richmond et al. 1993) provide some idea of the capabilities of BAIT in its current form.

9. FUTURE IMPROVEMENTS

We can make several changes to BAIT that should substantially increase its efficiency and the quality of its images. Replacing the current optics with a new set, based on a Ritchey–Chrétien design and a 0.75-m mirror, would double the collecting area. In addition, the Ritchey–Chrétien optics would yield much better images of guide stars far off-axis, allowing us to guide in even the star-poor regions near the Galactic poles. Moving the entire instrument to a site with better seeing would improve the accuracy of our photometry in crowded or complicated fields, and would enable us to study fainter objects or those requiring higher spatial resolution. Although we may install the new optics by the summer of 1993, BAIT will probably remain at Leuschner Observatory for a number of years. The cost of moving is unfortunately large.

The automatic control and scheduling system we have designed could be placed on other telescopes, either locally (as we have done with a 0.75-m reflecter at Leuschner) or at other sites. A network of small optical telescopes which all used the same protocol for communications would be a very powerful instrument, especially if the master computers at each site were able to communicate with each other. If there were bad weather at one site, the computer there could send urgent requests to other telescopes; these might inform each other after making each observation to avoid redundant exposures of the same objects. A network of sites covering a large range in longitude would enable astronomers to collect continuous streams of photometric data over many days, as is done with human operators using the Whole Earth Telescope observing network (e.g., Nather et al. 1990; Clemens et al. 1992). Adapting the BAIT control system to another telescope might not be prohibitively difficult, since a large part of the site-specific information is collected into the central configuration file.

Using the BAIT software, or similar systems, one might be able to automate other kinds of telescopes: those with near-infrared detectors are very similar to optical telescopes and would present few new problems. We have considered placing a fiber-fed spectrograph at the focus of one of our telescopes and creating an automatic spectroscopic instrument, but it is clear that we would need to make significant changes to the hardware, and some to the software as well.

This project could not have begun without the NSF Presidential Young Investigator Award (grant AST-8957063) to A.V.E. We are very grateful to Russ Genet for a generous donation from AutoScope Corporation. We have also received equipment donations from IBM, Photometrics Ltd., and Sun Microsystems, as well as financial support from NSF Cooperative Agreement AST-8809616 (Center for Particle Astrophysics, U.C. Berkeley), NSF grant AST-9115174, CalSpace grant CS-96-91, and the University of California at Berkeley. We have had helpful conversations with numerous people, especially Russ Genet and Jack Borde. Many of our developments have utilized experience and knowledge gained by the Berkeley Automated Supernovae Search Team, which was led by Richard Muller, Carl Pennypacker, and Saul Perlmutter. In particular, Ned Hamilton developed many details of the telescope control, and Robbi Smits contributed much to the modernization of the 0.5-m telescope. We thank Jerome Hudson for his insights into telescope control philosophy, Suzanne Jacoby for providing the NOAO keywords,

© Astronomical Society of the Pacific • Provided by the NASA Astrophysics Data System

and Bruno Leibundgut for comments on a draft of this paper. M.W.R. is grateful to Joseph C. Shields for many useful discussions. The excellent machining of the CCD camera by Lockwood Associates is appreciated.

REFERENCES

Allen, C. W. 1973, Astrophysical Quantities, 3rd ed. (London, Athlone)

Barrett, J. D., Thomas, R. C., and Baruch, J. E. F. 1991, PASP, 103, 1218

Benn, C. R., and Martin, R. 1987, QJRAS, 28, 481

Bessell, M. S. 1990, PASP, 102, 1181

Clemens, J. C., et al. 1992, ApJ, 391, 773

Code, A. D. 1992, in Robotic Telescopes in the 1990s, ed. A. V. Filippenko (San Francisco, ASP), ASP Conf. Ser. 34, p. 3

Colgate, S. A. 1982, in Supernovae: A Survey of Current Research, ed. M. J. Rees and R. J. Stoneham (Dordrecht, Reidel), p. 319

Colgate, S. A., Moore, E. P., and Carlson, R. 1975, PASP, 87, 565

Filippenko, A. V. 1992a, ed., Robotic Telescopes in the 1990s (San Francisco, ASP), ASP Conf. Ser. 34

Filippenko, A. V. 1992b, in Robotic Telescopes in the 1990s, ed. A. V. Filippenko (San Francisco, ASP), ASP Conf. Ser. 34, p. 55

Filippenko, A. V., and Richmond, M. W. 1991, in Advances in Robotic Telescopes, ed. M. A. Seeds and J. L. Richard (Mesa, Fairborn), p. 89

Filippenko, A. V., Richmond, M. W. Branch, D., Gaskell, C. M., Herbst, W., Ford, C. H., Treffers, R. R., Matheson, T., Ho, L. C., Dey, A. Sargent, W. L. W., Small, T. A., and van Breugel, W. J. M. 1992, AJ, 104, 1543

Ford, C. H. et al. 1993, AJ, 106, 1101

Genet, R. M. 1986a, in The Study of Variable Stars Using Small Telescopes, ed. J. R. Percy (Cambridge, Cambridge University Press), p. 235

Genet, R. M. 1986b, in Automatic Small Telescopes, ed. D. S. Hall, R. M. Genet, and B. L. Thurston (Mesa, Fairborn), p. 1

Genet, R. M., and Hayes, D. S. 1989, Robotic Observatories (Mesa, AutoScope Corp)

Greisen, E. W., and Harten, R. H. 1981, A&AS, 44, 371

Grosbol, P., Harten, R. H., Greisen, E. W., and Wells, D. C. 1988, A&AS, 73, 359

Harten, R. H., Grosbol, P., Greisen, E. W., and Wells, D. C. 1988, A&AS, 73, 365

Honeycutt, R. K., and Turner, G. W. 1992, in Robotic Telescopes in the 1990s, ed. A. V. Filippenko (San Francisco, ASP), ASP Conf. Ser. 34, p. 77

Jenkner, H., Lasker, B. M., Sturch, C. R., McLean, B. J., Shara, M. M., and Russell, J. L. 1990, AJ, 99, 2082

Johnston, M. D. 1988, in Very Large Telescopes and their Instrumentation (Garching, ESO), p. 1273

Lasker, B. M., Sturch, C. R., McLean, B. J., Russell, J. L., Jenkner, H., and Shara, M. M. 1990, AJ, 99, 2019

Maran, S. P. 1967, Science, 158, 867

McCall, J. F., Miedaner, T. L., and Code, A. D. 1968, AJ, 73, 756

Nather, R. E., Winget, D. E., Clemens, J. C., Hansen, C. J., and Hine, B. P. 1990, ApJ, 361, 309

Pearce, E. C. 1986, in Automatic Photoelectric Telescopes, ed. D. S. Hall, R. M. Genet, and B. L. Thurston (Mesa, Fairborn), p. 139

Perlmutter, S., Muller, R. A., Newberg, H. J. M., Pennypacker, C. R., Sasseen, T. P., and Smith, C. K. 1992, in Robotic Telescopes in the 1990s, ed. A. V. Filippenko (San Francisco, ASP), ASP Conf. Ser. 34, p. 67

Richmond, M. W. 1990, in Proceedings of the Electronics Oriented Astronomy Seminar 3, ed. J. Sanford (Costa Mesa, OCA), p. 68

Richmond, M. W. 1991, IAPPP, 45, 80

Richmond, M. W., and Filippenko, A. V. 1991, in Advances in Robotic Telescopes, ed. M. A. Seeds and J. L. Richard (Mesa, Fairborn), p. 81

Richmond, M. W., Treffers, R. R., and Filippenko, A. V. 1992, in Robotic Telescopes in the 1990s, ed. A. V. Filippenko (San Francisco, ASP), ASP Conf. Ser. 34, p. 105

Richmond, M. W. et al. 1993, AJ, submitted

Russell, J. L., Lasker, B. M., McLean, B. J., Sturch, C. R., and Jenkner, H. 1990, AJ, 99, 2059

Schmidt, E. G. 1992, in Robotic Telescopes in the 1990s, ed. A. V. Filippenko (San Francisco, ASP), ASP Conf. Ser. 34, p. 73

Smith, C. K., Crawford, F., Muller, R., Pennypacker, C., Perlmutter, S., Sasseen, T., Williams, R., and Treffers, R. 1988, in Automatic Small Telescopes, ed. D. S. Hayes and R. M. Genet (Mesa, Fairborn), p. 47

Treffers, R. R., Richmond, M. W., and Filippenko, A. V. 1992, in Robotic Telescopes in the 1990s, ed. A. V. Filippenko (San Francisco, ASP), ASP Conf. Ser. 34, p. 115

Wells, D. C., Greisen, E. W., and Harten, R. H. 1981, A&AS, 44, 363

© Astronomical Society of the Pacific • Provided by the NASA Astrophysics Data System