Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Physics, Computer Science & Mathematics Division

Presented at the North American Treaty Organization/ Advanced Study Institute on Supernovae, Cambridge, England, June 28 - July 10, 1981

THE BERKELEY AUTOMATED SUPERNOVA SEARCH

Jordin T. Kare, Carlton R. Pennypacker, Richard A. Muller, Terry S. Mast, Frank S. Crawford, and M. Shane Burns

June 1981

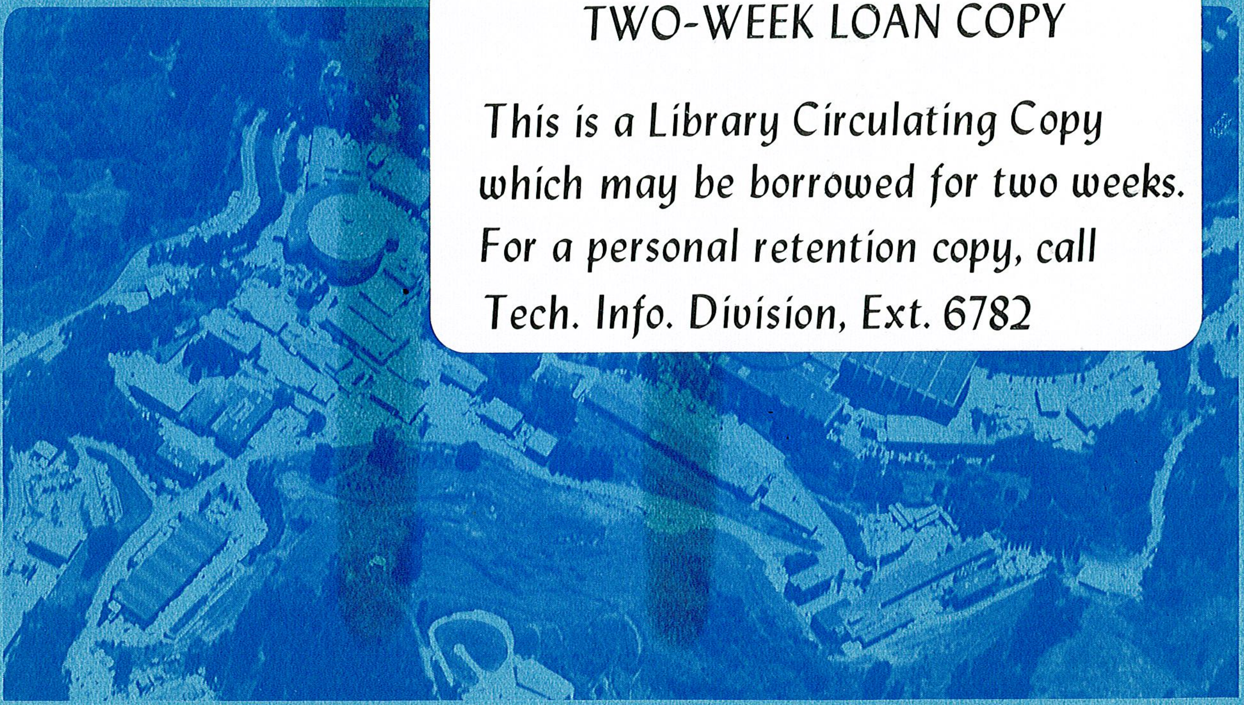

Prepared for the U.S. Department of Energy under Contract W-7405-ENG-48

THE BERKELEY AUTOMATED SUPERNOVA SEARCH

Jordin T. Kare, Carlton R. Pennypacker, Richard A. Muller,
Terry S. Mast, Frank S. Crawford, and M. Shane Burns

Lawrence Berkeley Laboratory
University of California, Berkeley 94720

The Berkeley automated supernova search employs a computer controlled 36-inch telescope and charge coupled device (CCD) detector to image 2500 galaxies per night. A dedicated minicomputer compares each galaxy image with stored reference data to identify supernovae in real time. The threshold for detection is m_v=18.8. We plan to monitor roughly 500 galaxies in Virgo and closer every night, and an additional 6000 galaxies out to 70 Mpc on a three night cycle. This should yield very early detection of several supernovae per year for detailed study, and reliable pre-maximum detection of roughly 100 supernovae per year for statistical studies. The search should be operational in mid-1982.

1. INTRODUCTION

Our observational knowledge of supernovae is still quite limited. Of the several hundred supernovae identified, less than 40% have even been classified as to type. Fewer still have good light curves or spectra, and virtually none have been observed well before maximum light. Yet the number of supernovae found per year is declining, as photographic surveys are terminated. The reason is that such surveys, while valuable, cannot gather the data needed to produce a qualitative increase in our knowledge of supernova explosions.

New searches are needed which can i) catch a few supernovae very early, for detailed study, ii) find a large number of nearby supernovae in a known sample of galaxies, and iii) find a very large number of distant supernovae for statistical and cosmological purposes.

This manuscript was printed from originals provided by the author.

The Berkeley automated supernova search is uniquely qualified to satisfy the first two of these goals, and is complementary to several other automated searches in achieving the third. Based on ideas proposed by Stirling Colgate, our system uses a computer controlled telescope and solid state imaging detector to collect images of up to 2500 galaxies per night. Images are processed in real time, to allow prompt detection and confirmation of supernovae and elimination of most false alarms. When fully operational, we should detect typically 100 supernovae per year well before maximum light, and several nearby supernovae per year at a fraction of a percent of maximum. Broadband and possibly standard filter (UBV) photometry will be performed directly, and cooperating observatories, both ground and space based, will be able to obtain spectral information at all frequencies from very early times.

2. A BRIEF REVIEW OF PAST AND PRESENT SEARCHES

The original search for extragalactic supernovae was implemented by Zwicky (1) on the 18-inch Palomar Schmidt in 1937, and extended using the 48-inch Schmidt in 1959. The search on the 48-inch has stopped, but photographic searches are continuing on the 18-inch and at other observatories (2,3). Schmidt plates of selected regions of the sky are exposed at regular intervals and compared with older plates using a blink comparator or film overlay technique. A large area of sky is covered by each plate, but the low quantum efficiency of photographic plates implies fairly long exposure times. The comparison process is tedious, and its detection efficiency is unknown — some supernovae are likely to be missed. The limited dynamic range of photographic plates leads to the loss of supernovae occurring against the bright cores of galaxies. However, the major limitation of photographic searches has been the long interval between successive plates, plus the long time required for analysis. In many cases, supernovae are not found until after they have faded into invisibility, precluding any further observations.

Hynek (4) employed an image tube detector and video comparison system to improve both sensitivity and speed of detection, and succeeded in finding several supernovae. However, the actual identification of supernovae was still done by eye. This search has been discontinued.

In recent years, several groups have turned to computers to analyse images. By digitally subtracting images, the sudden increase in brightness representing a supernova can be made conspicuous, even to a computer. Kibblewhite (5) has retained the photographic plate, but uses automated plate analysis, and has recently begun finding supernovae. This system should find large numbers of distant supernovae near maximum light, but currently has turnaround times (including plate shipment) of several days. Another system, under development by Angel and Sargent, employs a pair of charge coupled device (CCD) detectors on a fixed transit telescope to provide very deep scans of a narrow strip of sky every night (6). This system should be ideal for finding supernovae at high red shifts for cosmological purposes, but covers too small an area of sky to find nearby supernovae.

As early as 1971, Stirling Colgate began work on an automated telescope which could swing rapidly among a large number of foreground galaxies (7). An image tube supplied images to a remote computer for immediate processing. This system was quite ambitious, and proved to strain the limits of the available technology and manpower; work on it is still proceeding but it has not yet gone into operation.

Our own project is a direct descendant of the Colgate search, stimulated by the recent spectacular progress in solid state electronics. High performance CCD detectors, minicomputers cheap enough to dedicate to the search, and microcomputers for telescope control have all become available in the past few years. With the aid of limited funding from a large number of sources, we have assembled a group including one full time and several part time physicists, two graduate students, and a full time senior programmer. We have designed and built most of the special hardware required, and in collaboration with the Monterey Institute for Research in Astronomy (MIRA) have upgraded the MIRA 36 inch telescope drives for computer control. We are currently testing our hardware and writing software for the search, and we expect to be operational in mid 1982.

3. OVERVIEW OF THE BERKELEY SYSTEM.

Figure 1 is a diagram of the major components of our system. Here we outline the properties of the various items.

A telescope in the 24 to 36 inch range is needed to provide reasonable sensitivity for the search. We have been fortunate in arranging to use the MIRA 36 inch telescope, an equatorially mounted f/10 Cassegrain. The telescope Guidance/Acquisition Package (GAP) contains a moveable diagonal which allows rapid switching among several detectors — in particular, between our CCD and a Reticon-based spectrophotometer for "instant" spectra of bright objects.

The telescope employs a roller drive in right ascension, and has a maximum slew rate of 11 degrees per second. In collaboration with MIRA, we are installing a digitally controlled servo system which will slew the telescope smoothly over 1-2 degrees in under 3 seconds. The main slew motors drive the telescope to within roughly 1 arc minute of the desired position; then pneumatic clutches engage the stepping motor drives to bring the telescope within a few arc seconds. The eventual aim is to have the telescope controller (a HEX microcomputer) calculate all refraction and flexure corrections and provide a "virtual telescope" capable of arc second pointing accuracy anywhere in the sky. Since our software calculates the telescope position on every field, errors do not accumulate, and a much more modest pointing accuracy of about 15 arc

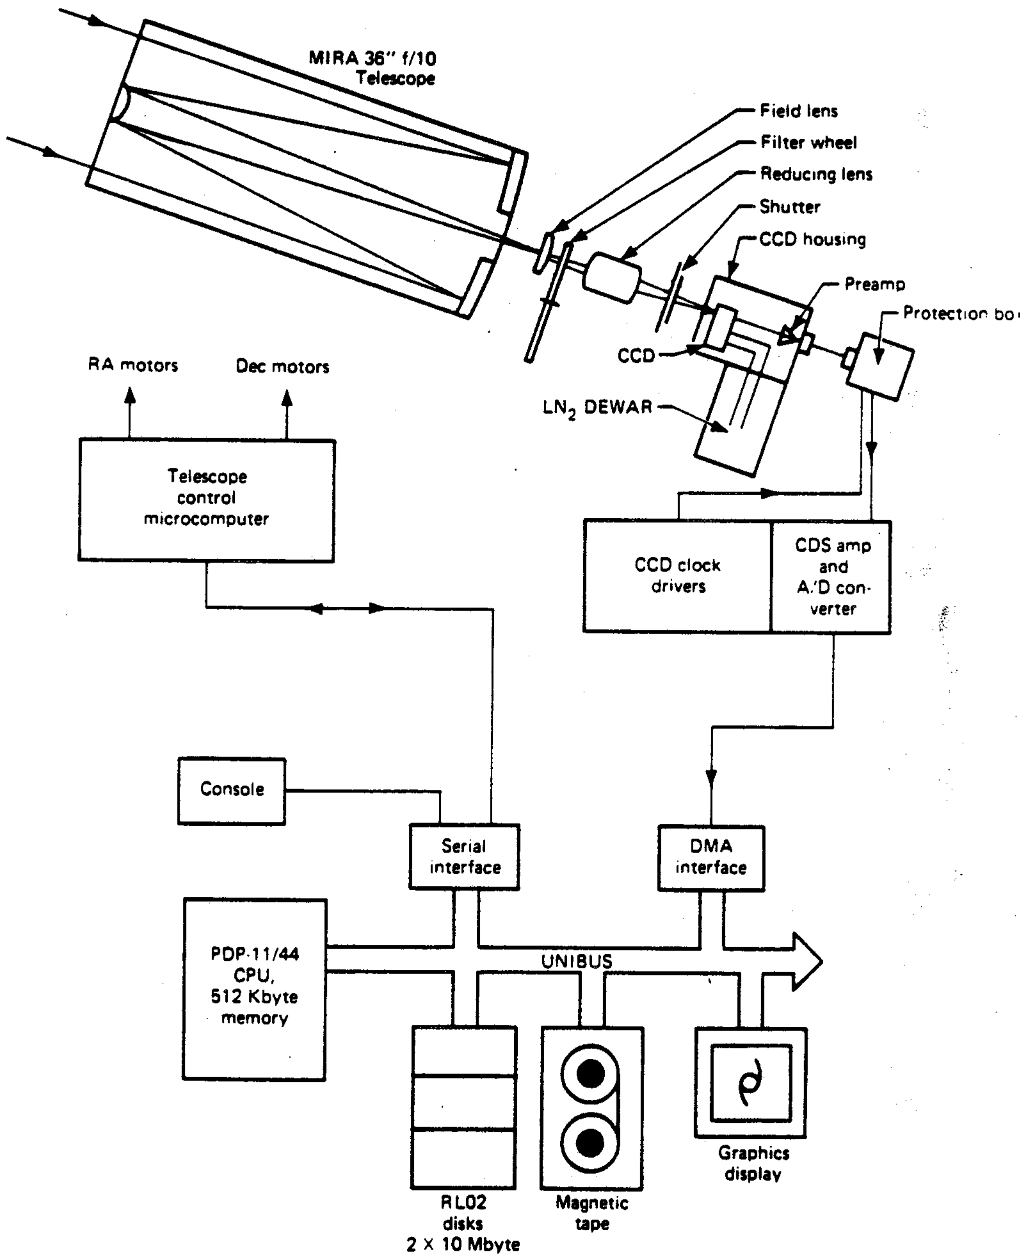

Figure 1 — The Berkeley Automated Supernova Search System

seconds over a 2 degree move is sufficient. The telescope control system is diagrammed in figure 2. In consultation with Frank Melsheimer, the telescope designer, we have determined that telescope wear will not be a problem, being concentrated mainly in easily replaced components such as teflon o-rings.

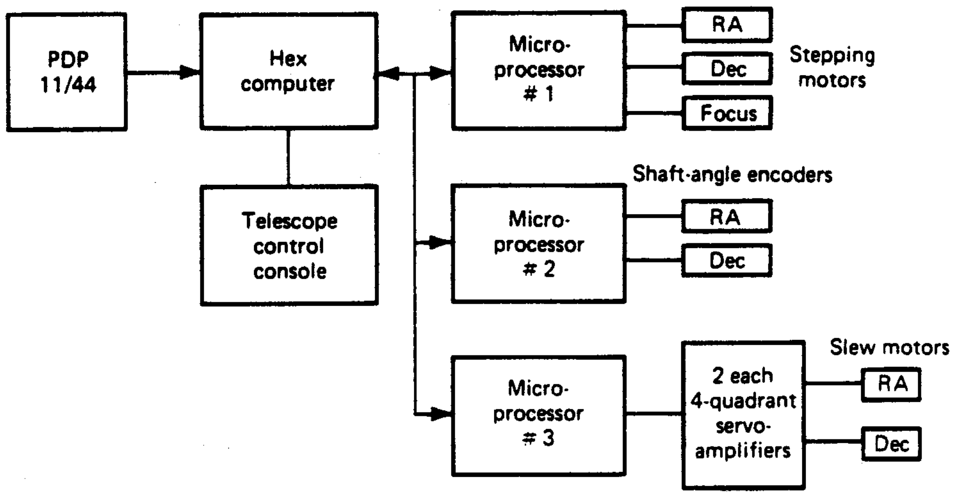

Figure 2 — MIRA Telescope Control System

The heart of the system is the CCD detector, RCA type number SID 53612. It is 1.5 cm by 1 cm, and is organized as an array of 512 by 320 photosensitive elements (pixels). Each pixel is 30 microns square. Photons striking the CCD generate photoelectrons which accumulate in each pixel until read out. Our CCD is a buried channel, thinned device which provides extremely high quantum efficiency and good blue response: the spectral response is shown in figure 3. At the end of each exposure, the collected electrons are shifted out one pixel at a time in a raster scan pattern, and detected by an on-chip FET amplifier. The CCD is cooled to -112 C by liquid nitrogen, with a feedback control system to maintain a

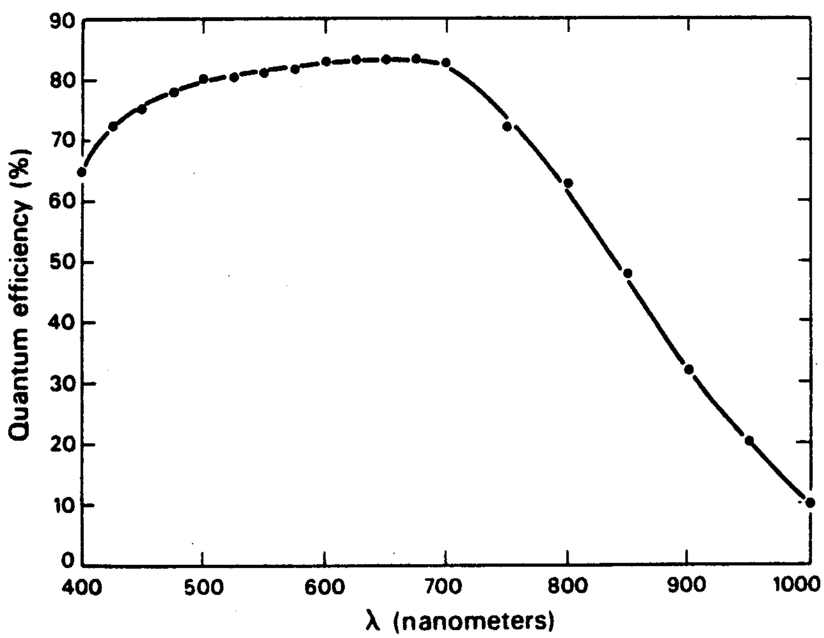

Figure 3 — Quantum Efficiency of RCA CCD

constant temperature. An evacuated housing provides insulation and protection from condensation. At such temperatures there is no detectable dark current, even over several hour exposures. Noise is introduced primarily by the output amplifier. The particular CCD we plan to use has a typical noise level equivalent to 40 electrons RMS per pixel, with substantial variation from chip to chip. Lower noise amplifiers are being developed; for now, we use the figure of 40 electrons. The pixel well depth, or charge capacity before overflow, is roughly 200,000 electrons, yielding a dynamic range approaching 5,000 to 1.

The telescope image is focused onto the CCD using a field lens and a Nikon reducing lens. The final effective f/number is f/2.25, which corresponds to a "CCD scale" of 3 arc seconds per pixel, or a total field of 16 by 25 arc minutes. This scale can be reduced as needed; see discussion below. A filter wheel can carry several filters for photometric purposes. Normally, a short-pass filter removes wavelengths longer than 750 nm to eliminate atmospheric OH lines which form interference fringes in the thin CCD chip. A long pass filter can also be installed to reduce the transmission of scattered moonlight on nights of full moon. A shutter, controlled by the computer, covers the CCD between exposures.

The output of the CCD is fed to a preamplifier inside the vacuum housing, and then to an electronics package (card cage) which houses a correlated double sampling amplifier. The correlated double sampling process subtracts noise associated with the CCD amplifier reset pulse from the signal. An analog to digital converter translates the amplified signal from each pixel to a 14 bit serial binary word. The card cage also contains the circuits which generate the precise clock pulses needed to drive the CCD. Signals from the card cage are routed through a "protection box" containing fuses and Zener diode clipping circuits, to reduce the risk of CCD damage due to circuit failure or transients. We have copied the circuitry used by John Geary of the Harvard-Smithsonian Observatory (8) to minimize the problems and risks of getting the CCD operating. Geary's CCD has been in use at the Mt. Hopkins observatory for over a year, so the design is well-tested.

The digital data from the card cage are fed into our computer via a custom-built buffer board and a standard Direct Memory Access (DMA) interface. We selected the PDP-11/44 computer from Digital Equipment Corporation as providing a large main memory space in a small physical space. The computer proper occupies a single foot-high chassis; the complete system of processor, disks, and power supplies fits easily into two 4 foot racks. We have 512 kilobytes of main memory, sufficient to store an entire image (160,000 pixels or 320 kilobytes) along with analysis software; this greatly reduces the time spent accessing disk memory. Two 10 megabyte disks provide mass storage for program development and for a limited number of reference images (see software discussion below); a magnetic tape drive for additional storage and data logging will be acquired soon. A large Grinnell graphics display with gray scale (512 x 512 pixels, 8 bits/pixel) allows monitoring of images for a variety of purposes.

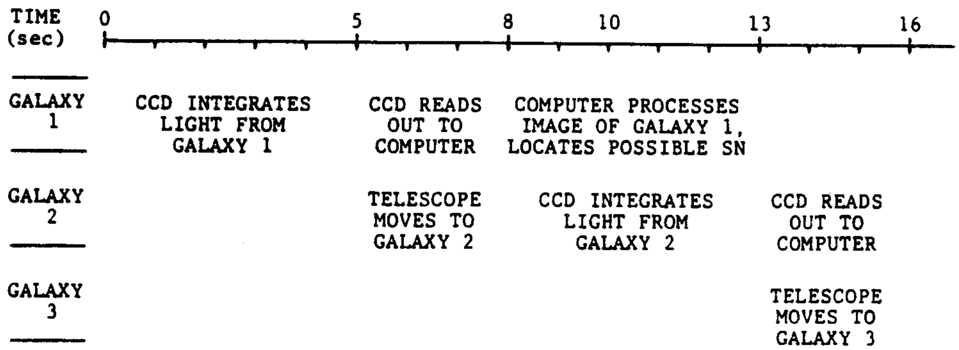

Figure 4 — Automated Search Sequencing

4. IMAGE ANALYSIS

Figure 4 shows the sequence of events during automated searching. The telescope is pointed at the first galaxy, and the CCD is exposed for a few seconds. As the telescope moves to the next galaxy, the CCD is read out and the image stored in the computer. During the new exposure, the stored image is processed to identify possible supernovae. If any are found, the computer can instruct the telescope to return to the galaxy in question and take a second look; otherwise the image is discarded and a new image read in, while the telescope moves to yet another galaxy.

The actual processing of the image is outlined in figure 5. An initial scan through the image locates bright pixels and collects them into a list of bright objects, mostly foreground stars brighter than 14th magnitude, of which there are typically 10 in each field. These stars are then matched against a reference list of stars to determine the telescope "offset" from the exact position desired. (This list, and all other reference data, are generated from CCD images acquired in the early stages of the search, rather than, for example, from catalogs). This offset is fed back to the telescope control software, so that pointing errors do not accumulate. However, the offset is used primarily to align the current image with a stored reference image for subtraction of constant features. To this end, simple fitting routines locate the position of each bright star to typically 0.2 pixels (0.6 arc sec). By averaging over 10 to 20 stars, the offset can be found to better than ±0.1 pixels.

Once the offset is known, the location of the galaxy of interest on the CCD is known also. Since the angular size of even nearby galaxies is typically less than the field of view of the CCD, substantial computing time is saved by processing only a rectangular section of the image containing the galaxy (with some margin to catch supernovae beyond the visual "edge" of the galaxy). Typically 10% of the full CCD field would be processed, although as the software is optimized such restrictions should become unnecessary. This processing involves:

-subtraction of amplifier bias
-subtraction of sky background, measured on a "dark" patch of sky in each field
-removal of faulty pixels and overflow from very bright stars

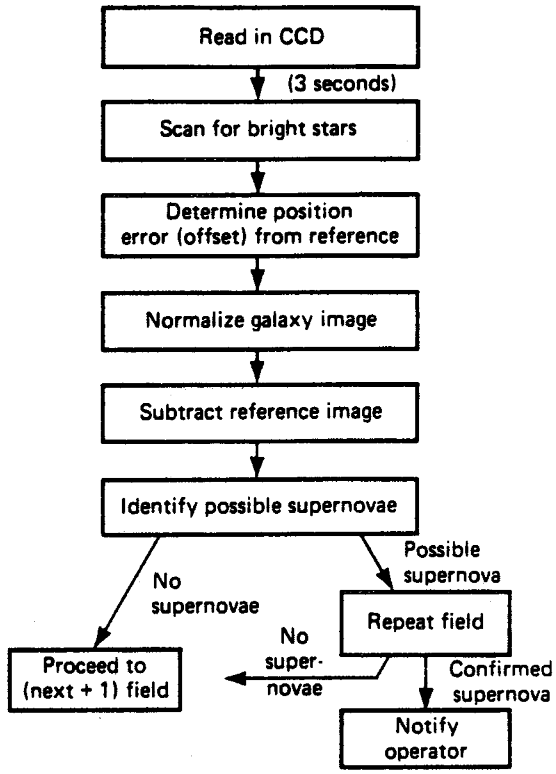

Figure 5 — Automatic Search Program

-scaling for pixel quantum efficiency variations (flat fielding)
-scaling for sky transmission variations, calculated from foreground star brightness
-alignment to the nearest whole pixel with the reference field.

Once this processing is complete, the system can search for supernovae in the active field. A variety of techniques can be used to store and compare images. We originally intended to simply store an image of the entire galaxy for reference, but this consumes a large amount of storage and is highly sensitive to slight misalignments. Since only the central portions of galaxies, and foreground stars, are detectable above the sky background with our system, we currently plan to store a list of detectable point objects, of which there may be a few hundred, plus very small maps of tens or hundreds of pixels around extended features such as galaxy cores or jets. The software will thus process three types of pixels:

-"ordinary" pixels, with a fixed threshold set roughly 5 standard deviations above mean sky background,
-"stars", which will be either disregarded as possible supernovae (thus eliminating variable stars) or given individual thresholds.
-"blobs", including galaxy cores, from which reference maps will be subtracted.

Based on simulations using 1-dimensional scans across galaxies obtained from J. Geary, we estimate that even very simple interpolation between pixels allows good subtraction with less than a factor of 2 increase in threshold, except in the exact cores of galaxies, which must be handled separately. A number of more sophisticated techniques may be introduced once the search is operational. For example, the telescope optics can be easily shifted to provide a "high resolution" reference picture with a pixel size of 1 or 1.5 arc seconds. Such a picture could be processed using Fourier transform techniques to match both fractional offsets and seeing variations, and would allow very accurate subtraction.

5. SYSTEM PERFORMANCE

5.1. Sensitivity and parameter selection

The noise level of our detection system is set by two factors: the CCD readout noise, and the statistical $n^{1/2}$ fluctuations in the number of photons received. The CCD noise statistics are assumed to be Gaussian, with a standard deviation (RMS value) of n_{rn} electrons. For our current CCD, we take n_{rn}=40. Away from the cores of galaxies, most of the background photons come from diffuse sky light, which at the location of the MIRA telescope is equivalent to one 21st magnitude star per square arc second. By comparison with a known photometric system (8), we calculate that we collect approximately 30 photoelectrons (pe) per second from a star of m_v=20, so we expect 12 pe/sec/arc sec^2, or N_{sky} = 108 pe/sec/pixel. These noise sources are independent, so our total noise for a single pixel is

$$n_{tot} = (n_{rn}{}^2 + N_{sky} \cdot t)^{1/2} \qquad (1)$$

We choose an integration time t of 5 seconds. Shorter times clearly make inefficient use of telescope time, since the telescope requires 3 seconds to move between galaxies. Longer times allow fainter supernovae to be detected, but reduce the number of galaxies observable each night. The gain in sensitivity is decreased once the sky noise dominates the CCD noise. For our current system, $(n_{rn})^2$ is 1600, so the CCD noise dominates for exposures up to 15 seconds, but we hope eventually to obtain quieter CCD's. Meanwhile, 5 second exposures provide quite good sensitivity.

The signal from a supernova (or other object) of magnitude m (disregarding color variations) would be

$$N_{supernova} = 30 \text{ pe/sec} \times 10^{(20-m)/2.5} \times t \qquad (2)$$

If a fraction f of this light falls into a single pixel, the "signal to noise" ratio for a single pixel would be

$$S/N = N_{supernova} \times f / n_{tot} \qquad (3)$$

In order to make the fraction f large, it is desirable to make the "pixel scale" as large as possible, at least until the pixel size is equal to the expected seeing disk. Beyond that point, raising the pixel size increases the sky background noise while increasing the mean value of f (i.e. decreasing the chance that the supernova blur disk will fall on a pixel boundary) only slowly. However, as long as the CCD noise dominates the total noise, larger pixels are desirable. Large pixels also minimize computing time and increase our field of view. Our chosen value of 3 arc seconds/pixel is limited by vignetting in the MIRA telescope optics and by the f/number obtainable with standard optics. The expected seeing at the current telescope site is 2-3 arc seconds, and when the telescope moves to its permanent site on Chew's Ridge, the seeing should typically approach 1 arc second. The mean value of f is therefore close to 1. For the following, we take a conservative value of 0.5, so that half the light falls in one pixel.

Because of the large number of pixels involved, relatively rare statistical fluctuations can cause unacceptable numbers of false alarms. Setting a detection threshold of 5 standard deviations would, assuming ideal Poisson statistics, yield 1 false alarm per 1.7×10^6 pixels, or, at 10^4 pixels per galaxy, one per 170 galaxies. Since such false alarms are caught on the second look, this is a reasonable rate. In practice, we will vary the threshold to obtain a suitable false alarm rate. However, using S/N = 5, we can find the limiting magnitude for the search:

$$(20 - m_{limit})/2.5 = \log_{10} (S/N \times n_{tot})/(f \times t \times 30 \text{ pe/sec}) \tag{4a}$$

$$m_{limit} = 18.8 \tag{4b}$$

5.2 False alarms

In addition to false alarms caused by random noise, there will be a number of distinct events which can "look like" a supernova. Many of these can be eliminated by a prompt second look at the field, something possible only with a system of this type.

Cosmic rays will generate traces on our CCD an average of 1.5 to 3 times per minute, but only 1/10 of these will fall in the area of the galaxy. Thus the cosmic ray false alarm rate will be 1 per 25 to 50 fields. This figure may be reduced by adding simple software tests to rule out extended or improbably bright "supernovae."

Variable stars will be listed as stars if they spend much time above our threshold; we will probably monitor them as interesting objects in their own right. Cataclysmic variables which brighten drastically will cause false alarms on their first eruptions, but should be very rare. In the ecliptic, asteroids occur with a surface density of roughly 6 per square degree to 19th magnitude, enough to trigger a false alarm once in 40 fields. Unfortunately, Virgo is in the ecliptic; outside Virgo 10% to 20% of possible fields are affected. The proper velocity of asteroids is not enough to distinguish them promptly, but after 10 minutes they will move one or more pixels, and can be rejected by a third, delayed check. It may also be possible to identify them by their reflected solar spectrum.

At full moon, scattered moonlight raises the sky background to about 900 photoelectrons/pixel/second, increasing the system noise level by a factor of close to 2 for a 5 second exposure. The light level drops sharply away from full moon, however, so we lose only 2 to 3 days per month. A 500 nm long pass filter can be used to block as much as 90% of the scattered light with a loss of less than half of the supernova signal: this may allow low thresholds even during full moons.

5.3 Search strategies

With t = 5 seconds, and a telescope move time of 3 seconds, we can examine 450 fields per hour, or 2700 per six hour night. Allowing for 200 repeated fields (a 6% false alarm rate) we can check 2500 fields per night. There are three options available, corresponding to the three areas where data are needed:

-search 2500 closest galaxies every night to catch close supernovae very early,

-search 2500 x n nearby galaxies every n^{th} night (n   5) to catch a moderate number of supernovae well before maximum, or

-search distant-limited sample of rich clusters, with typically 10 galaxies/CCD field, to obtain large numbers of distant supernovae near maximum light.

We plan initially to employ a mix of the first two options. We will examine galaxies in the Virgo cluster and closer (5-700 fields) every night, and an additional sample of 6000 galaxies at an average distance of 50 Mpc (H_o = 100 in all calculations) on a three night cycle. The third option depends critically on the system sensitivity, and requires additional software. In addition, other searches are competitive with ours in this area. We therefore ignore the "cluster strategy" for the moment.

We assume that the peak luminosity of a Type I supernova is -18.5, and that extinction due to dust in our galaxy and the supernova parent totals 0.2 magnitude. The observed magnitude at peak in Virgo (20 Mpc) is then m = 13.3. We can thus detect supernovae in Virgo 5.5 magnitudes below maximum, or at 0.6% of maximum light — lower if we choose to extend exposure times to 10 seconds. Type II supernovae are roughly 1.5 magnitudes fainter, and could be detected at 2-3% of maximum.

Since we check Virgo galaxies every night, an average of 12 hours elapses between the time a supernova crosses our threshold and the time we see it. The shape of the light curve is completely unknown at these early times, so the resulting increase in brightness at detection is uncertain; a factor of 2 in 24 hours is a reasonable guess, so the average brightness at detection may be a factor of 1.5 (0.4 magnitude) brighter than the minimum. Supernovae occurring in or near galaxy cores will not be detected until somewhat later, depending on the quality of our subtraction, but they will be detected well before maximum light.

The supernova rate in various types of galaxies is still uncertain. For convenience we assume 1 per galaxy per 50 years, a reasonable estimate. If we search 500 galaxies in Virgo, we would expect to find 10 supernovae per year. This number is sufficient to provide work for spectroscopists, satellite observers, and others interested in detailed observations, but too small to add appreciably to the statistical data base. We therefore also observe 6000 nearby galaxies selected from the Nilson and other catalogs. Thus we expect to see over 100 supernovae per year located in a known set of galaxies. This fixed sample should greatly reduce the statistical problems of brightness limited samples.

In these more distant galaxies, our detection thresholds are higher: at 50 Mpc we detect Type I supernovae only 3.5 magnitudes before maximum, and in the worst case a type II supernova at 70 Mpc would be detected only 0.5 magnitudes before maximum. The longer interval between scans of a given galaxy would also raise the average brightness at detection. Still, all supernovae not heavily obscured should be detected, and virtually all (especially type I's) detected several days before maximum.

Because the entire search, including the selection of galaxies, is governed by easily modified software, it will be straightforward to modify the search strategy as we learn more, or to meet specific requests. For example, if galaxies of a particular type or orientation prove interesting, a larger portion of the search can be devoted to them.

5.4 Additional observations

The primary function of our system in connection with the closest supernovae will be to serve as an alarm system, allowing other observatories to collect optical spectra and information outside the visible range at very early times. The very low level at which we detect supernovae may lead to entirely new observations; for example, detection of underluminous supernovae, up to 100 times fainter and 10 times rarer than type I's, which have been proposed to explain Cas A.

We will make nightly photometric measurements on all previously discovered supernovae, and on other objects of interest, such as variable stars. We can achieve reasonable accuracy through the use of foreground stars in the same field as standards. A limited number of fields could be examined repeatedly, using standard filters to obtain UBV data, or even using a grating to obtain rough spectra. Such photometry will be of particular interest in following the very early part of the light curve — for example, in catching the "glitch" in the Type I light curve which might occur when the shockwave reaches the stellar surface.

Finally, on more distant supernovae of primarily statistical importance, we expect to be able to provide light curves sufficient for classification by type, and approximate position within the parent galaxy. Position measurements accurate to about 1 arc second should be achievable by averaging over several observations. This should suffice for most applications, but a photographic backup will be desirable for maximum precision.

6. CURRENT STATUS

As of 1 August 1981, our hardware has been largely completed, and is being tested. The CCD driver electronics have been completely built and tested with an electronic CCD simulator. The computer has been operating for several months, and the the CCD interface has been debugged. Images are being read from the CCD simulator and displaved on our graphics display.

The CCD detector head and dewar have been tested for vacuum tightness and thermal control, and we anticipate plugging in our current "loaner" CCD chip very soon. The telescope drive servo amplifiers have been designed and breadboarded, and final versions are being built. The microprocessor controllers are currently being programmed by MIRA.

The magnitude of the required software effort is indicated by figure 6; the real time analysis software alreadv discussed must be surrounded with support programs. A substantial amount of software, between one and two man years' worth, has been written. However, we estimate that an even larger amount remains to be written. When hardware testing is complete, we will have three to four people programming essentially full time.

We have a commitment from MIRA for one night of telescope time per week for initial testing. Once our hardware is operating, we are guaranteed 3 nights per week for six months, sufficient to determine the efficiency of our search and find a significant number of supernovae.

Our timetable is as follows:

1 September 1981
: Hardware complete; CCD testing in lab

1 January 1982
: Telescope drives ready; CCD moves to telescope. Image data transferred to Berkeley via videotape for software development.

1 April 1982
: Computer moves to telescope. Minimal software complete. Automatic scanning tests begin.

4 July 1982
: SN 1982

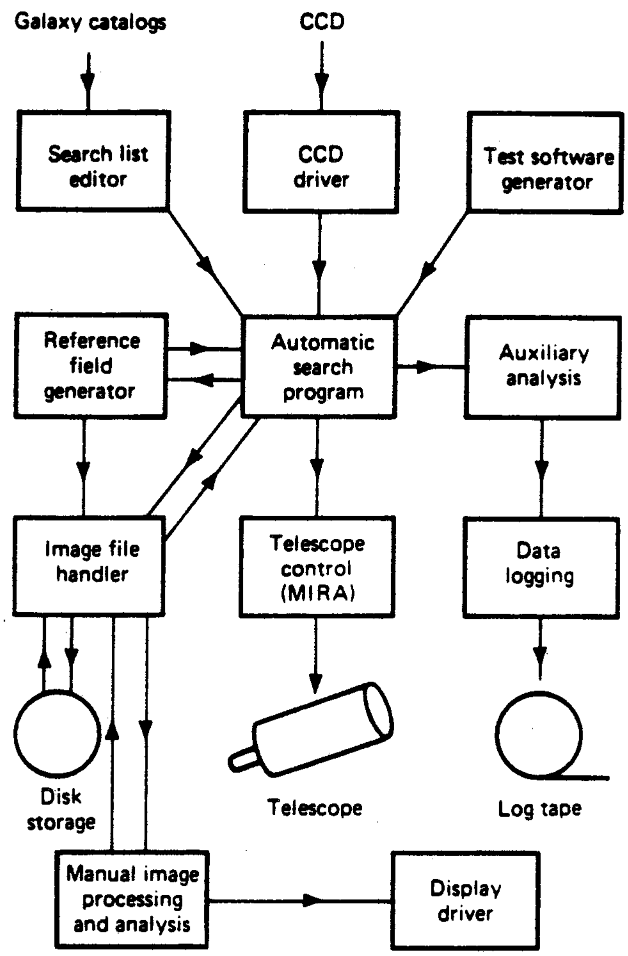

Figure 6 — Supernova Search System Software

Acknowledgements

We would like to thank Stirling Colgate for advice and strong encouragement. Roger Williams, our senior programmer, has been invaluable in developing our software, and Robert G. Smits and George Gibson of the LBL engineering staff are responsible for much of our progress in hardware. Undergraduates John Culver, David Smith, and Geordie Zapalac have contributed substantially to the project. John Geary at Harvard has been very helpful in getting our CCD operational. Luis Alvarez stimulated our original interest in this problem, and has been an enthusiastic supporter of our work. Finally, we would like to acknowledge the creative participation of Bruce Weaver and the other astronomers of MIRA in this project.

This work was supported in part by contracts and grants from the Department of Energy (contract W7405ENG48), the California Space Institute, the DuPont Corporation, the Research Corporation, the National Aeronautics and Space Administration, and the Alan T. Waterman Award of the National Science Foundation.

References

1. W. L. W. Sargent, L. Searle, and C. T. Cowal, in Supernovae and Supernova Remnants, ed. C. B. Cosmovici, p. 33 (Reidel, Dordrecht 1974).

2. L. Rosino, in Supernovae, ed. D. Schramm, p. 1 (Reidel, Dordrecht 1977).

3. J. Maza, in Proceedings of the Texas Workshop on Type I Supernovae, ed. J. C. Wheeler, p. 7 (U. Texas, Austin 1980).

4. J. R. Dunlap, J. A. Hynek, and W. T. Powers, in Advances in Electronics and Electron Physics, eds. J. B. McGee, D. McMullan, and E. Kahn, 33B, p. 789 (1972).

5. E. Kibblewhite, in Proceedings of the 1981 NATO Advanced Studies Institute on Supernovae, (Reidel, Dordrecht 1981), in preparation.

6. J. T. McGraw, J. R. P. Angel, and T. A. Sargent,in Applications of Digital Image Processing to Astronomy, ed. D. A. Elliott, p. 20 (SPIE, Washington 1980).

7. Information on this project is available from Stirling Colgate at the Los Alamos Scientific Laboratory. The automated telescope is described in:
S. A. Colgate, E. P. Moore, and R. Carlson, Publications of the Astronomical Society of the Pacific, 87, pp. 565-575 (1975).

8. H. Gursky, J. Geary, R. Schild, T. Stephenson, and T. Weeks, in Applications of Digital Image Processing to Astronomy, ed. D. A. Elliott, p. 14 (SPIE, Washington 1980).