Title
AN AUTOMATED SEARCH FOR SUPERNOVA EXPLOSIONS

Author
Kare, J.T.

Publication Date
1987-07-01

LBL-23782 Preprint

Lawrence Berkeley Laboratory

UNIVERSITY OF CALIFORNIA

Physics Division

Submitted to Review of Scientific Instruments

An Automated Search for Supernova Explosions

J.T. Kare, M.S. Burns, F.S. Crawford, P.G. Friedman,
R.A. Muller, C.R. Pennypacker, S. Perlmutter,
R. Treffers, and R. Williams

July 1987

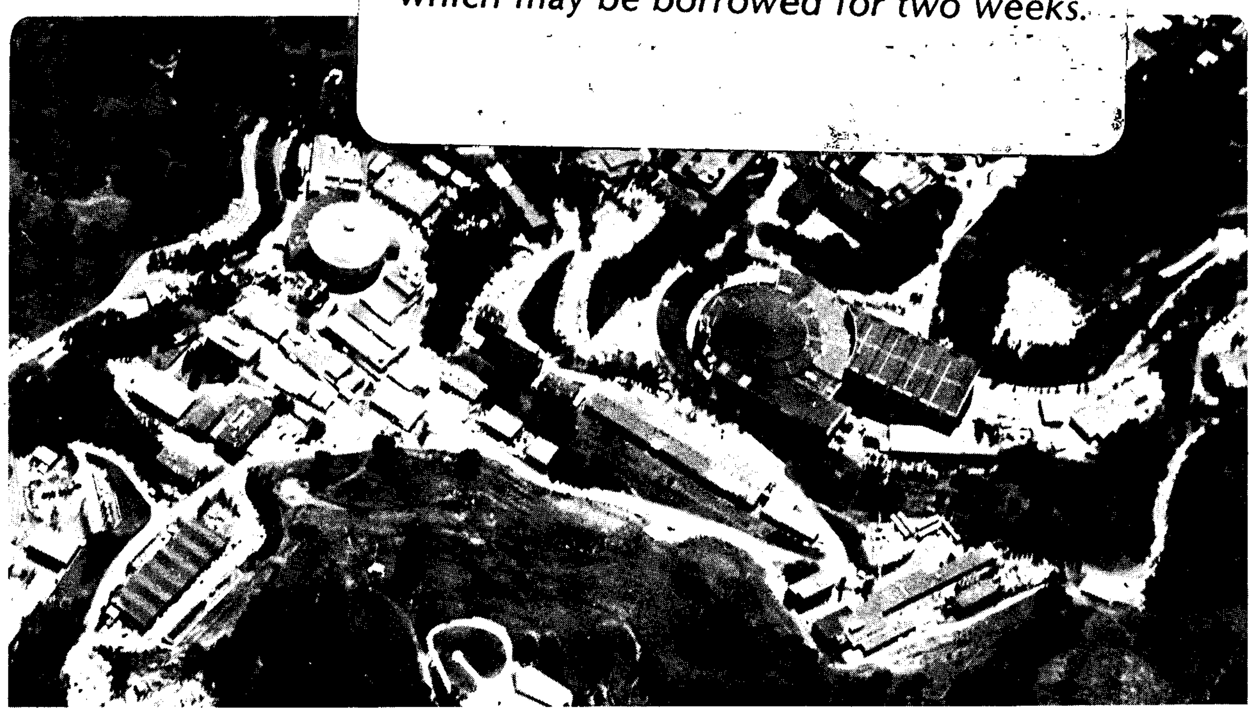

Prepared for the U.S. Department of Energy under Contract DE-AC03-76SF00098

DISCLAIMER

An Automated Search for Supernova Explosions

Jordin T. Kare*, M. Shane Burns**, Frank S. Crawford, Peter G. Friedman***, Richard A. Muller, Carlton R. Pennypacker, Saul Perlmutter, Richard Treffers, and Roger Williams

Lawrence Berkeley Laboratory
University of California
Berkeley, California 94720
USA

Abstract

We describes the design and construction of a new search for supernovae, based on the use of current computer and detector technology. This search uses a computer-controlled telescope and charge coupled device (CCD) detector to collect images of hundreds of galaxies per night of observation, and a dedicated minicomputer to process these images in real time. The search system has a detection threshold of 17^{th} visual magnitude and discovered its first supernova on 18 May 1986, and at this writing has found two additional supernovae, one at $m_V = 15.5$. It currently is being upgraded with an improved detector and a more powerful computer. At full speed and sensitivity, the system will be able to search 2500 galaxies per night with a limiting magnitude of 18 or fainter.

*current address: Lawrence Livermore National Laboratory Livermore, L278, Livermore, CA, 94550.

**current address: Dept. of Physics,University of Wyoming, Laramie, WY, 82071.

***current address: ATT Bell Laboratories, 6F002, 600 Mountain Ave.,Murray Hill, N.J., 07974.

Introduction

Supernovae are among the most important, as well as most spectacular, of astrophysical events. However, because of their transient nature, supernovae are difficult to observe. The goals of the Berkeley automated supernova search are 1) to detect a few nearby supernovae ($D \approx 20$ Mpc or less, $H_0 = 50$) per year at very early times ($\approx 1\%$ of maximum brightness for type I supernovae) for detailed study at all wavelenths, and 2) to detect a large number of supernovae at intermediate distances ($D \approx 20$-100 Mpc) and moderately early times (10 - 50% of maximum brightness) for photometric and statistical analyses. The Berkeley search is superior in many respects to manual photographic searches, and complementary to other automated searches [1,2] which are most efficient at finding distant ($D > 100$ Mpc) supernovae near maximum light.

The Berkeley system, shown in figure 1, uses a computer controlled, fast-slewing telescope and charge coupled device (CCD) camera to collect images of galaxies; these images are compared to stored reference images by a dedicated minicomputer system which identifies and confirms supernova candidates. Figure 2 shows our real-time system, described in later sections, that is soon to be operational. In this system we update ideas proposed by Stirling Colgate in the early 1970's [3], and employed in his automated search effort [4]. Major differences between Colgate's system and ours include the use of a CCD detector in place of an intensified image tube, and the use of a relatively small dedicated computer system at the telescope site for image processing, with a separate microcomputer for telescope control.

At present, the system computer directs the telescope computer to collect images of a preset list of galaxies; the CCD exposure is 30 seconds per galaxy, and several hundred galaxies can be imaged each night. The images are recorded in digital form on videotape, and fed into the system computer the following day for processing. The computer subtracts a reference image (taken with the same CCD and a 120 second exposure) from the current image to generate a difference image. Currently this difference image is stored on disk, and scanned for supernova candidates by a human operator using a graphics display terminal.

In the near future, the computer itself will scan the difference image for likely supernova candidates, and will automatically acquire and process additional images to confirm possible supernovae. By moving the computer to the telescope site and reading images directly into its memory, we will be able to process images and confirm supernovae in real time. With a better

CCD, improved optics, and upgraded hardware and software, we expect to be able to collect and process an image in as little as 8 seconds, allowing us to scan 2500 galaxies per night. Finally, by moving to a better telescope and observing site, we can improve our detection threshold to better than 18th magnitude, and detect supernovae earlier and more reliably.

Hardware

Charge Coupled Device (CCD)

The heart of the supernova search system is the cooled CCD camera which collects the images. The uses and advantages of CCD detectors in astronomy have been widely discussed [5,6]. The primary advantages of CCDs for the supernova search are high quantum efficiency, good linearity, wide dynamic range, and mechanical and electrical stability. We currently use an RCA device (SID53612-X0); a thinned, back illuminated CCD with 320 x 512 pixels. Each pixel is 30 microns square. This is a three-phase buried channel device. A description of the properties of this type of CCD is given by Fowler et al. [7] along with results of laboratory tests of one device. A second identical CCD camera using an RCA SID501 CCD was built to take advantage of the lower noise and higher efficiency of the SID501. Unfortunately, our device is unusable due to a high nonuniform background level. The background is highest near the on chip amplifier, and is roughly proportional to the image integration time; we believe it is due to electroluminescence in the output amplifier, as seen by other experimenters [8].

CCD Housing and Electronics

Our CCD driver electronics are based on designs by J. Geary of the Harvard-Smithsonian Center for Astrophysics [9]. The output of the on-chip FET amplifier is boosted by a preamplifier in the CCD housing and transmitted via shielded twisted pair to a correlated double sampling (CDS) amplifier and a 14 bit serial output A/D converter. Our current readout rate is 1 pixel every 50 microseconds, or one full image in 8 seconds. The output of the A/D converter is transmitted to the computer or videotape data recorder via a Hewlett Packard fiber optic data link using a simple pulse-width modulation scheme. This link is very cheap (< $50) but the fiber lenght is limited to ≈ 20 meters by the high attenuation of the plastic-core fiber. It isolates the CCD from computer electrical noise and also minimizes the risk of damage due to inductive effects of lightning strikes. The CCD is cooled to -110 C with liquid nitrogen, to reduce dark current essentially to zero. A feedback circuit driving resistive heaters maintains a constant temperature at the CCD. A vacuum housing insulates the CCD from the environment. Using a 2-stage mechanical pump, the vacuum system reaches pressures of 18 - 20 millitorr, with a hold time of several days. Additional pumping comes from a cryopump: a stainless steel tube containing about 3 cm^3 of Linde Molecular Sieve (mixed types 4A and 5A), which is immersed in the liquid nitrogen reservoir. Our liquid nitrogen dewar is a large consumer thermos bottle ("pump-pot"), with a capacity of 1 liter and a hold time of approximately 15 hours. We flow 0.5 cu. ft. per hour of dry nitrogen through the area between the window and the shutter mechanism to prevent condensation on the quartz window of the CCD housing. Further details of both the CCD and our particular driver circuits are available in the senior honors thesis of J. Culver [10].

Measured CCD Properties

The detailed properties of our current CCD are discussed in the thesis of M. Shane Burns [11]. The noise level and gain of our device were measured by the standard technique [7] of measuring the variance of the output signal at different levels of illumination. The noise is 86 ± 15 photoelectrons/pixel. The effective gain of the CCD output amplifier and correlated double sampling circuits is such that one count of our analog to digital (A/D) converter corresponds to 25 ± 5 photoelectrons detected. Transfer inefficiency in the horizontal direction is $\varepsilon = (2 \pm 1) \times 10^{-4}$, and in the vertical direction $\varepsilon = (3 \pm 1) \times 10^{-4}$, at -100 C, where ε is the fraction of stored charge left behind after each transfer. Transfer inefficiency in the horizontal direction increases to 5×10^{-4} at -125 C. Our normal operating temperature is between -120 C and -110 C.

The transfer inefficiency of our CCD is significantly worse at low light levels. We have made detailed measurements of this effect, and find that the inefficiency varies nonlinearly with the pixel value, and is essentially constant only above approximately 250 A/D counts (≈ 6000 photoelectrons). This has been only a minor problem for the supernova search, as we currently collect about 100 A/D counts worth of background in our 30 second standard exposures, and are not much affected by small transfer losses. However, dim star images are significantly distorted, and astrometric use of our CCD (as for the solar companion search discussed later in this paper), or shorter supernova search exposures, may require "flashing" the CCD with a small light source to improve transfer efficiency.

The sensitivity of our CCD to external noise pickup is high, and is critically dependent on the adjustment of the positive voltage pulse applied to the gate of the reset transistor which clears the CCD output gate between pixels (ϕ-reset). There is significant (5-10 A/D counts) 120-Hz pickup if this voltage is adjusted as little as 0.1 volt away from its nominal level of +0.9 volts. The manufacturer's claimed quantum efficiency for our CCD is 70 - 80% from 400 to 700 nm. We have measured the quantum efficiency at one wavelength (633 nm), using a helium-neon laser source and a calibrated silicon photodiode as a reference. The efficiency is 31 ± 8 per cent, or a factor of 2.5 lower than claimed. We have not measured the spectral response directly. At a wavelength of 500 nm, the overall system sensitivity is a further factor of 2 below the value we calculate from known losses, including the 30% quantum efficiency. Some of this difference may be due to even lower CCD quantum efficiency at this wavelength.

The standard deviation of individual pixel quantum efficiencies is approximately 1%. Because we do not require very accurate quantum efficiency correction, we have not attempted to measure or correct for the color dependence of quantum efficiency variations.

New CCD

Because of the relatively poor performance of our current CCD, especially its high noise and low quantum efficiency, we have obtained a new CCD, and are currently constructing a camera assemby around it. The new CCD, an 800x800 pixel format Texas Instrument device (model #VPCCD) has a readout noise of about 10 electrons (rms) per pixel. The blue response of this new CCD is higher than the RCA device.

Telescope

The telescope in use at this writing is a 30" (75 cm) f/8 Ritchey-Chretien reflector, located at the Leuschner Observatory of the University of California in Lafayette, California. The Leuschner telescope has been equipped with a fully automatic control system over the last four years, largely through the efforts of Richard Treffers [12]. The telescope position is sensed by Inductosyn absolute angle encoders with a least count accuracy of 4 arc sec. The encoders are coupled via an IEEE-488 bus interface to the control computer. Currently this is an IBM PC with 256K memory and dual floppy disk drives; however, the telescope has operated automatically with a 24K Commodore PET computer with one disk drive. The IBM computer is programmed in C. The control computer sends pointing commands to a closed-loop digital servo system constructed by R. Smits and G. Gibson of LBL. This system uses torque motors coupled to optical shaft encoders to slew, guide, and track the telescope. The smallest step of the shaft encoders corresponds to approximately 1/15of an arc second of telescope motion. The slewing speed of the telescope is limited by the counting rate of the shaft encoders to roughly 100,000 steps, or just under 2 degrees, per second. Faster slewing is possible with different encoders, and ramping of the telescope speed to minimize vibration is done automatically. The telescope can move between galaxies (typically 2 degrees apart) and settle in 3 seconds.

The telescope can be pointed with a short term, relative pointing accuracy of about 10 arc seconds. Absolute pointing accuracy and night to night repeatability are approximately 40 arc seconds, largely due to a periodic error in the angle encoders. We could add further pointing corrections to reduce this, but do not plan do so soon.

The Leuschner control computer communicates with the main system minicomputer via an RS-232 serial port and dial up modem at 300 baud. A 1200 baud modem and telephone line provide an operators terminal for the PDP-11 at the telescope site. In addition to running the telescope, the control computer operates the CCD shutter, controls the CCD readout, and produces both a printed log and logging information for the videotape recorder audio track. The control computer can (and currently does) operate independently of the system computer, with coordinates selected either manually or from a stored "search list".

Another telescope, used for some testing and which has performance like others to be used for the full speed search, is the 36" f/10 Cassegrain reflector belonging to the Monterey Institute for Research in Astronomy (MIRA) and located on Chew's Ridge near Monterey, California. The MIRA telescope provides larger collecting area and considerably better observing conditions, but is less conveniently accessible for development work and until recently was located at a very poor observing site.

Optics

The angle of sky subtended by one CCD pixel is the "pixel scale". The optimum pixel scale is roughly equal to the seeing disk; larger pixels increase the sky background noise in each pixel, while smaller pixels result in the light of a supernova being spread over many pixels (and also increase the computer processing time used). An additional factor is that smaller pixels allow more accurate subtraction of images, and higher precision in astrometry; the latter is important in the parallax search for a solar companion star ("Nemesis") [13] which is also being conducted with the supernova search hardware.

Our current optics are a compromise between the requirements of the supernova search, the Nemesis search, and available standard lenses. A field lens directs the light from the telescope focal plane into the exit pupil of a Bronica 6 cm format camera lens. The field lens is placed slightly out of the telescope focal plane so that surface dust and imperfections are not in focus on the CCD, and also to partially correct the focal plane curvature produced by the rest of the system. The Bronica lens, with a 200 mm focal length, is run "backwards" to transform light from the telescope focal plane to (nearly) parallel light; a Nikkor 50 mm focal length lens focuses this light onto the CCD. A computer controlled eight position filter wheel will allow us to insert filters automatically into the parallel light region between the Bronica and Nikkor lenses. This system gives very good image quality and focal plane flatness, with negligible pincushion distortion, at a scale of 2.4 arc seconds per pixel. The major optical problems for the initial embodiment of the supernova search are severe vignetting at the edges of the CCD, and relatively low total transmission even on axis. The next generation system with the Texas Instrument CCD will use no reducing optics.

Computer System

The main computer system for the supernova search is a Digital Equipment Corporation (DEC) PDP-11/44 with 454 + 160 Mbytes of Fujitsu winchester disk storage, floating point processor hardware and (currently) 2.5 Mbytes of main memory. This system is located at the Lawrence Berkeley Laboratory. In view of the problems we have encountered in using a 16 bit computer, we have purchased a DEC MicroVAX II, and will use this full 32 bit machine for most future system development. The MicroVAX is installed at the Leuschner Observatory telescope site. Graphic output for the 11/44 is provided by a Grinnel GMR-270 (512 x 512 pixels by 8 bits/pixel gray scale) and Tektronix dry silver process graphics printer.

At the telescope and the laboratory, we use IBM AT computers to store and manipulate CCD images. We use a Tecmar Graphics Master Board (200037, rev.E1), to hold 4 bits (16 gray levels) of the entire CCD image.

CCD data, either direct from the CCD or recovered from videotape, are read into the PDP-11/44 via a direct memory access (DMA) interface (MDB DR11-B). A separate card contains a serial to parallel converter and a 128 word by 16 bit first in first out (FIFO) buffer, which accumulates CCD data when the computer bus is not available for DMA.

Data for the MicroVAX are buffered through an IBM PC, which provides a simple hardware and software interface, and communicates with the Microvax over an Ethernet link.

Videotape Storage Unit

Since the main supernova search computer is not at the telescope site (and until recently, no large computer was available there), we store each night's images for later processing. The CCD data rate is considerably too high (40 Kbytes/sec for 8 seconds) for transmission via telephone line, and no inexpensive alternative (such as line-of-sight microwave transmission) is available. Even with the MicroVAX available for real time processing, it is desirable to store the raw CCD data for archival purposes. We have therefore built a digital interface to a consumer videotape recorder, which provides extremely compact and inexpensive storage of very large amounts of data: less than 2 cents per Mbyte stored, and $500 per \&"drive". This interface stores two 14 or 16 bit words per TV horizontal line (60 microsec.), 228 lines per field, for a data rate of 27,360 words per second and a tape capacity of 591 Mbytes per 3 hour tape. We use Beta format recorders, slightly modified to disable the internal "drop out compensation" circuit (which would otherwise replace data lost due to tape dropouts with incorrect data). Beta recorders provide faster switching between modes (play, fast forward, and rewind) than VHS recorders, and have slightly higher video quality, but VHS format could be used. The recorder system uses simple redundancy and error checking to achieve pixel error rates of a few times 10^{-5} -- sufficient for the supernova search, although not for most computer applications. A 2-bit count, recorded with each pixel, allows the recorder interface to detect dropouts up to 3 pixels long and fill in the image with dummy (zero) pixels; this prevents "jogs" in the recovered image.

We store logging information on the audio track of the video tape with a standard 300 baud modem; the telescope control computer writes the time, date, telescope pointing data, and exposure time for each picture. The PDP-11/44 can read this information via a serial port, and operates the videotape recorder via a simple interface which allows the computer to "push buttons" on an infrared remote control by sending characters out the same port. A simple video frame store (scan converter) built into the videotape interface has been used to display the CCD image at the telescope site. Although the scan converter can store and display only a 128 x 128 pixel region of the CCD at a time, it is sufficient for aligning the telescope and monitoring the operation of the CCD. A videotape recorder system similar to ours, also designed for CCD data storage, has been described in the literature [14]. Further data on our system is available in reference [15].

Software

Computer System Software and Utilities

Operating System and Languages

The PDP-11/44 runs the RSX-11M operating system, version 3.2. Individual programs (tasks) in RSX-11M are limited to a 16 bit address space; we work around this restriction mainly by sequentially spawning multiple tasks, each performing a single function. Overlays, which allow different program subsections to share the same task address space, are difficult to construct efficiently for FORTRAN programs, and have been found to lead to subtle run-time errors. Images are stored in named data regions, which are accessed through a series of utility routines we have written. These routines directly access the PDP-11's memory mapping registers, which allow access to the full 22 bit (4 Mbyte) address space of the processor. The MicroVAX computer runs MicroVMS.

The bulk of the supernova search software is written in DEC FORTRAN IV-Plus, with a small amount of MACRO-11 (PDP-11 assembler) code used for device drivers. We considered several other languages and operating systems, including UNIX/C, PASCAL, and FORTH, for our purposes. UNIX does not provide appropriate real time facilities, and available C and PASCAL packages were not guaranteed to use the facilities of RSX-11 properly. None of the members of the supernova search project were sufficiently familiar with FORTH to consider using it for a large project. The total supernova search software base now contains roughly 200,000 lines of FORTRAN code, and represents some 10 man-years of programming effort.

Utilities

We have written our own routines for a variety of low level tasks, including reading in the CCD, controlling the videotape recorder, and driving the Grinnell graphics display. The region access routines mentioned above allow us to manipulate several (on occasion as many as 10) images and subimages in main memory, with a minimum of disk access. We have also accumulated a collection of basic image processing functions, such as tasks for subtracting two images or doing a convolution, which can be called either by a program or directly from a keyboard. Images can be stored on disk, either directly (2 bytes per pixel) or in a compressed form using delta modulation (storing the difference between successive pixels) with an effective storage density of about 1.1 bytes per pixel.

We considered using any of several existing image processing software packages available when we began our work, but chose not to even though we have undoubtedly reinvented several wheels. We expected to need a relatively limited set of functions heavily optimized for speed, rather than interactivity; this has not proven to be entirely the case. One of the chief advantages of the MicroVAX environment has been the availability of VISTA [16], a convenient image processing package developed by Lick Observatory, which provides many of the utility routines we need, and which has greatly accelerated the construction of our second generation of software. As the number of galaxies processed by the supernova search has grown from a few to several hundred per night, we have discovered a vital need for data base management software to keep track of the status of each galaxy. Lacking actual database software, we have written several file-sorting and updating programs in FORTRAN, but we have not yet created a fully satisfactory system for keeping track of (at the moment) 2000 galaxies and some tens of thousands of exposures. Accurate exposure records will be particularly important for statistical calculations of, e.g., the supernova rate.

Automated Search Software

Introduction

The two major programs which make up the the supernova search software are the Reference Builder and the Image Cycle. The Reference Builder automatically generates disk files of reference information from galaxy images stored on videotape. These files include the actual reference picture of the galaxy, plus lists of bright objects in the CCD field and various log files. The Image Cycle is the set of programs which collects search images, compares them to the appropriate reference images, and identifies, logs, and confirms possible supernovae. The major subsections of the Image Cycle are the Control program, the Image Subtractor, and the Supernova Finder. The Control program reads in a CCD image -- currently from videotape, eventually directly from the CCD -- and supplies the image and necessary additional information to the Image Subtractor. The Subtractor processes the image to match the reference picture as closely as possible, and subtracts the reference image from the current image to produce a difference image. These programs are currently operating.

In the future, the difference image will be passed to the Supernova Finder, which will distinguish supernova candidates from false alarms and request confirming exposures from the Control program. Currently the functions of the Supernova Finder are done visually, by a human inspecting the difference image on a graphics display. Although including a person in the loop may lead the reader to wonder why we bother with any image processing at all, it is in fact much easier (and faster and less fatigueing) to identify supernovae on the difference image than on the un-subtracted current image. Various preliminary versions of the Supernova Finder are now run in parallel with ("shadowing") the visual scanning, so that we can estimate and improve their performance. We find that the human eye, assisted by the computer, is still better than either the eye or the computer alone.

Reference Builder

The reference builder operates on reference CCD exposures to produce a series of reference files for each galaxy: the reference subimage array (.RSA or, in compressed format, .RSC) file, the object catalog (.OBC), and the reference subimage parameters (.RSP) file. A chronological log (.CLG) file and an observation log (.OLG) file are also created by the reference builder and updated by the subtraction system. The name of each file consists of a 9 character "galaxy ID" code giving the approximate galaxy coordinates plus a sequence number, followed by the appropriate three character extension. This approach creates many small files, with some waste of storage space, but allows all the work of locating data for a particular galaxy to be given to the operating system. In operation, the reference builder starts the videotape recorder and reads the audio track until it finds a valid galaxy ID, then reads the following image into memory. It locates fiducial stars by convolving this reference image with a "peak filter" (a horizontal row of pixels with coefficients -.25, -.25, 0, 1, 0, -.25, -.25) and grouping adjacent bright pixels together to generate a list of objects. Bad CCD pixels, and the edges of the CCD, are masked out.

The reference builder then processes a 16 x 16 pixel subimage around each object to find its properties (brightness, centroid, and width). We have found experimentally that, given the vignetting of our optics and the substantial variations in quantum efficiency over the CCD, we must correct (flatten) the field locally around each star to obtain consistent brightnesses. Each pixel is corrected by subtracting the value of the corresponding pixel in a dark field, and dividing by a pixel quantum efficiency determined from a "quantum efficiency" field (an exposure of twilight sky, minus a dark field) which also compensates for vignetting in the optical system. We have found that in practice this process leaves smooth variations of up to 5 percent of the raw (uncorrected) background level in the corrected image. We therefore measure the local background by averaging pixels from the corners of the star subimage after flattening and before finding the brightness and centroid.

The properties of the 10 brightest objects found which do not saturate the CCD are stored in the .OBC file. This file also has space to list other objects (such as variable stars) that lie anywhere in the image, and that we might want to monitor the brightness of; currently we do not keep track of any such "special objects".

The reference builder can identify the galaxy by location (usually within a few pixels of the center of the CCD) and width; although the width calculation is crude, it is (usually) sufficient to distinguish a galaxy from a star. In practice, our standard subimage is large enough (150 x 150 pixels) that the reference builder simply sets the boundaries of the subimage to fixed coordinates, and the galaxy is almost always found within them. The .RSP file is created, containing the pixel coordinates for the reference subimage, plus other parameters (such as galaxy core width and brightness) that we do not currently use. The actual subimage is stored in the .RSA file. The total disk space used for one galaxy is about 50,000 bytes; most of this is in the .RSA file. We are slowly converting existing .RSA files to compressed format .RSC files to save both disk space and access time. The .CLG chronological log file for each galaxy contains a one-line entry for each operation on that galaxy, including reference building and subtraction. This provides an "audit trail" for the observations on that galaxy. The .OLG file contains a detailed record of the most recent observations, for use by the Supernova Finder; we do not yet use it.

Image Cycle

Control Program

In the full supernova search, the Control Program will operate in real time, sending exposure requests to the telescope control computer and reading in the resulting CCD images. It will select galaxies to be observed from a master list of galaxies, and will be able to accept requests from the Supernova Finder for special exposures to confirm supernova candidates and identify false alarms. In the current system, selection of galaxies is done by the system operator, using a simple program to make cuts (in, e.g., RA, Dec, and brightness) on our master list of galaxies. The resulting search list(s) are written on IBM PC-compatible diskettes for use by the telescope control computer; they can also be transmitted to the telescope control computer by modem during each night's observing. During image analysis, the following day, the Control Program simply reads a galaxy ID and an image from videotape and starts the Image Subtractor directly.

Image Subtractor

The Image Subtractor, once given a current image of a galaxy, first retrieves the list of fiducial stars from the file (identifier).OBC, then searches a box (currently 50 x 50 pixels) around the reference location of the brightest fiducial, and calculates an approximate offset assuming the brightest pixel found to be the desired star. Given this rough offset, it locates the brightest pixel near the expected location of each fiducial star, and processes a 16 x 16 pixel subimage around each such bright pixel, as in the Reference builder, to get a centroid, brightness, and width for each presumed star. A mean offset and mean brightness ratio relative to the reference list of fiducials is calculated, with missing or poorly aligned stars discarded. If too few "good" stars are found, the subtractor starts over with a 50 x 50 search for the next-brightest fiducial; if no match can be found for any of the fiducials the image is flagged as bad and skipped.

We do not currently correct for optical distortions, magnification, or rotations of the image; ideally these produce errors over the galaxy subimage of much less than a pixel. Unfortunately, we have discovered that our CCD camera mounting allowed the camera to rotate slightly ($\approx$ 2 degrees) relative to the telescope, so our reference images are sometimes rotated significantly relative to the current images. We now align the CCD to the telescope axes (by "streaking" star images across the CCD) whenever the camera has been removed from the telescope, and check the alignment nightly; we will eventually replace the misaligned reference images. Meanwhile, we include a simple fit for rotation in the offset calculation, so that the offset will be correct at the center of the CCD.

Once the current offset is found, the software locates the portion of the CCD containing the galaxy and copies this portion to a separate region in memory; all further processing is done on this subimage. The Subtractor flattens the subimage, finds and subtracts a local sky background level, and multiplies the subimage by a constant to correct for atmospheric transmission and exposure variations (as determined from the fiducial star brightnesses).

The reference subimage is retrieved from disk and "jogged" (shifted by a fraction of a pixel) by convolution with a 2 x 2 pixel matrix. The jogging program can also blur the reference subimage by convolving it with a gaussian to correct for poorer seeing in the current image, but in practice we do not measure the seeing accurately, so we set the blurring to zero; this also reduces the processing time. Once the current and reference subimages are matched, the reference subimage is subtracted from the current image, and the resulting difference subimage stored in a separate region in memory for processing by the Supernova Finder. The current software copies the subimage to disk for later display.

Supernova Finder

The Supernova Finder program is the only major software section that is planned but not yet operating. The purpose of the Supernova Finder is to distinguish supernova candidates from various kinds of false alarms, and to alert the system operator when a (probable) supernova is found. When the supernova search was planned, we expected to encounter several types of false alarms: statistical fluctuations in the CCD output, cosmic rays passing through the CCD, asteroids and variable stars in the field of view, and subtraction errors caused by CCD misalignment and the finite CCD pixel size.

The simplest automatic supernova Finder simply checks for bright pixels that exceed a threshold, nominally five standard deviations above the local background, and requests a second confirming exposure if any are found. False alarms due to rare events -- cosmic rays, statistical fluctuations, etc. -- which do not appear in a second image are trivially eliminated by the confirmation process. Asteroids are eliminated by delaying the confirming exposure by at least a few minutes; the asteroid image will move several pixels. Based on our experience, it will also be possible to discard many false alarms with simple algorithms; for example, most cosmic ray events stand out because they occur in one or two pixels, while typical starlike objects near our sensitivity threshold occupy at least nine pixels significantly above background; a cut based on the width of an object will reject most cosmic rays. The major problem with a simple Finder, however, is that it cannot recognize false alarms caused by poor subtraction of the galaxy core and foreground stars. This problem is exacerbated by the relatively crude offset calculations and fractional-pixel jogging algorithm used by the Image Subtractor, and by some currently-uncorrected sources of misalignment (specifically, rotations introduced in some of our reference images by mismounting of the CCD camera on the telescope).

The human eye is very good at recognizing the pattern of errors produced by a faulty subtraction: foreground stars appear as adjacent bright and dark spots (supernovae paired with "antisupernovae") or as bright spots surrounded by dark rings. It is difficult to make a simple program that will recognize these patterns, but several options exist. A program can check around each bright pixel for a nearby dark pixel (antisupernova) and rule the candidate out if one is found. This increases the signal to noise ratio required for detection, however (since the program is essentially adding together the signals from two or more pixels, only one of which contains the supernova). This check also becomes difficult for large misalignments or bad seeing, when the dark pixel may be many pixels away from the bright one, and it may fail if the dark area is spread over many pixels. The supernova-antisupernova test may, however, be adequate as our subtraction techniques improve.

A more powerful technique is to identify in advance the pixels likely to produce false alarms, and present the Supernova Finder with a list of pixels to ignore (or treat with care). This can be done either by processing the reference image through a gradient-finding program which locates and lists foreground stars, or (more generally) by processing several test images (assumed to contain no supernovae) and listing all the pixels where false alarms appear. Again, this fails if the misalignment of images is too large, producing too many false alarms, but for misalignments of one pixel or less (already achievable in the absence of rotations) it should be very reliable. Both of these techniques assume that there will be some subtraction errors, and neither requires "Nyquist sampling" (pixels smaller than 1/2 the seeing disk) to eliminate false alarms.

When the Supernova Finder is operating, it will normally check that a supernova candidate appears in at least three CCD images, one delayed long enough (typically 1 hour) to rule out asteroids, before alerting the operator. With three observations, the last two with extended exposures, we can be confident that the observed signal is real. The major remaining source of false alarms will be variable foreground stars, including ordinary novae, which have long or irregular periods. We cannot eliminate all such variables, but we do not expect to see very many that cannot be distinguished from supernovae. We have not had any cases where human observers could not distinguish a variable star from a supernova, although as we lower our detection threshold and increase our sensitivity, we undoubtedly will. For occasional events, prompt spectroscopy will distinguish supernovae from variables. When we are finding enough events in distant galaxies to make prompt spectroscopy of all of them impractical (and redundant), we can distinguish supernovae by tracing light curves for a short time before reporting detections. For interesting candidates in nearby galaxies, though, we will request spectra from other observers as soon as our own prompt tests (including, e.g., color difference measurements) say "supernova!"

System Performance

Search Sensitivity

The ultimate limit on the sensitivity of the supernova search is set by the signal to noise ratio. If we consider only a single pixel (i.e. no summing of adjacent pixels), the signal to noise ratio is

$$S/N = \{n_{sn}\, f\, t_{exp}\} / \{ ([n_{sky} + n_{gal}]\, a_{pixel}\, t_{exp} + N_{CCD}^2)^{1/2} \}$$

where

n_{sn}	=	photoelectrons/sec from supernova
f	=	fraction of supernova light in 1 pixel
t_{exp}	=	Exposure time
n_{sky}	=	photoelectrons/sec/arc sec^2 from sky

n_{gal} = photoelectrons/sec/arc sec^2 from galaxy
a_{pixel} = area of 1 pixel (arcsec^2)
N_{CCD} = CCD noise, rms, photoelectrons per pixel

The value of n_{sky} is fixed by the sky brightness and f/number of the optical system; it depends only indirectly on the telescope aperture. Our measured value for n_{sky} with our current optics (f/3.5) is 70 photoelectrons per arc sec^2 per second at Leuschner, which has relatively bright sky (estimated as equivalent to one magnitude 17.5 star per arc sec^2); the sky brightness at MIRA is approximately 5 times lower. n_{gal} is typically much lower than this outside of the galaxy core (5-10 arc sec radius), and we neglect it here. n_{gal} will raise the detection threshold in galaxy cores when accurate subtraction routines are used.

N_{CCD} is 86 ± 15 for our current CCD. Our current higher performance CCDs has noise of ~10 photoelectrons/pixel. The random variable f depends on the seeing, and on the position of the supernova image with respect to CCD pixel boundaries. In the limit of perfect seeing, f is 1.0. For typical seeing at the Leuschner site and 2.4 arc second square pixels, we find a mean value for foreground stars (brightest pixel/total brightness) of 0.18. Figure 3 is a plot of f for 17 stars. Note that we can actually improve our sensitivity slightly by checking blocks of 4 or 9 pixels -- f(4 pixels) is typically 0.5, while the noise level is twice that for one pixel, so the effective value of f is 0.25.

Current and Future Threshold for Detection

The conversion from the number of detected photo-electrons (n_{sn}) to limiting magnitude depends on the spectrum of the supernova, the transmission of the entire optical system including the atmosphere, and the quantum efficiency of the CCD. The conversion factor we use is

$$n_{sn} = 6.5 \times 10^{(20-m_v)/2.5)}$$

This value is based on observations of standard stars of known visual magnitude using our current CCD at the Leuschner telescope (figure 4). Figure 5 assumes the increased collecting area and lower sky background available at the MIRA or other 36 inch telescope at a good site, and improved reducing optics with 75% transmission rather than our current 50%. Improved seeing and better pixel scale are assumed to increase f to 0.5. Note that for short exposures, CCD noise dominates over sky noise, so that one gains sensitivity linearly with increasing exposure. At longer exposure times, the noise level rises as the square root of the exposure time, so the effective sensitivity rises more slowly.

Search Speed

Telescope Performance

The Leuschner telescope currently requires approximately 10 seconds to point over a typical inter-galaxy distance of two degrees. This includes sufficient time for the telescope to settle. With an 8 second delay for readout (we currently do not read out during moves, due to software timing problems rather than to noise pickup), the total time spent per galaxy is approximately

$$t = t_{move} + t_{readout} + t_{exp} = 18 \text{ sec} + t_{exp}$$

The MIRA telescope can move between galaxies in under 3 seconds. We have tested this by trailing a star image across the CCD using both Right Ascention and Declination stepping motors, then stopping one motor abruptly. No overshoot or ringing is observable on a time scale of a few hundred milliseconds. We expect to achieve comparable performance at the Leuschner telescope by fine-tuning the upgraded drive system. We have tried reading out the CCD in times as short as 3 seconds, and while the telescope drive was operating, and have not found any increase in readout noise. By reading out directly into the system computer, or by buffering the CCD data in a microcomputer memory, we can eliminate constraints due to the videotape recorder. We should therefore be able to read out the CCD while the telescope is moving, giving

$$t = t_{move} + t_{exp} = 3 \text{ sec} + t_{exp}$$

Figure 6 plots the number of galaxies observable per (6-hour) night versus t_{exp} for these two cases.

Processing Speed

The total computer time required for one image cycle is currently about 60 seconds for a 150 x 150 pixel galaxy subimage. Nearly half of this time is spent in the jogging subroutine. By revising the algorithm used in this routine, and by simple optimization of other routines (for example, minimizing disk accesses by buffering reference files in memory) we expect to reduce this time to under 30 seconds, which is sufficiently fast to keep up with our current CCD. However, achieving a mean processing time per galaxy of under 8 seconds (including CCD readout) would certainly require careful hand-coding of several routines. We expect to have the second-generation system running on the MicroVAX computer, which has significantly higher performance, before this becomes necessary.

Search Strategies

The first goal of the Berkeley search is to discover some modest number of supernovae as soon as possible after the explosion, and at as small a distance as possible, to permit detailed study of the explosion process. The second goal is to collect as many supernovae as possible for statistical measurement of light curves and distributions, with detection at less than 20% (type I) or 50% (type II) of maximum light being sufficient to keep the probability of detection high.

There is a complex tradeoff between exposure time, detection threshold, cycle time (between successive pictures of a given galaxy), and number of galaxies searched. In addition, the Berkley system can select individual galaxies to search, and can thus concentrate on relatively massive galaxies, or on specific types of galaxies which have been found to produce large numbers of supernovae. The maximum range for detection of type I supernovae at a fraction f_{max} of maximum light is roughly

$$D_{Mpc} = f_{max} L^{1/2} \times 10^{\{(m_{det} - M_v)/5\}} \times 10^{-5}$$

For $M_v = -19.4$ (characteristic of type I supernovae)

$$D_{Mpc} = f_{max}^{1/2} \times 10^{(m_{det} - 5.6)/5}$$

for a detection threshold of magnitude m_{det}. To the extent that galaxies are uniformly distributed in space, the number of observable galaxies is roughly

$$N = 4/3\, \pi\, D_{Mpc}^3\, d_{gal}\, f_{sky}$$

$$N = 0.04 \times 10^{3/5(m_{det} - 5.6)\}}\, f_{max}^{3/2}$$

where d_{gal}, the density of galaxies, is roughly 0.02 per Mpc^3 [17], and f_{sky}, the fraction of the sky accessible from our site in one night, is approximately 0.5. Currently, our sensitivity limits the total number of galaxies accessible ($f_{max} = 0.2$, or 0.5 for a type II supernova) to a few thousand, assuming a 30 second exposure time ($m_{det} = 17$).

In practice, we have generated our lists of galaxies from the ZCAT catalog, produced by Huchra at SAO [18], and available in machine-readable form. We generate our nightly search lists with cuts in right ascention and declination to suit the area of sky accessible from our site at a particular time of year. Galaxies are selected by recessional velocity (z value), and by apparent visual magnitude, favoring bright, nearby galaxies. We can also select by (estimated) absolute magnitude, to favor intrinsically bright galaxies with (we hope) many supernovae. The ZCAT catalog includes roughly 5000 galaxies within the limits z<5000 km/sec, $m_V \leq 16$, and declination > -30 degrees; we have reference images for less than half of these.

As our sensitivity increases (with a better CCD or improved optics), our threshold for the nearest galaxies will decrease, and the total number of accessible galaxies will increase; however, we will also reduce our exposure times and the total time spent per galaxy. With our design system, capable of searching 2500 galaxies per night, we could search 7,500 galaxies in a three night cycle, which would exhaust the Z-cat catalog. When we reach this level of performance, we will be able to do selective scanning; in particular, we will scan several hundred nearby galaxies (primarily in the Virgo cluster) every night to detect a few supernovae at the earliest possible time.

Current Status, Summary, and Conclusions

Current Status

We are now able to search for supernovae in approximately 300 galaxies per night, although we are limited by weather and logistics to the equivalent of about 2 full nights per week We take 60-second exposures, giving a supernova detection limit of 16th magnitude. We have reference images of approximately 1000 galaxies. With the aid of human scanners, we are able to process galaxy images in two to three minutes each, sufficient to keep up with our observations, although not to get far ahead.

Despite the limits of the current system, we have found several supernovae that were also discovered by others [19,20]. On 18 May 1986 we analyzed an image from the previous night and discovered supernova 1986i, a 14th magnitude type II supernova in NGC-4254 [21]. Since this first discovery, we have discovered two more supernovae, 1986n, and 1986o. A detailed report of observations of 1986i, by our group and others, will be published shortly [22].

Following our initial discovery, we expected to find more supernovae at a rate of roughly one per month with the current system; however, our actual discovery rate has been considerably lower. Since we are searching nearby galaxies, it seems unlikely that we would be completely overlooking supernovae within our search set; we might fail to observe them at early times, but they would stand out conspicuously at maximum light. We will need to observe many more galaxies, however, before we can rule out possible faults in our system or simply the statistics of small numbers as causes, and draw any conclusions about the supernova rate. Because we expect substantial performance improvements with the new computer and CCD, we are now concentrating our efforts on upgrading the system, rather than on maintaining the maximum search rate of the current system.

We have had only a few false alarms that were sufficiently convincing to require more than a cursory second look to rule them out. No false alarms have required spectroscopy or other detailed measurements to discard them. This suggests that we will be able to use a fairly simple supernova finder, plus a minimum of human effort, to identify supernovae reliably.

Summary and Conclusions

The Berkeley automated supernova search is an operating system for automated astronomy. It has discovered three supernovae, and is now searching for more. While the current performance falls short of the design goals, we fully expect to be able to reach those goals in time: to search 2500 galaxies per night with real time processing and a detection threshold fainter than 18th magnitude. However, even in its current state, the Berkeley search can make significant contributions to our observational knowledge of supernovae.

Acknowledgements

We thank the many people who have contributed to the development of the Berkeley Automated Supernova Search, including particularly Stirling Colgate, John Geary, , and, at Lawrence Berkeley Laboratory, Robbie Smits and George Gibson. Luis Alvarez has provided many suggestions. This work has been supported by a number of sources of funding, including contracts and grants from the California Space Institute, the DuPont Corporation, the Research Corporation, the National Science Foundation, the National Aeronautics and Space Administration, and the Alan T. Waterman award of the National Science Foundation. One of us (JTK) thanks the Fannie and John Hertz Foundation for fellowship support during this research. Portions of this project have been conducted under contract with the U.S. Department of Energy (Contract DE-AC03-76SF00098).

REFERENCES

1. M.G.M Cawson and E.J. Kibblewhite, Automated Supernova Search from Photographic Plates, in SUPERNOVAE: A SURVEY OF CURRENT RESEARCH, M. J. Rees and R. J. Stoneham, eds., (D. Reidel, Dordrecht, 1982).
2. J.T. McGraw, J.R.P. Angel, and T.A. Sargent, A Charge-Coupled Device (CCD) Transit Telescope Survey for Galactic and Extragalactic Variability and Polarization, in APPLICATIONS OF DIGITAL IMAGE PROCESSING TO ASTRONOMY, Ed. D.A. Ellitoo, (SPIE, Washington, 1980).
3. S.A. Colgate, E.P. Moore, and R. Carlson, PUBL. ASTRON. SOC. PACIFIC, 87, 565 (1975).
4. S. A. Colgate, Automated Supernova Search, in SUPERNOVAE: A SURVEY OF CURRENT RESEARCH, M. J. Rees and R. J. Stoneham, eds., (D. Reidel, Dordrecht, 1982).
5. J.D.E. Beynon and D.R. Lamb, CHARGE COUPLED DEVICES AND THEIR APPLICATIONS, (McGraw Hill, 1980).
6. J. C. Geary and D. W. Latham, Eds., SOLID STATE IMAGERS FOR ASTRONOMY, (SPIE vol. 290, 1981).
7. A. Fowler, P. Waddell, L. Mortara, Evaluation of the RCA 512 x 320 Charge Coupled Device (CCD) Imagers for Astronomical Use, in SOLID STATE IMAGERS FOR ASTRONOMY, J. C. Geary and D. W. Latham, Eds., (SPIE vol. 290, 1981).
8. J. R. Janesick, Jet Propulsion Laboratory, private communication.
9. H. Gursky, J. Geary, R. Schild, T. Stephenson, and T. Weekes, Astronomical Performance of the 512 x 320 RCA Charge Coupled Device (CCD), in APPLICATIONS of DIGITAL IMAGE PROCESSING TO ASTRONOMY, D. A. Elliott, Ed., (SPIE vol. 264, 1980).
10. J. Culver, A CCD Camera for an Automated Supernova Search, Senior Honors Thesis, University of California, Berkeley (1982).
11. M. Shane Burns, Development of a CCD Camera for an Automated Supernova Search and Observations of a Supernova in NGC 5033, Ph.D. Thesis, Dept. of Physics, U.C. Berkeley, LBL#20648, Nov., 1985.
12. Treffers, Richard R., The Leuschner Observatory Automated Telescope, Pub. Astron. Soc. Pac., May 1985.
13. M. Davis, P. Hut, and R.A. Muller, Extinction of Species by Periodic Comet Showers, NATURE, 308, #5961, 715-717, 19 April 1984.
14. S. Ebstein, Digital Recording on Videotape, Review of Sci. Inst., 54, #7, 883-885, July

1983.

15. Kare, J. T., AN AUTOMATED SEARCH FOR SUPERNOVAE, Ph. D. Thesis, University of California, Berkeley, November, 1984.
16. VISTA is available from Lick Observatory, University of California at Santa Cruz, 1987.
17. C. W. Allen, ASTROPHYSICAL QUANTITIES, 3rd Edition (Athlone, London, 1973), p. 289.
18. Huchra, John, ZCAT, Harvard-Smithsonian Center for Astrophysics, 1985.
19. Metlova, N, Supernova in NGC 5033, IAU CIRCULAR, #4077, June, 1985.
20. Evans, R.E., Supernova in NGC 3367, IAU CIRCULAR, #4173, Feb, 1986.
21. Pennypacker, C.R., et al, IAU CIRCULAR, #4219, May, 1986.
22. Pennypacker, C.R., et al, "Observations of the Type II Supernova in M99," to be submitted to ASTRONOMICAL JOURNAL, July, 1987.

Figure Captions

FIGURE 1: Design of the Berkeley System -- Overview. The Berkeley system, shown in figure 1, uses a computer controlled, fast-slewing telescope and charge coupled device (CCD) camera to collect images of galaxies; these images are compared in the lab to stored reference images by a dedicated minicomputer.

FIGURE 2: Similar system as in Figure 1, except upgraded for real-time system comparisons of reference and new images at the observatory.

FIGURE 3: Fraction of total brightness in 1, 4, 9 pixels. The random variable f depends on the seeing, and on the position of the supernova image with respect to CCD pixel boundaries. In the limit of perfect seeing, f is 1.0. For typical seeing at the Leuschner site and 2.4 arc second square pixels, we find a mean value for foreground stars (brightest pixel/total brightness) of 0.18. Figure 3 is a plot of f for 17 stars. Note that we can actually improve our sensitivity slightly by checking blocks of 4 or 9 pixels -- f(4 pixels) is typically 0.5, while the noise level is twice that for one pixel, so the effective value of f is 0.25.

FIGURES 4 and 5: Current (4) and Future(5) Thresholds for Detection. We plot the number of photoelectrons detected versus t_{exp} for our current system and our nominal (design) system. The signal from various magnitudes of supernova is plotted, along with the system noise level (signal to noise level of 1) and detection threshold (s/n of 5) n_{sn}=6.5 x 10 {((20 - m sub v) / 2.5)} . This value is based on observations of standard stars of known visual magnitude using our current CCD at the Leuschner telescope. Figure 5 assumes the increased collecting area and lower sky background available at the MIRA or that typical of an observatory at a dark site, and improved reducing optics with 75% transmission rather than our current 50%. Improved seeing and better pixel scale are assumed to increase f to 0.5. .Note that for short exposures, CCD noise dominates over sky noise.

FIGURE 6: Galaxies per night vs. exposure time. The time per galaxy is the time for the telescope motion (t_{move}) plus the exposure time (t_{exp}), or t_{tot} = t = t_{move} + t_{exp} = 3sec + t_{exp} . Figure 6 plots the number of galaxies observable per (6-hour) night versus t(exp) for these two cases.

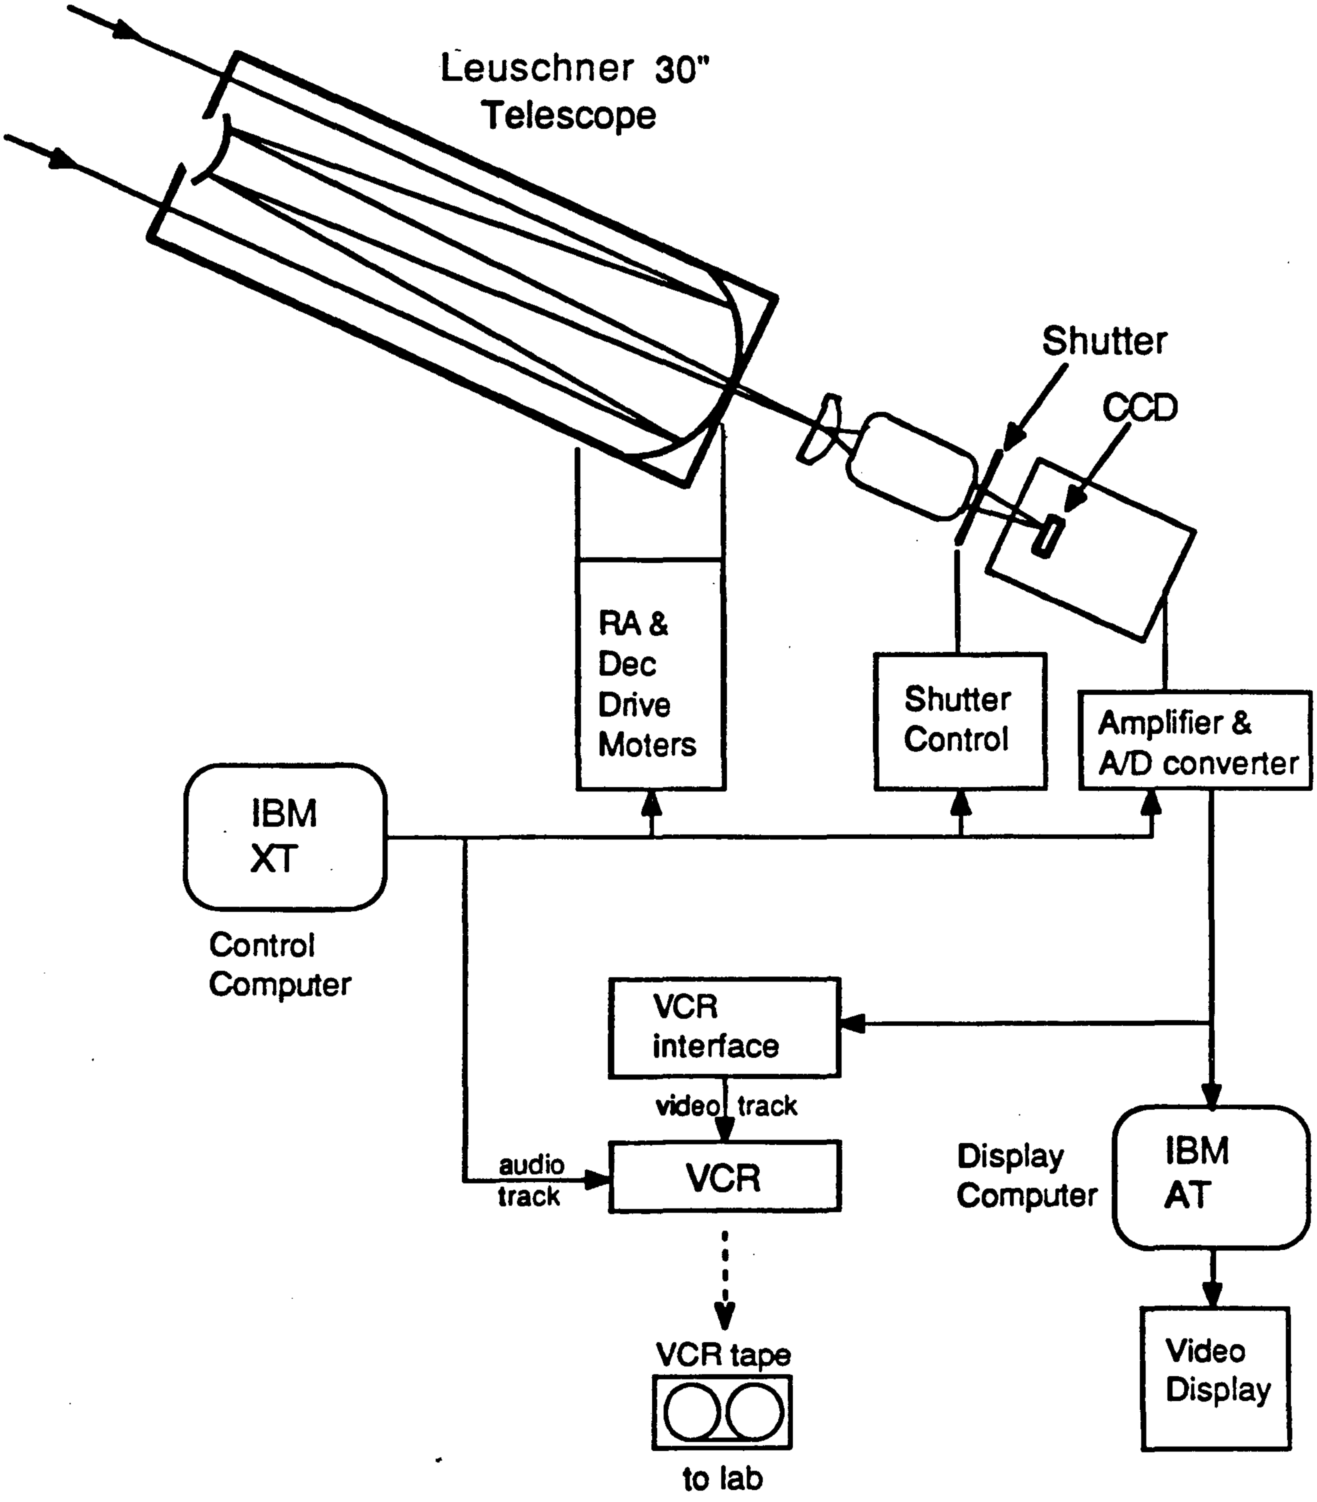

Figure 1: Current System

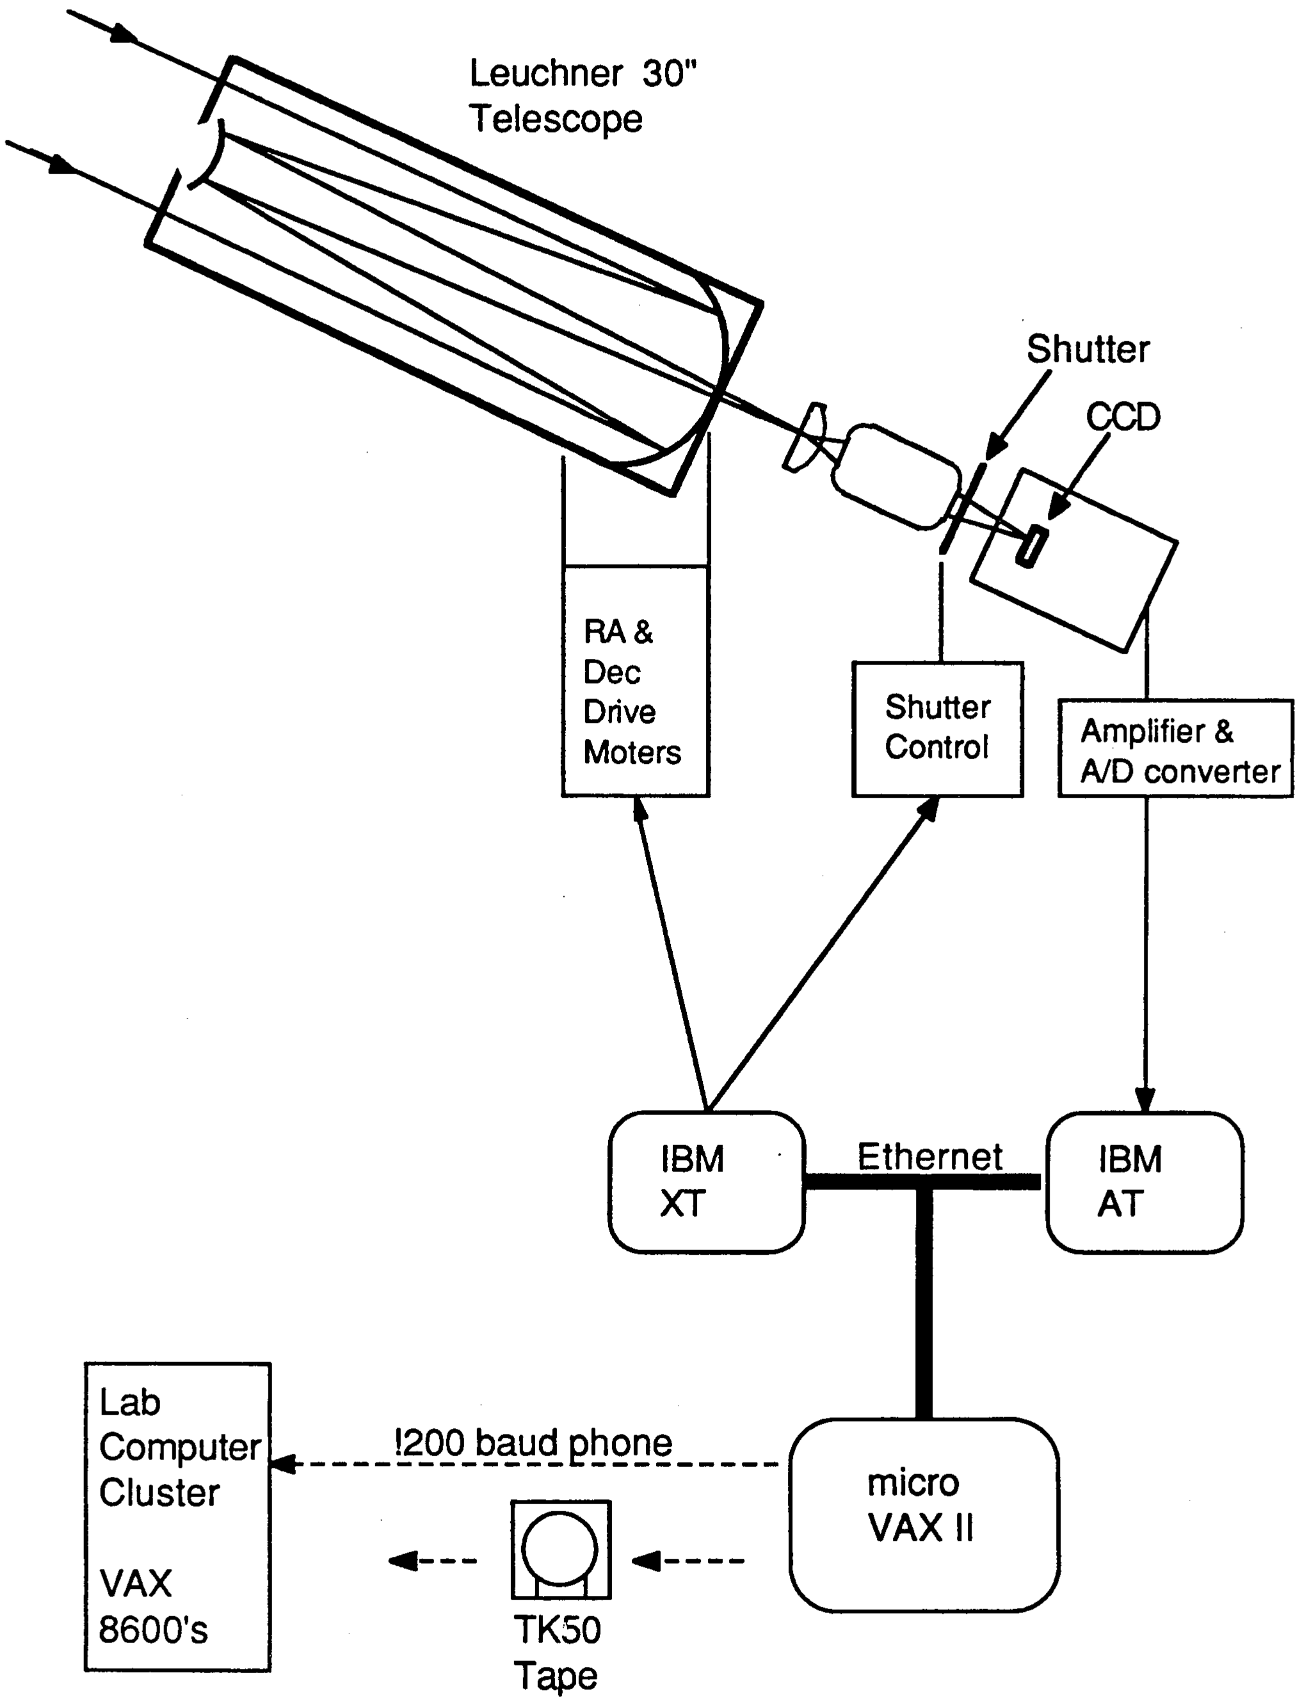

Figure 2: New System

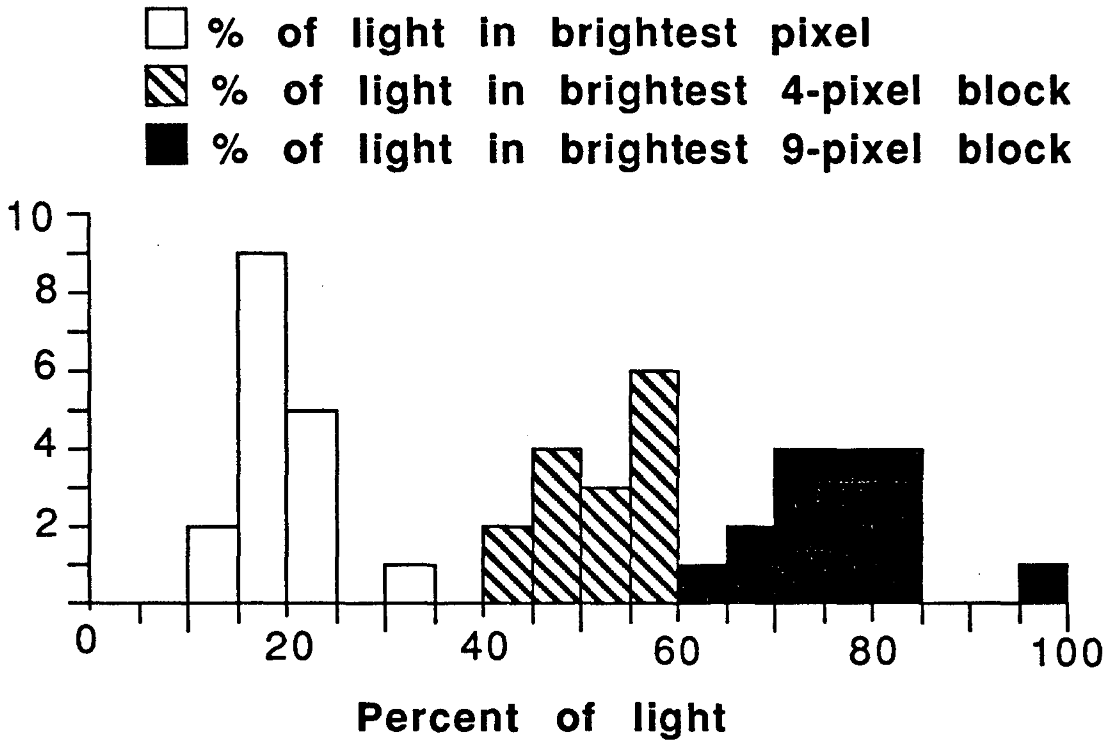

Figure 3: Starlight Distribution in Pixels

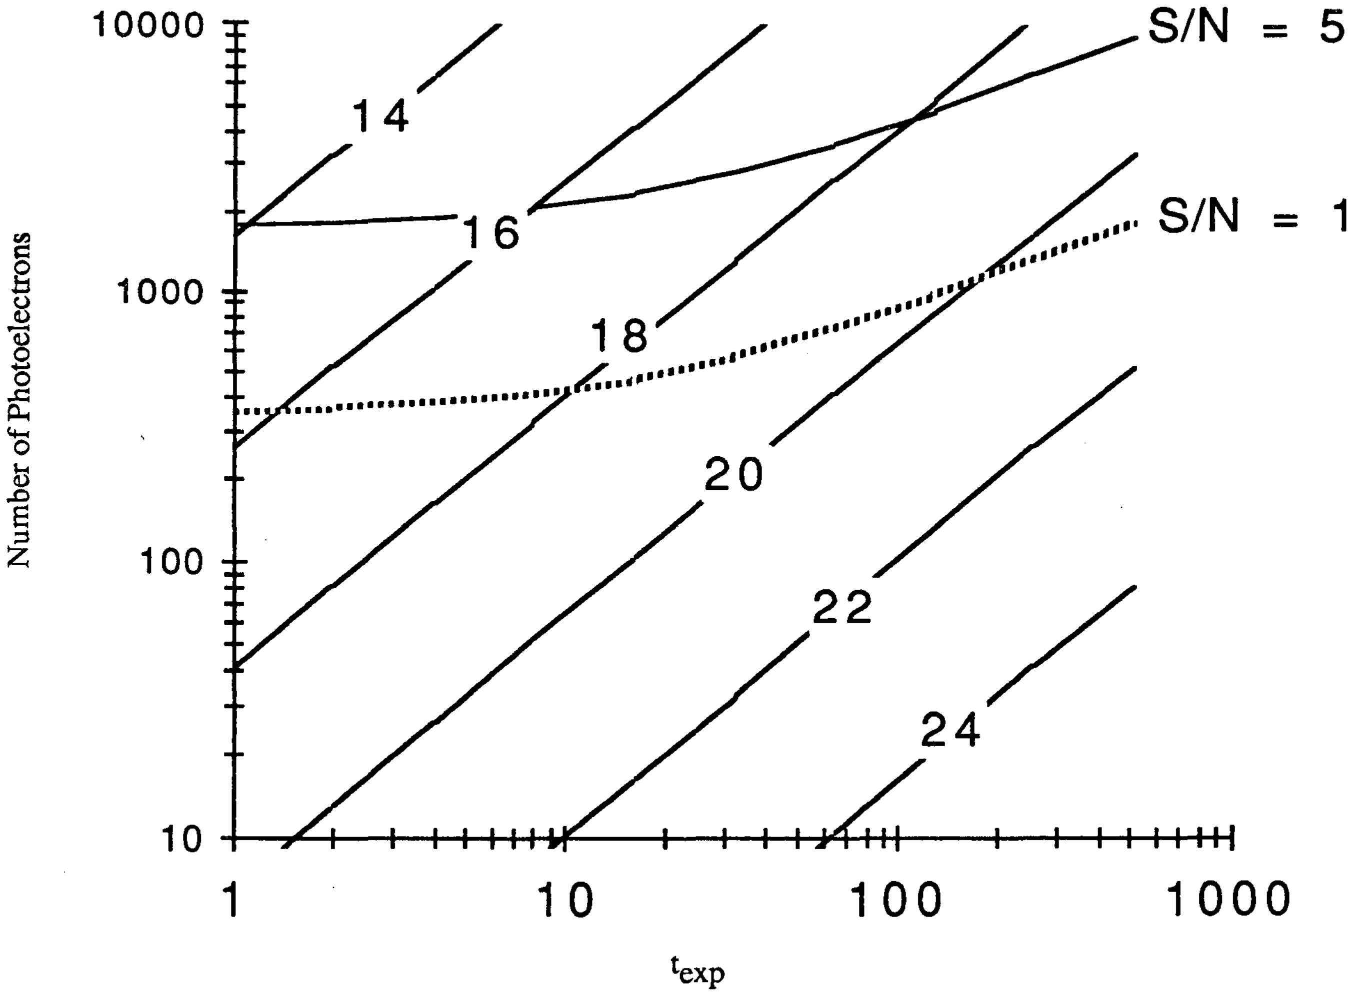

Figure 4: Leuschner Sensitivity

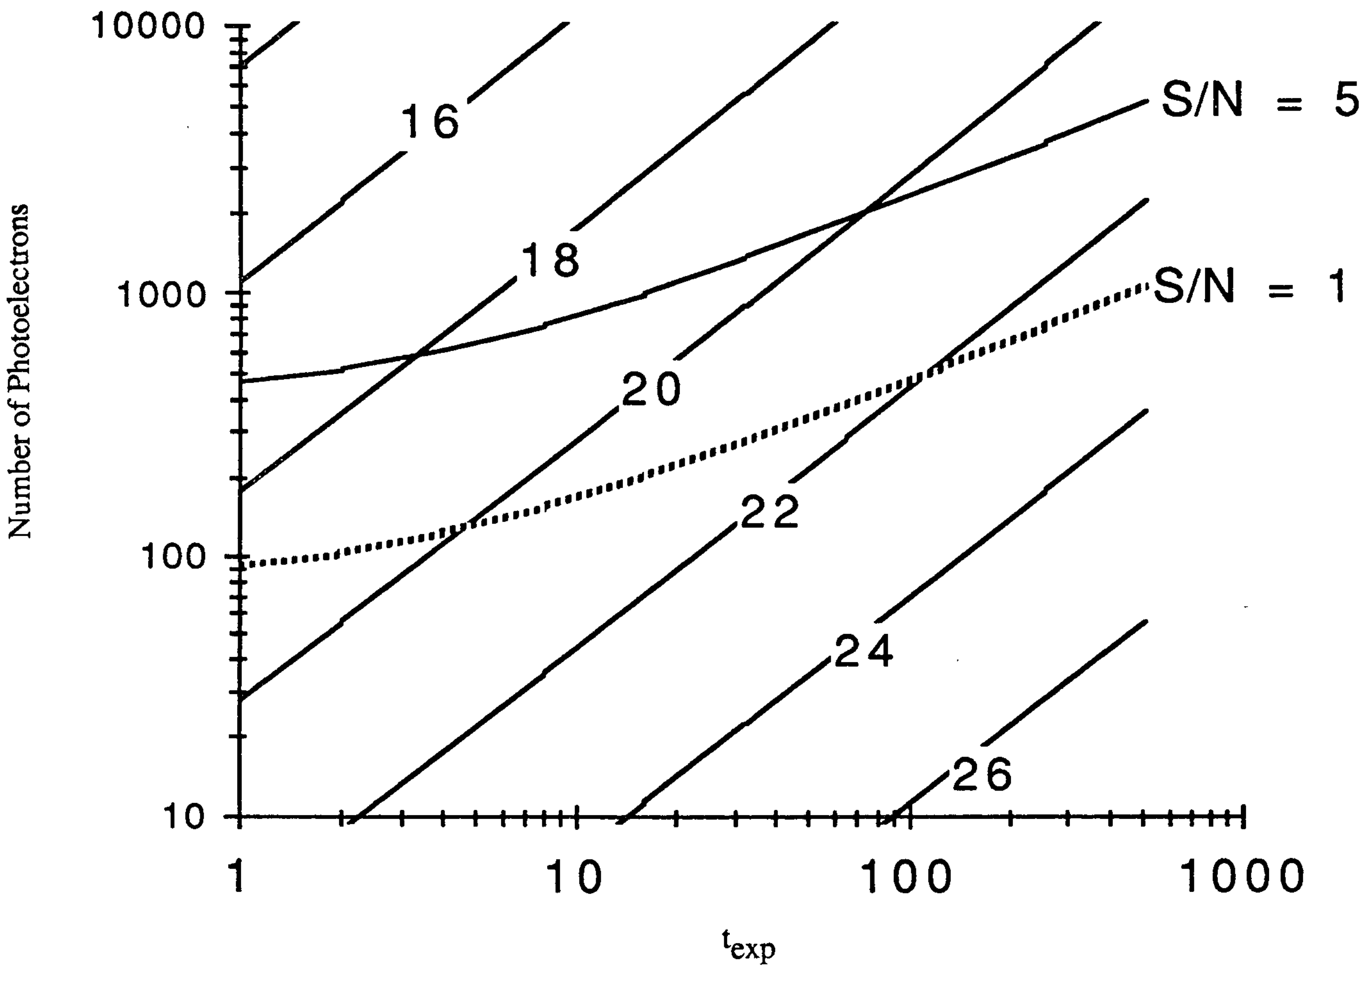

Figure 5: Sensitivity at Good Site

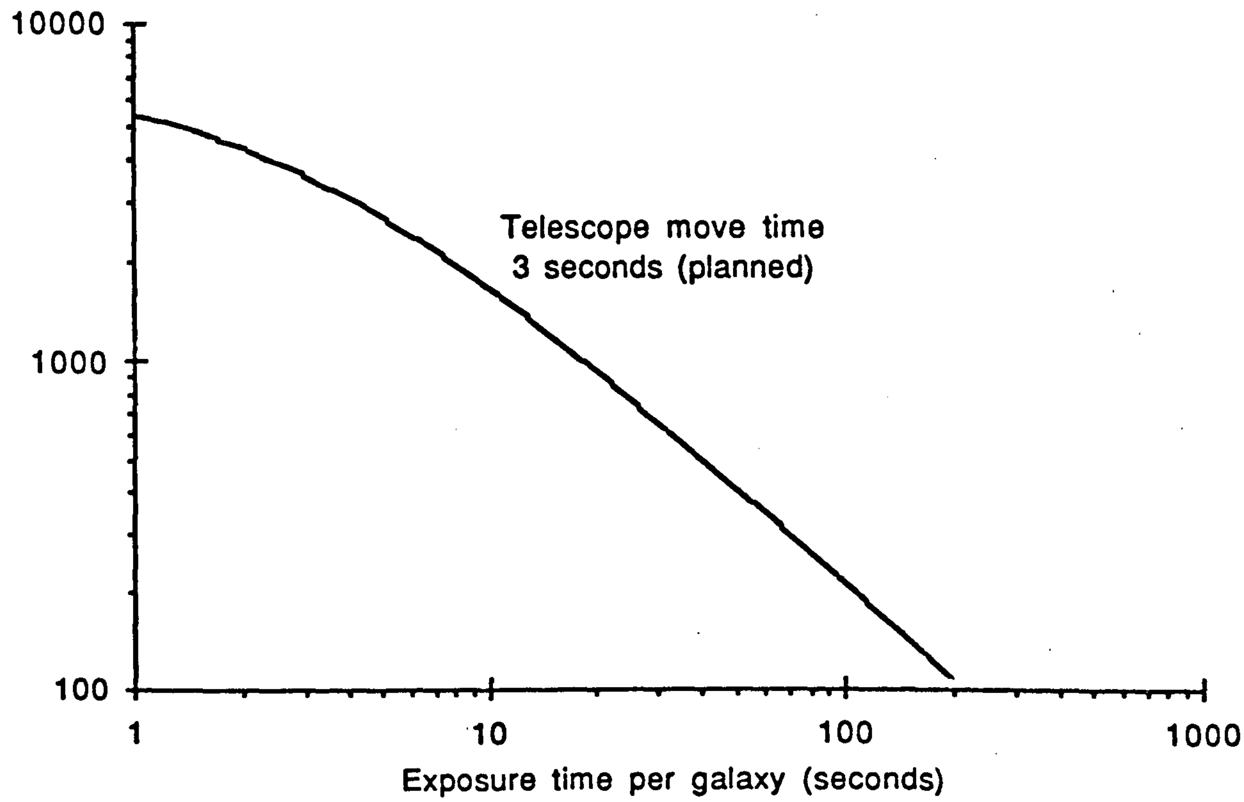

Figure 6: Galaxies per Night as a Function of Exposure Time

LAWRENCE BERKELEY LABORATORY
TECHNICAL INFORMATION DEPARTMENT
UNIVERSITY OF CALIFORNIA
BERKELEY, CALIFORNIA 94720