Lawrence Berkeley National Laboratory
Recent Work

Title
AN AUTOMATED SEARCH FOR SUPERNOVAE

Permalink
https://escholarship.org/uc/item/9809b94q

Author
Kare, J.T.

Publication Date
1984-11-01

eScholarship.org Powered by the California Digital Library
University of California

UC-34b
LBL-19340
c1

Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Physics Division

RECEIVED
LAWRENCE
BERKELEY LABORATORY
MAY 16 1985
LIBRARY AND
DOCUMENTS SECTION

AN AUTOMATED SEARCH FOR SUPERNOVAE

J.T. Kare
(Ph.D. Thesis)

November 1984

For Reference
Not to be taken from this room

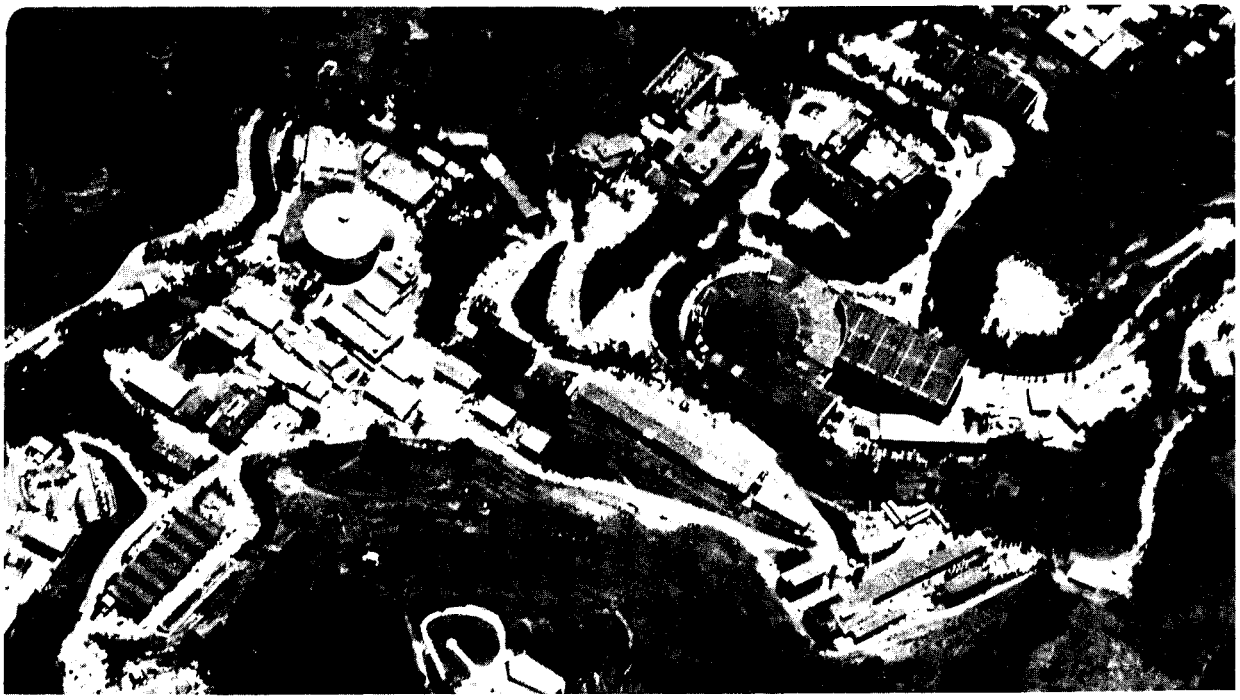

Prepared for the U.S. Department of Energy under Contract DE-AC03-76SF00098

LBL-19340
c.1

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

LBL-19340

An Automated Search for Supernovae

Ph.D. Thesis in Physics

Jordin T. Kare

Lawrence Berkeley Laboratory

And

The University of California, Berkeley

Berkeley, CA 94720

15 November 1984

AN AUTOMATED SEARCH FOR SUPERNOVAE

Jordin T. Kare

ABSTRACT

Supernovae are among the most important, as well as most spectacular, of astrophysical events. However, because supernovae are both unpredictable and short-lived, they are difficult to study in real time with traditional observational techniques. In particular, we have very little information about the development of supernovae in their first few days -- a few luminosity measurements, very few spectra, and no observations in the radio, infrared, X-ray, and gamma ray bands. There is thus a great need for a system capable of locating large numbers of bright supernovae at early times.

This thesis describes the design, development, and testing of such a search system 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 system is now collecting test images of up to several hundred fields per night, with a sensitivity corresponding to a limiting magnitude (visual) of 17. At full speed and sensitivity, the search will examine some 6000 galaxies every three nights, with a limiting magnitude of 18 or fainter, yielding roughly two supernovae per week (assuming one supernova per galaxy per 50 years) at 5 to 50 percent of maximum light. An additional 500 nearby galaxies will be searched every night, to locate about 10 supernovae per year at one or two percent of maximum light, within hours of the initial explosion.

Table of Contents

Acknowledgements

In my work on the Berkeley Automated Supernova Search, I have received help, advice, and assistance from a great many people. To begin, I would like to thank Richard Muller, my thesis advisor and mentor, and Luis Alvarez, in turn Richard Muller's advisor nad mentor, both of whom have taught me a great deal about what it means to do good physics, as well as about the details of how to get it done.

Carl Pennypacker, project leader for the supernova search, frequently assisted me with my portion of the search project, as well as shouldering his own share of the burdens. My fellow students, notably Shane Burns and Saul Perlmutter, have also been of great help, and Roger Williams, in his late night efforts to keep our software structures from collapsing of their own weight, has been invaluable. Thanks also to the various undergraduates who have passed through the supernova search, notably John Culver, whose senior thesis has been useful indeed.

I thank the entire LBL astrophysics group, and particularly Terry Mast, for insightful questions and discussions. Of those outside the group, I would like to thank the board of the Fannie and John Hertz Foundation, and particularly Lowell Wood and Tom Weaver, for their extended support, personal as well as financial. And finally, I thank Frank Crawford, without whose Hot Chocolate Effect I might all too easily have forgotten that physics can be fun.

The work described in this thesis has received support from numerous sources. The author has been supported primarily by a fellowship from the Fannie and John Hertz Foundation. Other sources of funding are 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.

Portions of this project have been conducted under contract with the U.S. Department of Energy (Contract DE-AC03-76SF00098).

"And this is my quest, to follow that star,
No matter how hopeless, no matter how far..."

--Man of La Mancha

1. Background

1.1. Introduction

Supernovae are among the most spectacular and energetic of astrophysical events. They are, however, both rare and short lived. Past searches, largely based on photographic techniques, and accidental discoveries have identified only a few hundred supernovae; detailed light curve or spectral measurements exist for only a handful. Measurements of the properties of supernovae before maximum light are limited to a few supernovae discovered slightly before maximum, and to supernovae captured by accident on pre-discovery photographs. Also, no mechanism has existed for discovering supernovae in nearby galaxies on a sufficiently regular basis to permit planning for studies of active supernovae outside of the visible spectrum.

The major goals of a modern supernova search are 1) to find nearby supernovae at the earliest possible time, for detailed study in all energy ranges, 2) to find and follow a large number of supernovae at intermediate distances, for collection of improved statistical information on light curves and distribution, and 3) to locate large numbers of supernovae at very large distances, for measuring supernova rates and cosmological distances. The Berkeley automated supernova search described herein has been developed to meet the first two of these goals, with performance far exceeding that of older photographic searches. It is complementary to other new automated searches which are most efficient at finding distant supernovae near maximum light.

1.2. The Nature of Supernovae

1.2.1. Properties of Supernovae

There are at least two varieties of supernovae, identified as Type I and Type II [1]. As shown in figure 1.1a [2], type I supernovae very roughly double in brightness every 24 hours, reaching maximum luminosity about 1 week after the initial explosion. At maximum, they reach an absolute luminosity of 10^{43} ergs per second, or 10^{10} that of the sun. The absolute magnitude (visual) at maximum is approximately -19.4 (from Tammann, assuming H_0=50) [3]. After maximum, Type I's drop in luminosity by 2 to 3 magnitudes in another 20-30 days, after which the luminosity falls exponentially with a characteristic 77 day decay time (56 day half life). The most distinctive feature of the type I spectrum is a complete absence of hydrogen lines. Type I supernovae are found in all types of galaxies, including ellipticals, suggesting that they are associated with old stellar populations.

Type II supernovae are typically 1 magnitude fainter than type I's; Tammann estimates M_v=-18.6 at maximum. Models indicate that the type II light curve (figure 1.1b [4]) reaches maximum light very quickly and has a relatively complex pattern of decay. Type II's have strong hydrogen lines in their spectra, and are associated with spiral arms; they are not found in elliptical galaxies. In this paper, the properties of a "supernova", without any qualifiers, will generally be those of a typical type I supernova.

1.2.2. Current Models of Supernovae -- Type I

One currently popular model for the predecessor of a type I supernova [5] is a white dwarf star in a binary system. The dwarf is composed of approximately 1.4 solar masses of carbon and oxygen. Hydrogen transferred from the binary companion

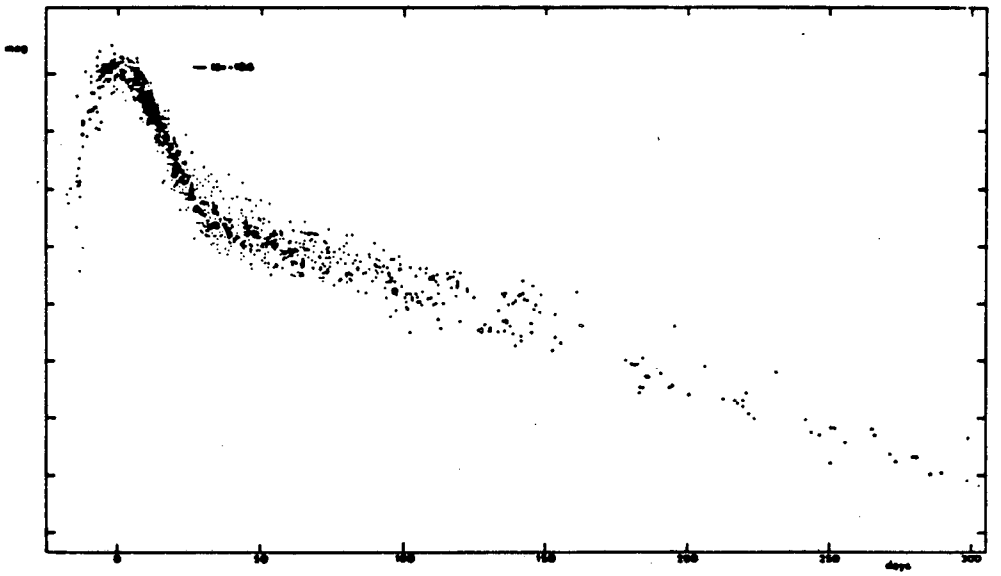

Fig. 1. Composite blue light curve obtained by the fitting of the observations of 38 type I supernovae. One magnitude intervals are marked on the ordinates.

Figure 1.1a (from [2])

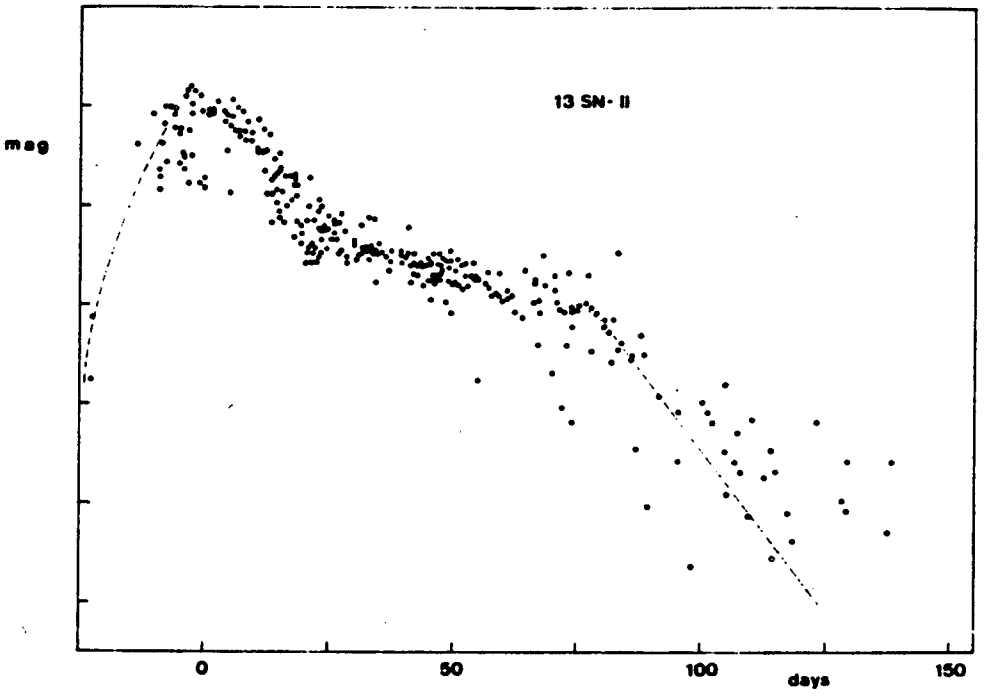

Fig. 1. Composite light curve obtained from 13 type II supernovae.

Figure 1.1b (from [4])

is burned stably to helium, and then to C or O. At a fixed critical mass, thermonuclear ignition occurs, although the details of the ignition and burning process are unknown. Essentially the entire mass of the star is converted to ^{56}Ni, the most stable (highest binding energy) isotope under the conditions prevailing during the explosion.

The shape of the early part of the type I light curve is determined by the cooling and expansion of the explosion debris, and by energy from the decay of ^{56}Ni into ^{56}Co. Initially, the black body temperature may be above 10^5 K, and the visible spectrum rises with frequency as ν^2, so pre-maximum type I supernovae are quite blue. Near maximum light, the temperature is of order 10^4 K, and as the debris continue to cool, the total visual luminosity falls.

The 77 day tail is now believed to be powered by beta decay of ^{56}Co into ^{56}Fe [6] with a half life of 56 days. Axelrod [7] has been able, using very detailed stellar atmospheric models, to roughly reproduce the evolution of the spectrum of SN 1972e, and to identify one particular peak as characteristic of the decay of ^{56}Co to ^{56}Fe.

1.2.3. Current Models of Supernovae -- Type II

The hydrogen-containing spectra and galactic distribution of type II supernovae suggest that they are associated with massive, short lived giant stars. Models for the evolution of such stars, particularly the detailed models of Weaver and Woosley [8] indicate that, once such a star has completed hydrogen burning in the core, progressively heavier elements are burned until a core of ^{56}Fe and nearby isotopes is formed. When the mass of this core reaches approximately 1.4 solar masses, it becomes unstable against gravitational collapse. The core collapses on a time scale of milliseconds, reaching nuclear densities and releasing enormous quantities of neutrinos. Through some combination of shock waves produced by the "core bounce" at nuclear densities, and the deposition of energy or momentum by outflowing neutrinos,

approximately 10^{51} ergs of the energy released by the collapse is transferred to the outer layers of the star, blowing them into space [9].

The luminosity and light curve of a type II supernova are determined by the extended envelope surrounding the collapsing core; the collapse itself is only observable through such currently unavailable means as neutrino or gravity wave detectors. The rapid rise in brightness is associated with the shock wave of the explosion reaching the stellar surface, and the total luminosity and the shape of the light curve depend on the details of the stellar envelope.

1.2.4. The Supernova Rate

The most detailed estimates of the supernova rate have been made by Tammann [10], who obtains typical rates of 1 per 20 years to 1 per 50 years per galaxy, depending on galaxy type. Table 1.1, reproduced from [10], gives estimates of the supernova rate per 10^{10} solar luminosities for different types of galaxies. More pessemistic estimates have been made [11], but those of Tammann seem to be relatively widely accepted.

For the purposes of this paper, and in the design of the Berkeley automated search, we have assumed a real supernova rate between 1 per 50 years and 1 per 100 years per galaxy searched. These figures can be shifted somewhat for the Berkeley search by deliberately selecting bright (and therefore massive) galaxies, and by concentrating on particular galaxy types, such as those identified by Tammann as having high supernova rates, but we take this as a reasonably conservative range.

1.3. A Summary of Past and Current Supernova Searches

The earliest recorded supernova observations are of supernovae occuring in the unobscured portion of our own galaxy, and thus visible to the naked eye. Records of

Adopted SN Frequencies in SNu (per 10^{10} $L_{B\odot}$[1]) per 100 yr) for Different Types of Galaxies

	All SNe	SNe I	SNe II	n_{SN}[2]
E	0.22	0.22	0	13
S0	0.12	0.12	0	6
S0a, Sa	0.28	0.28	0	9
Sab, Sb	0.69	0.37	0.32	38
Sbc, Sc, Scd, Sd	1.38	0.77	0.61	93
Sdm, Sm, Im	1.02	0.83	0.19	11
I0	undetermined			7

1) The B-luminosities are in the B_r-system of the RC2, they are corrected for galactic and internal absorption as outlined in Sandage and Tammann (1981).

2) n_{SN} is the number of SNe in sample B; these numbers give an estimator of the statistical error of the listed frequencies.

Table 1.1

Supernova rate vs. Galaxy type (from Tammann [10])

some 8 such events in the last 2000 years have been found [12], with varying amounts of detail and confirmation. The average rate for such local supernovae is thus about 1 per 250 years. Clark and Stephenson [13] estimate that the unobscured fraction of our galaxy is about 1/7, and that of order 1/2 of the visible supernovae are recorded in the records we have; this is compatible with a true supernova rate in this galaxy as high as one per 20 years.

Unfortunately, no such local supernovae have occurred since AD 1604, so all modern observations have been of extragalactic supernovae. The first such observation was in 1885 [14]. The first modern photographic search was begun by Zwicky in 1933 [15] using a 3.25" camera lens as a wide-field objective, and continued on the 18" and 48" Palomar Schmidt cameras. The Palomar 48" Schmidt searches detected a total of 281 supernovae [16] before being discontinued in 1975. Other Schmidt telescope searches have been conducted, for example at Asiago [17] and Cerro El Roble [18].

All of these searches have a common procedure: Plates of selected fields are taken at regular intervals, and pairs of similar plates are compared by superposition, or using a blink comparator. Depending on the degree of effort expended, the interval between plates for a single field is typically several days, and may be weeks or even months. Processing and scanning of plates may also require several days; thus, supernovae are rarely detected at very early times, and in many cases are not detected until well after maximum light. The limited dynamic range of photographic plates, combined with the limitations of human scanners, makes detection of supernovae in galactic cores difficult. Finally, as pointed out by Tammann [19] and others, extreme care is required to obtain useful statistics from photographic searches, since the set of galaxies sampled is poorly defined.

In the late 1960's, Hynek [20] constructed a system based on intensified vidicon detectors and video displays; comparison of fields was still done by eye, but supernovae could be detected promptly. This system found only one supernova before maximum light, and operated for only a short time before being shut down.

In an attempt to increase the speed and accuracy of photographic searches, Kiblewhite [21] and others are using automated plate measuring machines to rapidly locate supernovae on plate pairs. Such searches are very good at finding distant supernovae near maximum light, but, because of limited sky coverage and the logistics of handling plates, they are not very efficient at catching nearby supernovae at early times. Galaxy samples on such plates are still poorly defined, although automatic counting and classification of galaxies can be done. The very large amount of data present on wide field photographic plates is an additional problem; a sufficiently fast and powerful automated measuring machine system represents a very large investment, and generally cannot be dedicated to searching for supernovae full time.

In the early 1970's, Stirling Colgate suggested that a supernova search be constructed using a computer controlled telescope capable of examining a large number of galaxies, one at a time, each night. An electronic imaging detector would send images directly to a computer, which could compare each picture to a stored reference and immediately report supernovae. Colgate has constructed such a system [22], using an intensified image tube detector and a telescope constructed around a surplus radar antenna mount [23]. This system is still in existence, and although handicapped by relatively old technology and limited support, has been able to collect reference images. It has not at this writing begun finding supernovae.

1.4. The Berkeley Automated Supernova Search

1.4.1. History

In 1978, one of our group (Richard Muller) suggested that we follow Stirling Colgate's proposal, but apply the latest technology to the problem. In particular, he and Carl Pennypacker proposed taking advantage of large area, high quantum efficiency charge coupled device (CCD) detectors to obtain greater sensitivity than previously possible. Also, the declining price of minicomputers and microcomputers made purchasing a dedicated computer for image analysis, and additional hardware for automated telescope control, relatively easy. We therefore chose to proceed with a new automated search for supernovae.

Construction of the search began in 1979, and the system computer was purchased in early 1980. We obtained a "loaner" CCD from RCA, and the CCD camera was operational in 1981. The quality of this CCD was extremely poor, but we were able to test the camera and a limited amount of software. We obtained and installed an improved loaner CCD in the summer of 1983; this device is relatively noisy and has low quantum efficiency, but is basically usable. At this writing, RCA has

delivered a new CCD, but we have not had an opportunity to test it. This CCD is expected to have good quantum efficiency and transfer efficiency, but may still have relatively high (60 photoelectron) noise.

We spent most of the period from 1981 to spring 1984 writing software for the supernova search, and testing the CCD system on the Monterey Institute for Research in Astronomy (MIRA) 36" telescope. MIRA suffered delays in moving to their permanent observatory site, and in getting their computerized telescope control system operational, so our search was moved to the Leuschner Observatory 30" telescope in the spring of 1984. From the spring of 1984 to the present, work on the supernova search has consisted of further software development, interfacing of the Leuschner telescope control system to the search computer, and observations, with observing time currently divided between collection of referece galaxy pictures for the supernova search, and a search for a solar companion star.

1.4.2. Overview of the Berkeley Search System

The hardware for the Berkeley system is shown in figure 1.2. A small (by astronomical research standards) telescope collects light from a galaxy and focuses it through reducing optics onto the CCD detector, which operates in a vacuum housing at a temperature of -120 C. The shutter opens for, typically, 30 seconds. The telescope then moves, under computer control, to another galaxy, while the CCD is read out into a computer.

Once an image is acquired, processing proceeds as shown in figure 1.3. The computer determines the coordinates and brightness of the current image using bright foreground stars (fiducial stars); this information can be fed back to the telescope to correct for pointing error. The current image is aligned with a stored reference image of the same galaxy; if the two match, no action is taken.

THE BERKELEY AUTOMATED SUPERNOVA SEARCH SYSTEM

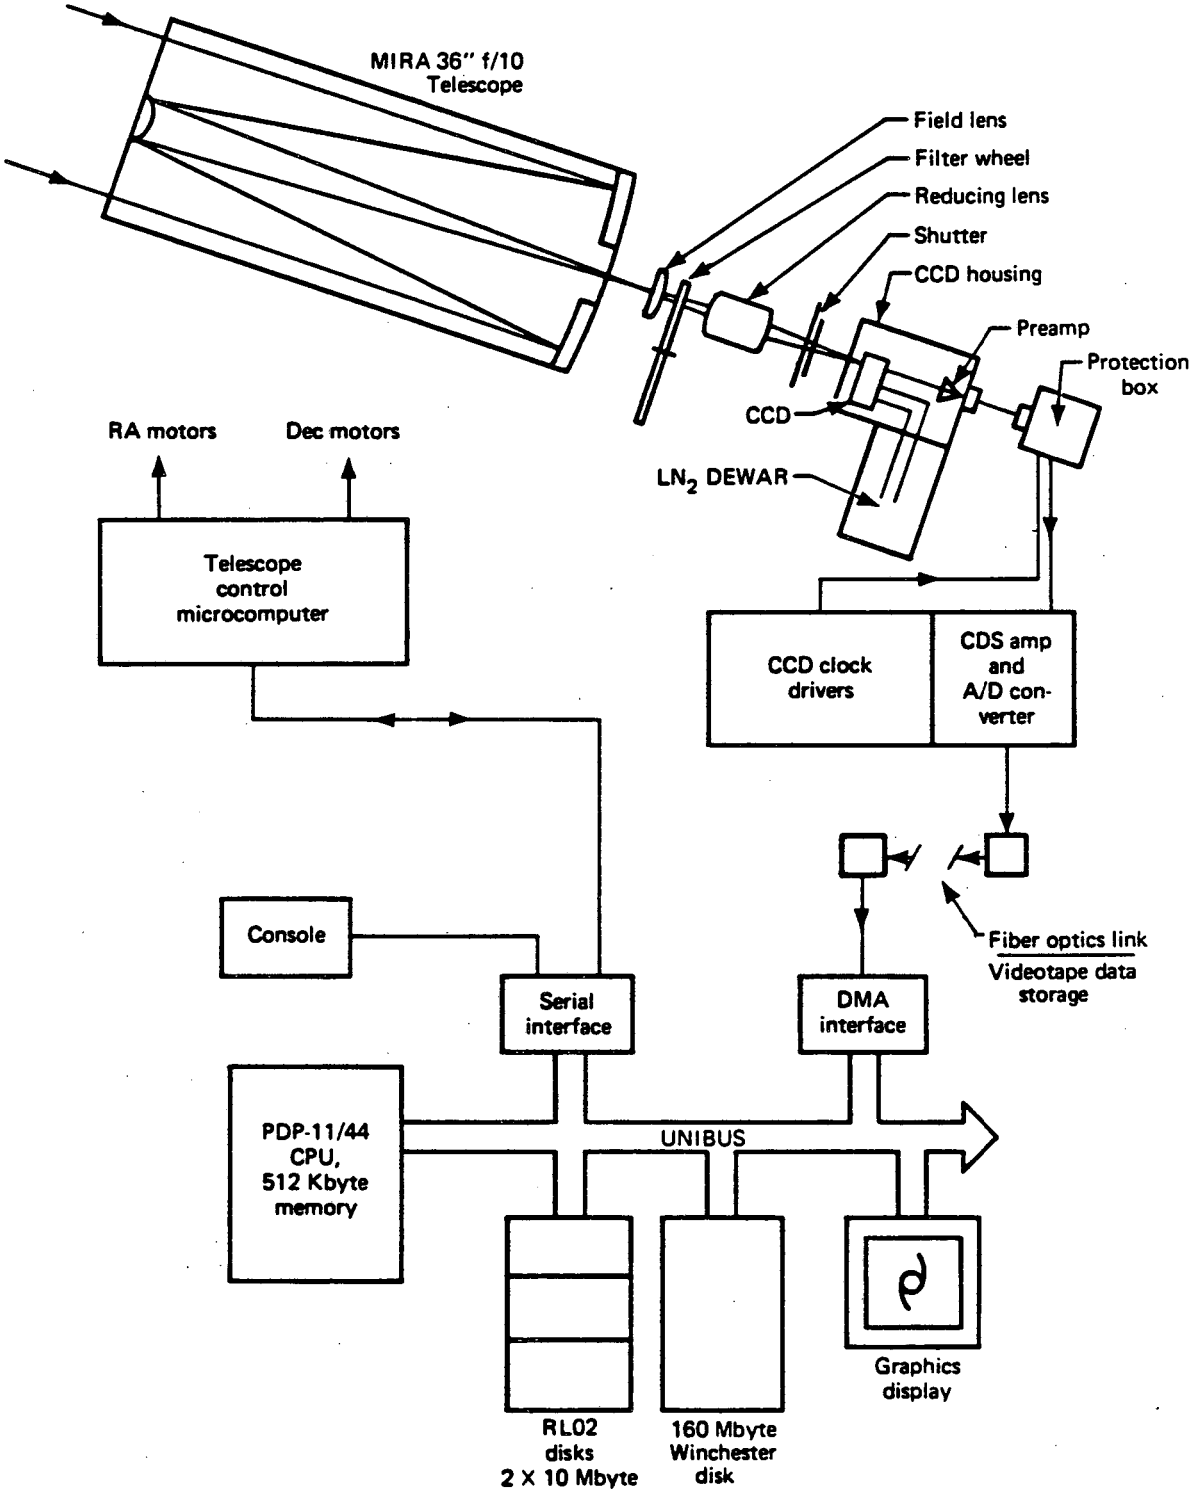

XBL 836-10439

Figure 1.2

AUTOMATED SEARCH SOFTWARE

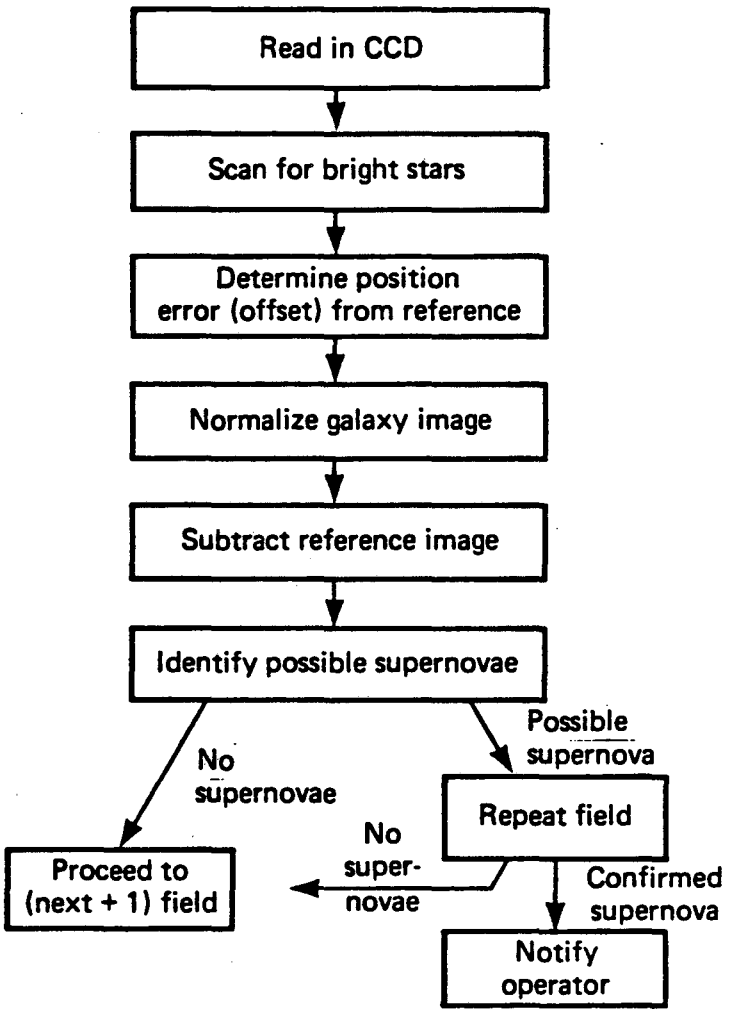

Figure 1.3

If the current image contains an extra point source that the software considers a possible supernova, the telescope is directed to collect a second image of that galaxy. If the same new source appears in this second image (and in a third, taken after a few minutes delay), the software reports a possible supernova to the system operator, who can inspect the images, request further observations, and, if satisfied, announce the supernova to the astronomical community.

The system can (and currently does) operate "offline", with CCD images stored in digital form on videotape for analysis the following day. Additional observations, including broad band and color photometry, can be made directly, while observations beyond the capacity of the search telescope can be made in cooperation with other observatories.

1.4.3. Goals of the Berkeley Search

The design goal for the supernova search is to be able to monitor 2500 galaxies per night, with a total time per galaxy of 8 seconds and a limiting magnitude of approximately 18.5, sufficient to detect a type I supernova in the Virgo cluster at less than 1% of maximum light. By searching approximately 500 nearby galaxies each night, we would expect to catch 5-10 supernovae per year at 1% of maximum, for detailed study. By searching an additional 6000 galaxies on a 3 night cycle, we would expect to catch 1-2 supernovae per week, for statistical measurements of light curves, luminosity, and distribution within and among galaxies.

In addition to the general utility of the search to the astronomical community, there are a number of specific questions which the search will address. A few of these are:

Early time spectra -- Because of the high sensitivity and rapid response time

of the Berkeley search, an occasional supernova will be detected at very early times. By collaborating with other observers, we hope to use "target of opportunity" time on large telescopes (and, eventually, satellites) to obtain detailed spectra of type I supernovae at very early times. These spectra will substantially extend the range of data available to modellers of supernova atmospheres, and should be of great value in constraining models of supernova explosions.

Early time anomalies — Figure 1.4 [24] is a computer-generated light curve for a particular (and relatively exotic) type I supernova model, starting with a 9 solar mass red supergiant which has lost most of its envelope. The spike in luminosity occurring immediately after the explosion is associated with the shockwave of the explosion reaching the stellar surface and heating the remaining low-density helium shell; it is bright enough, and long enough in duration, that it might occasionally be detected by the Berkely search in nearby supernovae. While this particular model is not likely to be accurate, it is at least possible that the very early time light curve will show interesting behavior detectable only with an automated search system.

Underluminous type II's — There is a substantial discrepancy between the number of type II supernovae observed, and the apparent birth rate of pulsars, which are presumed to be remnants of type II core collapse. One can account for this by postulating the existence of "dim" supernovae [25]; stars which suffer core collapse after losing most of their outer hydrogen envelope, and which therefore emit much less visible light. One candidate for such an event is the supernova which produced the Cassiopeia A remnant; the explosion should have been a conspicuous naked eye event, but is not mentioned in the historical record. If such dim supernovae do occur, the Berkeley search will be the first to be able to detect them reliably in nearby galaxies.

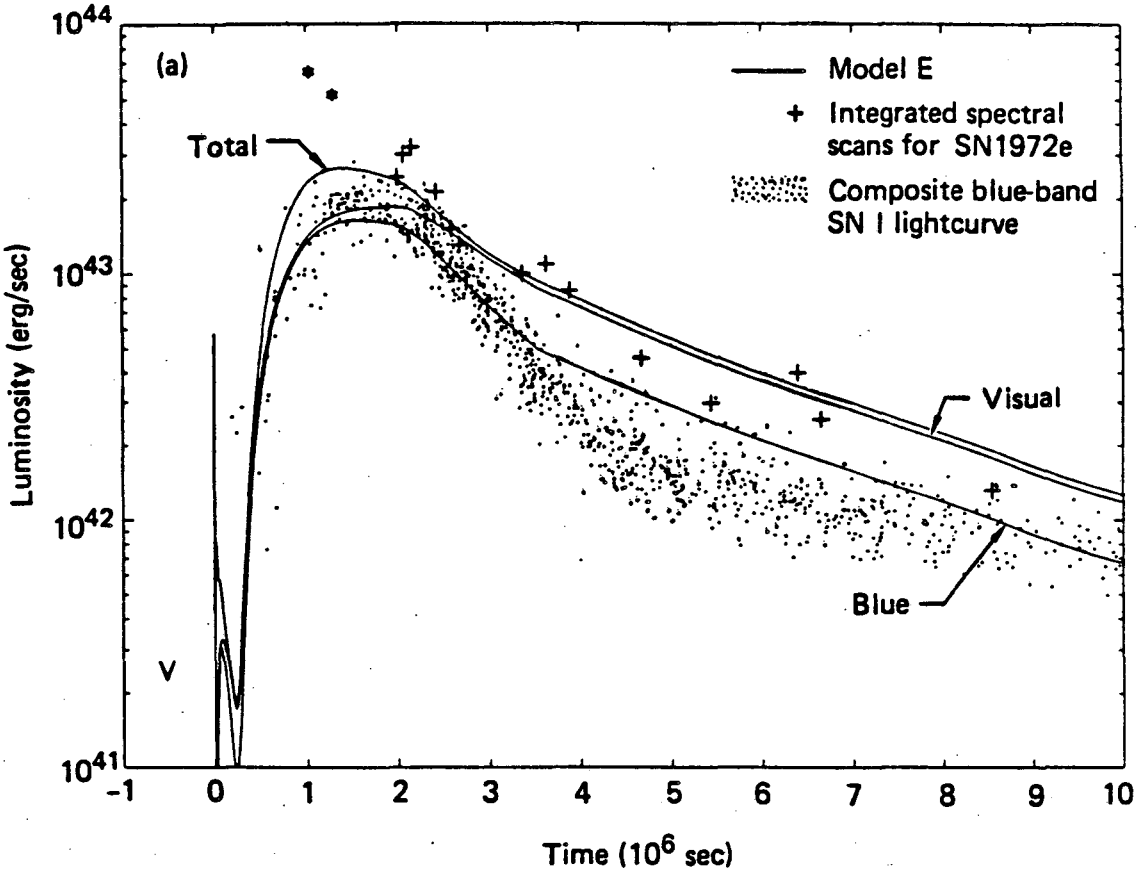

Figure 1.4

Computer-generated light curve for a type I supernova model (isolated massive star) showing early time spike [24].

Type I light curves -- As shown in figure 1.5 [26], some observers have divided type I supernovae into two subclasses, "fast" and "slow". It is still not known whether this is a real distinction, or simply an arbitrary division in a continuum. Because the Berkeley search should rapidly collect a great deal of data on type I light curves, especially at early times, it may be possible to resolve this controversy, or at least provide a better data base, relatively quickly.

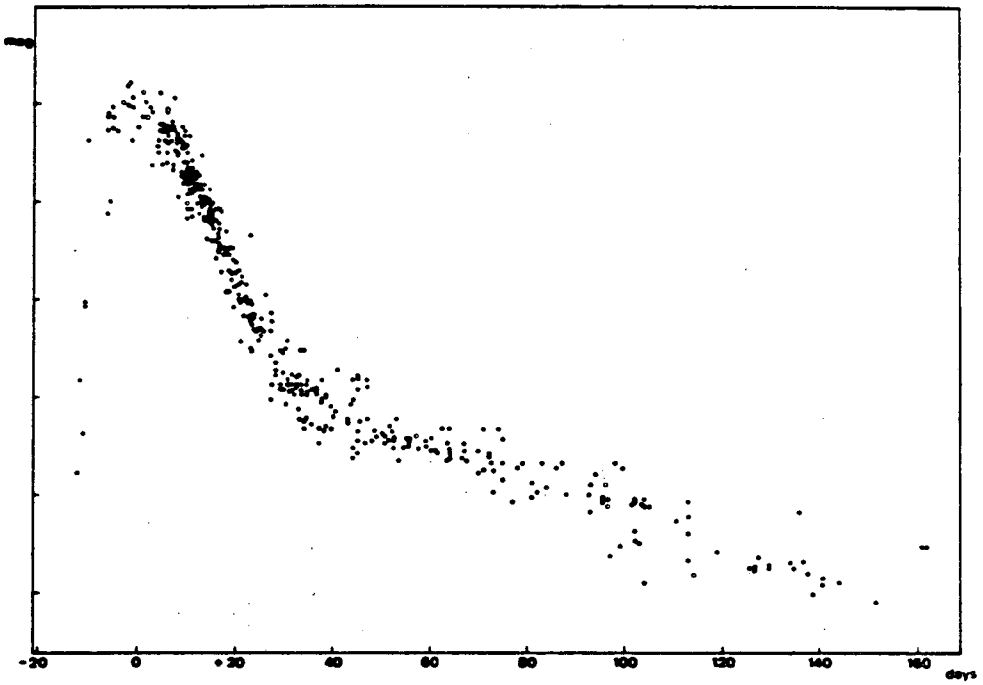

Fig. 2. Composite light curve of 11 'fast' SN-I.

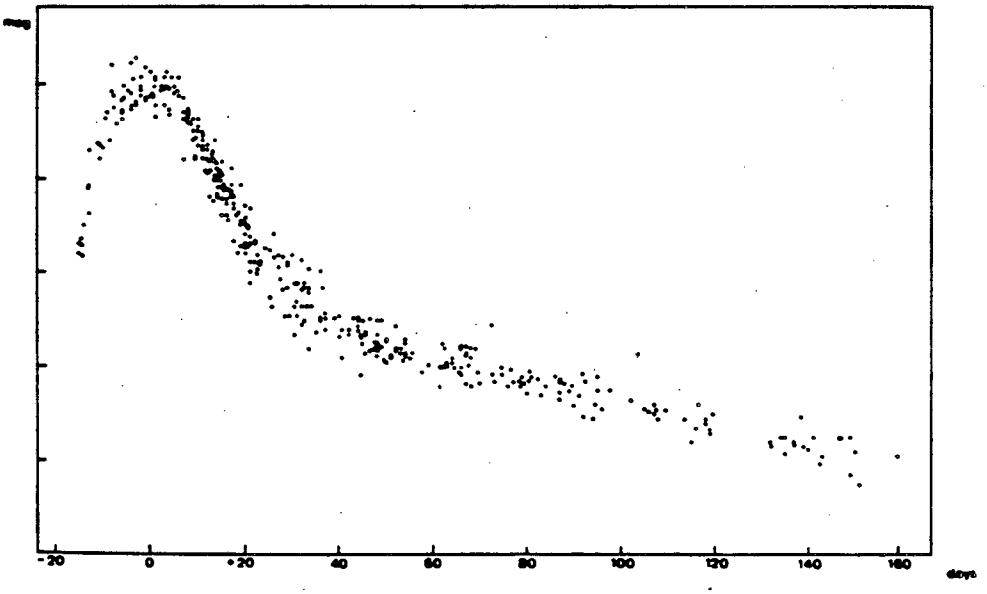

Fig. 3. Composite light curve of 15 'slow' SN-I.

Figure 1.5

Evidence for the existence of two varieties of Type I supernovae?

2. Design of the Berkeley System -- Hardware

2.1. Detector

2.1.1. CCD

A charge coupled device (CCD) transfers packets of electrical charge from place to place on a semiconductor using electric fields produced by insulated electrodes; one- and two-dimensional CCD arrays have been developed as analog and digital memories and various kinds of signal processors, as well as imaging detectors [27]. CCD detectors are desirable for astronomical purposes because of their high quantum efficiencies, low noise, inherent stability, and high resolution.

Our CCD is an RCA device (SID53612-X0) with 320 x 512 resolution elements (pixels). This device is thinned and back illuminated for improved quantum efficiency. Each pixel is 30 microns square, so the sensitive area is 0.96 cm x 1.54 cm. This is a buried channel CCD (the charge-collecting regions are below the silicon surface) with a transparent electrode structure, so the entire surface area of the chip is photosensitive. The only other astronomical-grade CCD commercially available at the start of this project, one made by Fairchild, had only 30% sensitive area, with the remaining surface devoted to shift register structures; this 3 to 1 difference in collecting area (effectively in quantum efficiency) led us to select the RCA chip.

2.1.2. CCD Electronics

The CCD driver electronics constructed for the Berkeley system are based on designs by by J. Geary of the Harvard-Smithsonian Observatory [28]. Figure 2.1 shows a block diagram of the circuit; details of both the CCD and the driver circuits are given in the senior honors thesis of J. Culver [29], which is attached to this paper as appendix 1. Three - phase timing pulses are derived from a master crystal clock

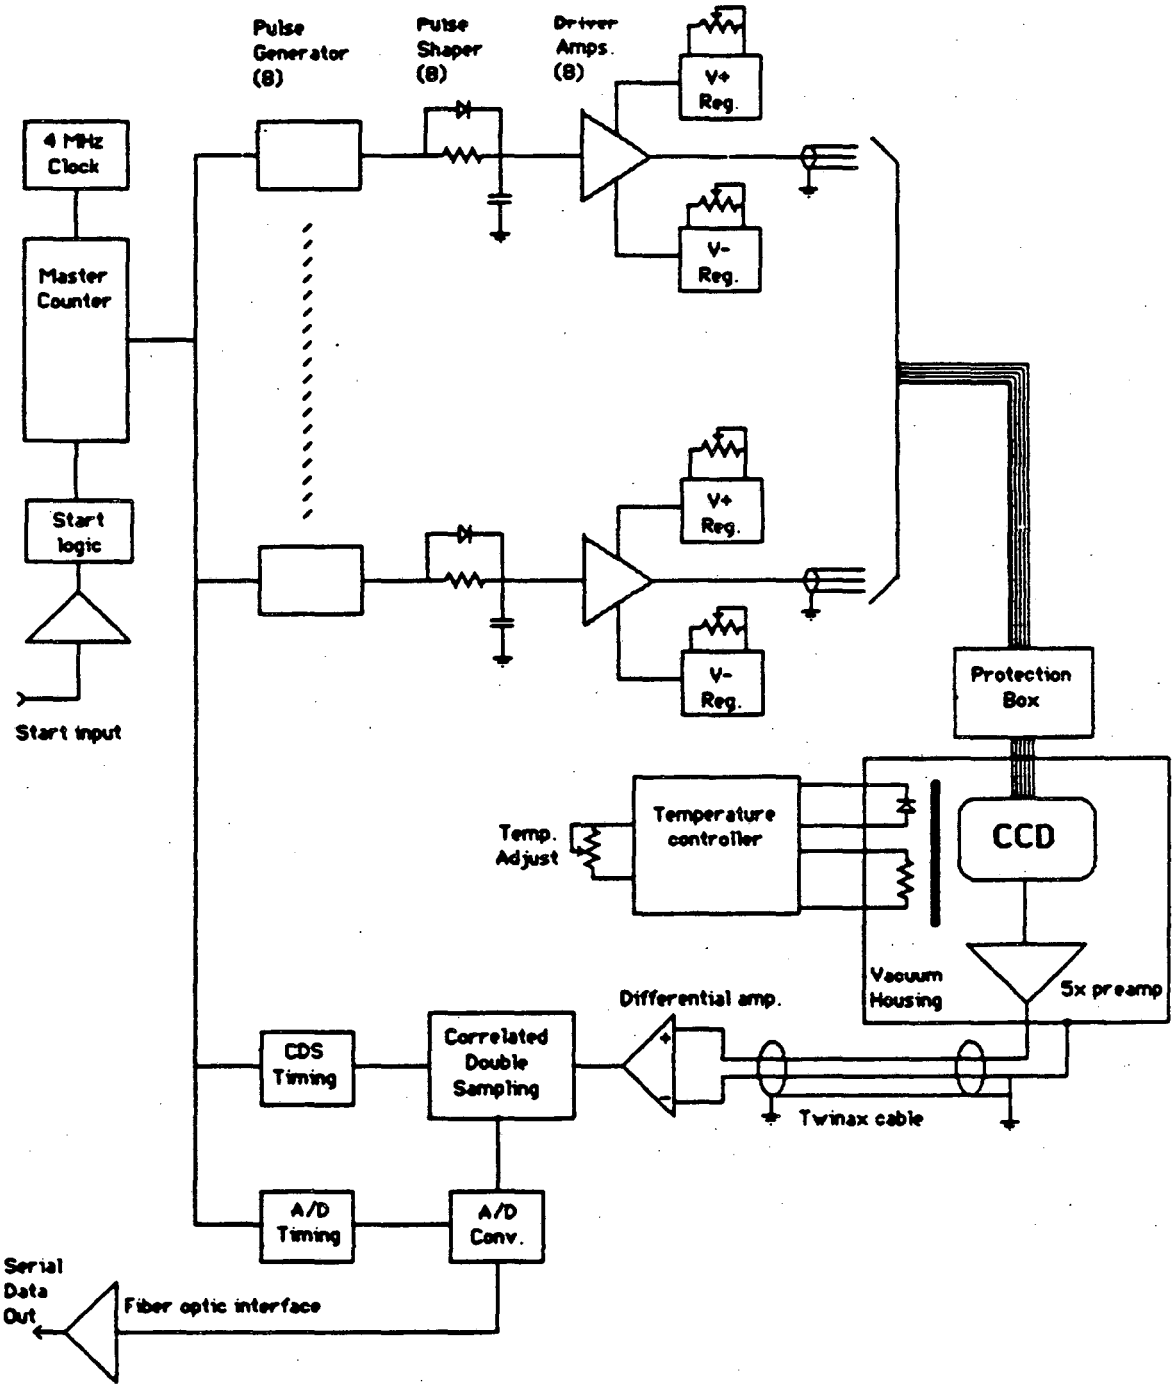

Figure 2.1 -- Drive Electronics for RCA CCD Detector

and counter. These are converted to drive pulses with adjustable voltage levels and controlled fall times by buffer amplifiers. Signals to the CCD pass through a protection box containing fuses and Zener diodes; this minimizes the chance of applying destructive voltages to the chip. Additional resistors inside the CCD housing drain off any static buildup when the drive electronics are not connected.

The output voltage from the on-chip FET amplifier is buffered by the preamplifier inside the CCD housing, then sent as a differential signal to the input amplifier of the correlated double sampling (CDS) board. The differential circuit, using twin-ax cable, picked up substantially less 60 Hz and higher frequency interference than the original single-ended circuit transmitting over 1 meter of coaxial cable. On the CDS board, the output voltage is sampled immediately after the CCD output capacitor is reset, and again after the next pixel is transferred to the output; the first sample is subtracted from the second to cancel out the thermal noise charge $\sqrt{CkT}$ stored on the output capacitor. The difference signal, proportional to the charge stored in the pixel, is digitized to 14 bit accuracy and transmitted as a serial bit stream to the computer or videotape recorder interface. Our normal readout rate is 1 pixel every 50 microseconds, or 8 seconds for the full CCD. The maximum readout rate with our current A/D converter is 1 pixel every 20 microseconds, or 3.2 seconds for the full CCD.

We have constructed a separate CDS amplifier box which mounts directly on the CCD housing and receives clock signals via fiber optics, to further reduce noise pickup and eliminate ground loops between the clock and output cables, but have not yet installed it.

2.1.3. Measurement of CCD Properties

Details of the CCD measurements will be included in a forthcoming thesis by M. Shane Burns. The following is a summary of the properties we have measured to date.

2.1.3.1. Noise and Gain

We have measured the noise and gain of our CCD using the standard technique [30] of measuring the variance of the output signal at different levels of illumination. The noise level of our current device is 88±10 photoelectrons/pixel, and the gain is roughly 28 photoelectrons per A/D converter count. Transfer inefficiency in the horizontal direction is $\epsilon=2\pm1\times10^{-4}$, and in the vertical direction $\epsilon=3\pm1\times10^{-4}$, 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, and our normal operating temperature is between -120 C and -110 C.

There is a significant "fat zero", or minimum filling charge in each pixel required for efficient charge transfer. We estimate the fat zero to be 2000 photoelectrons, or about 70 A/D counts. In normal operation at Leuschner, sufficient sky background is collected to exceed this fat zero level in 6 to 7 seconds, thus presenting no problems with our current 15 to 30 second exposures.

We have found that the sensitivity of our CCD to external noise pickup is high, and is critically dependent on the adjustment of the positive voltage of the ϕ-reset pulse. We get significant (5-10 A/D counts) 120 Hz interference if this voltage is adjusted as little as 0.1 volt away from its nominal level of 0.9 volts.

Our current device is a "loaner", with specifications not guaranteed by RCA. A replacement chip in a dual in line (DIP) package with an improved, lower noise

output amplifier has been delivered by RCA, but has not been tested at this writing.

2.1.3.2. Quantum Efficiency

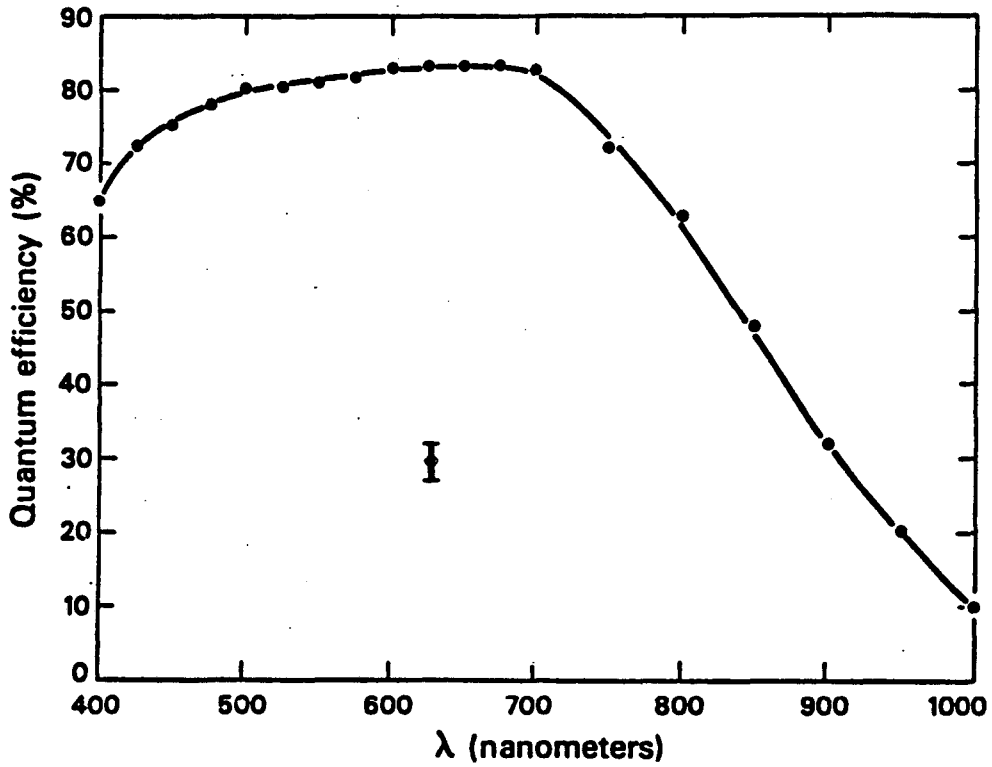

Figure 2.2 -- RCA CCD Quantum efficiency vs. wavelength
Curve -- manufacturer's data
\+ -- measured value at 633 nm

The manufacturer's typical quantum efficiency curve for our CCD is shown in figure 2.2. 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 we obtain is 30 $\pm$ 3 per cent, or a factor of 2.5 lower than the published curve. We do not as yet have any explanation for this low efficiency; however, this is typical efficiency for a standard "thick" CCD. We have not measured the spectral response directly. At a wavelength of 500 nm, the overall system response is a factor of 2 to 2.5 below our expectations, and some of this loss may be due to even lower CCD quantum efficiency at this wavelength.

The pixel to pixel variations in quantum efficiency are small. Figure 2.3 shows the relative efficiencies for a block of pixels near the center of the CCD; these values are derived from an exposure of the twilight sky, and scaled so that the mean over the whole CCD is 1000. The random noise level is less than one count. We do not require very accurately flat fields, so we have not tried to measure or correct for any color dependence of these quantum efficiency variations.

		Column									
		160	161	162	163	164	165	166	167	168	169
Row	250:	820	816	820	816	819	822	826	826	830	830
	251:	820	823	825	822	822	821	825	826	828	826
	252:	825	825	822	826	823	829	827	833	826	836
	253:	826	829	830	828	829	827	829	830	834	836
	254:	836	829	827	829	830	831	829	831	837	836
	255:	831	831	830	827	830	834	831	831	836	833
	256:	837	833	832	836	838	835	833	835	840	841
	257:	837	830	834	835	832	834	838	840	837	840
	258:	836	843	836	833	835	844	834	839	840	840
	259:	839	837	838	837	835	838	837	839	843	841
	Average:	831.43		Std. Dev.:		6.300					

Figure 2.3

1/(quantum efficiency), in arbitrary units, for 10 x 10 pixels. Pixel to pixel variations average 0.75% of the local mean.

There are several large-scale nonuniformities in the CCD response. Curved streaks of high sensitivity and a "dimple" of low sensitivity at the bottom center of the CCD are visible in figures later in this thesis. The dimple is visible as a mechanical defect in the CCD chip. A thin strip at the top of the CCD has very low response, some 30% below the CCD mean; this is the area farthest from the output

horizontal register. Variable bright patches at the bottom of the CCD may be due to light generation (electroluminescense) in the horizontal register [31]. There is only one known bad pixel, in the upper left of the CCD, which has a significant dark current and produces a "flare" on images read out after long exposures or delays. Several columns in the left portion of the picture are "cold" by a few A/D counts, the amount varying with temperature, but the effect is small enough to have little effect on bright foreground stars, and is fortunately in the wrong direction to produce false alarms.

2.1.4. CCD Housing and Dewar

Astronomical CCD's are operated cooled to approximately 150 - 200 K (-70 to -120 C) to reduce the dark current (signal proportional to time between readouts) to essentially zero. Our CCD is cooled with liquid nitrogen, using a feedback circuit driving resistive heaters to maintain a constant temperature. A vacuum housing insulates the CCD and minimizes condensation. The construction of the CCD housing and liquid nitrogen dewar is sketched in figure 2.4.

Using a standard 2-stage mechanical pump, the vacuum system reaches pressures of approximately 18 millitorr, with a useful hold time of several days. Our liquid nitrogen dewar is a large consumer thermos bottle (pump-pot), available from K-mart. It has a capacity of slightly under 1 liter of LN_2, and a hold time of approximately 15 hours. A cryopump consisting of about 3 cm^3 of Linde Molecular Sieve (types 4A and 5A) in a stainless steel tube is attached to the vacuum housing with an o-ring seal. The cryopump tube was originally epoxied into place, but the pump material must be replaced occasionally, and we had trouble maintaining a leak tight joint.

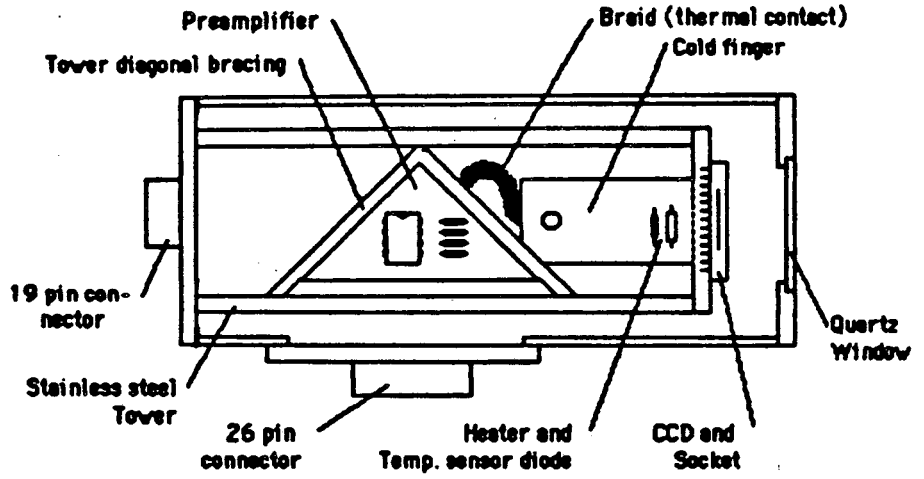

CCD Detector head
Side view

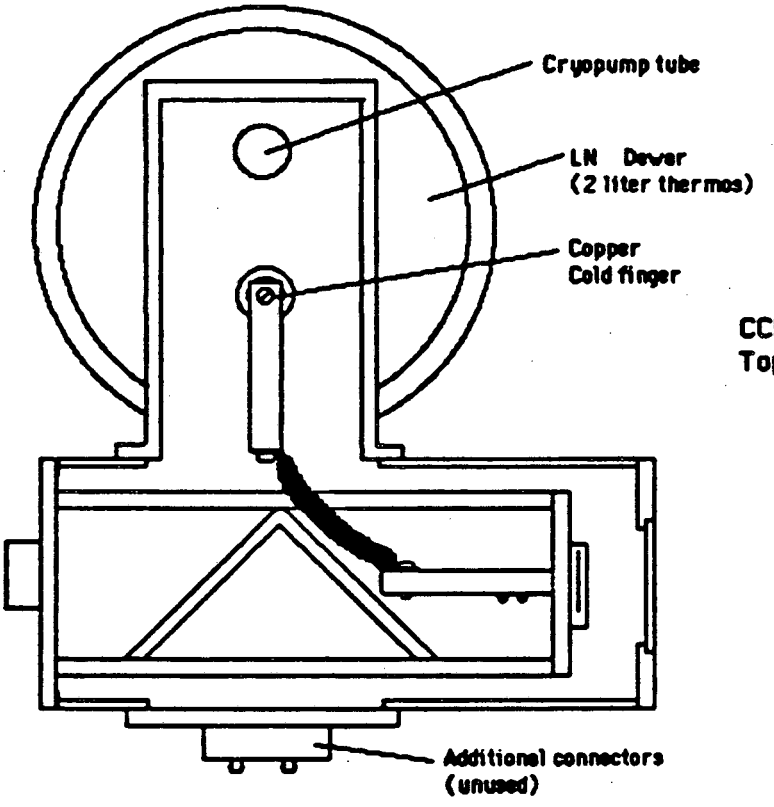

CCD Housing and Dewar
Top view

Figure 2.4 -- CCD vacuum housing and Dewar assembly

When operating in high humidity, particularly at the MIRA telescope site in Cachagua valley, we found that the quartz window of the CCD housing would fog up frequently. We now flow dry nitrogen through the area between the window and the shutter mechanism. 0.5 cubic feet of nitrogen per hour prevents fogging at any time.

2.2. Optics

2.2.1. Telescope

The telescope we are using now is a 30" (75 cm) f/8 Ritchey-Chretien reflector, on an equatorial fork mount, located at the Leuschner Observatory of the University of California in Lafayette, California.

Another telescope, used for some testing and possibly 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). This telescope has recently been moved to the MIRA Oliver Observing Station, on Chew's Ridge near Monterey, California. The MIRA telescope provides larger collecting area and faster pointing, but is less conveniently accessible for development work and until recently was located at a very poor observing site.

2.2.2. Reducing Optics

2.2.2.1. Optimum Scale

Because of the small physical size of the CCD, we must use additional optics to obtain an appropriate "CCD scale" of approximately 3 arc seconds per pixel. If the pixel scale is too small (e.g. 1.1 arc sec/pixel at the Leuschner focal plane), the light from a supernova is spread over many pixels, reducing the signal to (CCD readout) noise ratio. Also, more pixels must be processed to cover a given galaxy, thus increas-

ing computer requirements.

For pixels much smaller than one seeing disk, the supernova signal increases as the pixel area (square of the CCD scale). The sky background signal is also proportional to the area, so the sky background noise increases as the square root of the area, or linearly with the CCD scale. The signal to noise ratio thus goes as the CCD scale squared while CCD readout noise dominates, and linearly with the CCD scale when sky background noise dominates. When the pixels are much larger than one seeing disk, however, the signal to noise ratio will at best be constant, and will decrease with increasing CCD scale if the sky noise exceeds the CCD noise. The optimum CCD scale thus depends on several factors (CCD noise, sky brightness, typical seeing, and exposure time); however, pixels equal to or slightly larger than the typical seeing disk are generally near optimum.

Some information is lost when an image is recorded with pixels larger than approximately half the seeing disk (sampling rate less than twice the highest spatial frequency present). This loss of information results in imperfect subtractions around foreground stars and galaxy cores. However, as will be discussed in section 3.2, there are several ways to minimize these errors. We have no convenient way to accomodate different pixel sizes for search and reference images. We choose to favor large pixels, trading localized subtraction errors for increased signal, for the current search. As lower noise CCD's and/or increased computer power become available, we may switch to smaller pixels.

The maximum possible pixel scale is set by optical considerations; Liouville's Theorem (conservation of phase space) requires that the product of aperture and solid angle be conserved.

$$D_{telescope}\,\theta_{pixel} = l_{pixel}\,\Theta_{lens}$$

or, for small angles

$$D_{telescope}\,\theta_{pixel} = \frac{l_{pixel}}{f/no}$$

The maximum possible value of Θ is 2π (f/0.5 lens); for a 30 inch telescope (D = 0.75 meter) and a 30 micron square pixel, the maximum possible CCD scale is 2.5×10^{-4} radians, or 52 arc sec, square. For practical lenses, a limiting speed of f/2.0 yields approximately 4.2 arc sec per pixel (3.5 arc sec at the MIRA telescope). We actually use lenses as fast as f/1.4 (Nikkor 50 mm), but the effective f/number is limited by the demagnification ratio used and the telescope f/number to about f/2.

2.2.2.2. Installed Optics

The reducing optics which we have used are shown in figure 2.5. Figure 2.5a shows the original scheme, used for tests at the MIRA telescope. The field lens images the telescope aperture onto the reducing Nikon (Nikkor) lens. The MIRA focal length of 9 meters gives a telescope focal plane scale of 23 arc sec/mm (0.7 arc sec/pixel). Because of mechanical constraints on the size of our instrument, we used a very short focus (24 mm) f/2 lens, a standard lens for 35 mm cameras. Note that the CCD is somewhat smaller than a 35 mm film frame, so we expect good image quality over the entire CCD from such camera lenses. This gave an effective f/number of 2.2 and a demagnification ratio of 4.5, for a CCD scale of 3.1 arc sec/pixel. At Leuschner, the shorter telescope focal length gave a CCD scale of 4 arc sec/pixel with the same lens (demagnification ratio of 3.6), but we found some vignetting (loss of light transmission near the edges of the field).

To reduce the CCD scale and minimize vignetting, we switched to the optics in figure 2.5b. Our 24 mm lens was lost or stolen, and with no length restrictions at Leuschner, we chose a 50 mm f/1.4 replacement lens. The added achromat lens permits closer focusing of this lens, at the expense of increased curvature of field and

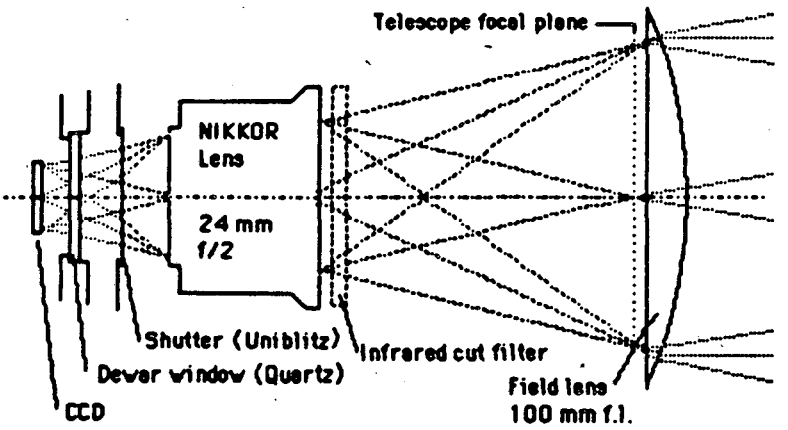

Figure 2.5

Reducing Optics

A. Original (MIRA telescope) Optics

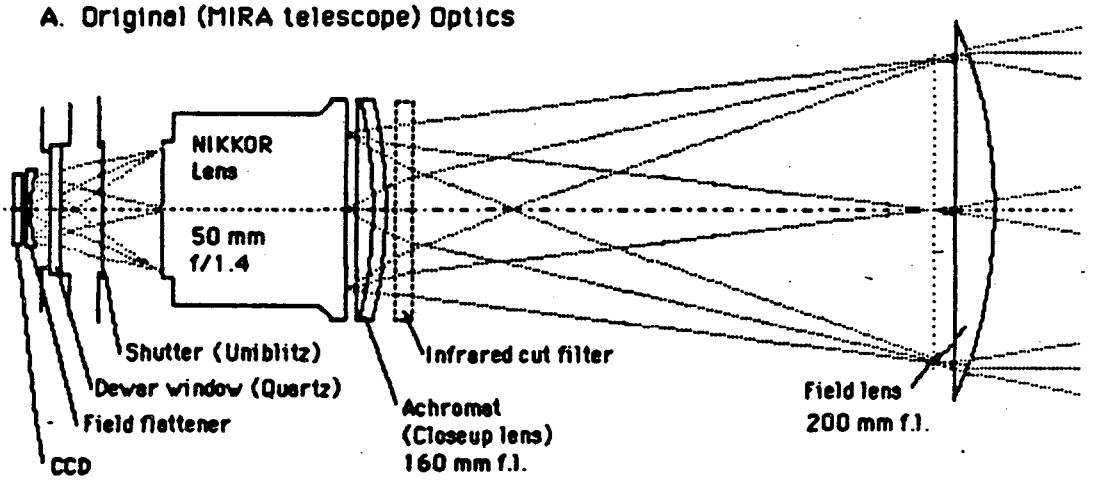

B. Leuschner Telescope

C. Leuschner Telescope -- Modified for reduced distortion

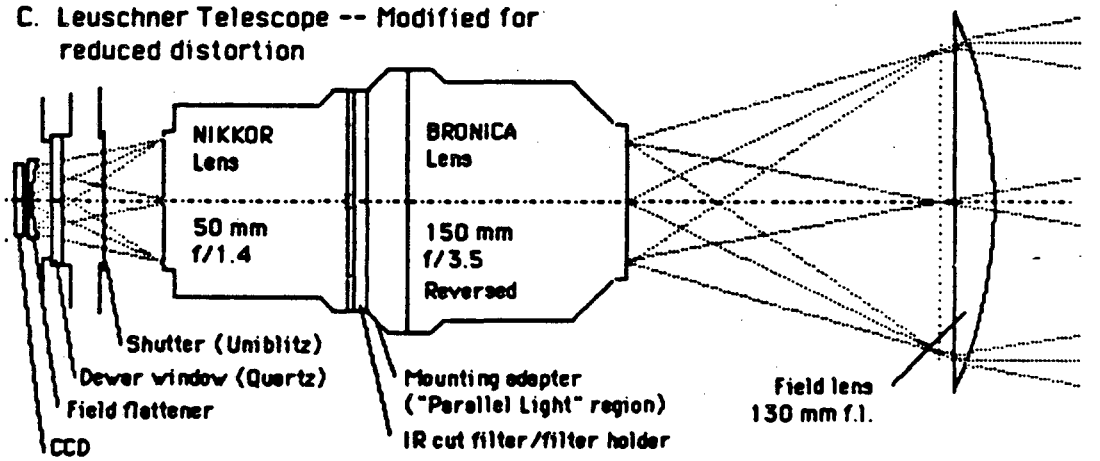

distortion; non-achromat closeup lenses were tried, but introduced unacceptable chromatic aberration. The field flattener (not related to the software process of field flattening) is mounted inside the CCD vacuum housing and partially compensates for the field curvature. This system provided a CCD scale of 2.8 arc sec/pixel at the Leuschner telescope, with minimal vignetting but severe distortion. The measured pincushion distortion (of the form $\delta r = \epsilon r^3$) is approximately $\epsilon = 1.8 \times 10^{-7} \frac{pixel}{pixel^3}$. This distortion shifts images by over 1.4 pixels, 200 pixels from the optical center of the CCD; this is enough to make accurate alignment of images difficult without very precise telescope pointing.

The current optics are diagrammed in figure 2.5c. The achromat closeup lens is replaced by a Bronica medium format camera lens, run "backwards" to transform light from the telescope focal plane to parallel light; the Nikkor lens focuses this light onto the CCD. A larger format lens is needed because the field of interest in the telescope focal plane is several centimeters across, larger than a 35 mm camera lens is designed to accomodate. The resulting CCD scale is 2.4 arc sec/pixel, with very good image quality and flatness of field, and reduced pincushion distortion ($\epsilon = 6.5 \times 10^{-8}$). However, the Bronica lens contains internal stops which produce substantial vignetting at the edges of the CCD. The Bronica also has relatively poor transmission. It is thus likely that these optics will eventually be replaced again. Figure 2.6 shows the vignetting produced by the current optics; the top and bottom of the CCD are lowered in sensitivity by about 30%, and the extreme corners by as much as 50%.

The transmission through various optical elements is given in table 2.1. Note that roughly half of our available light is lost in the reducing lens system; this results from our compromise among image quality, light loss, space constraints, and use of inexpensive standard lenses.

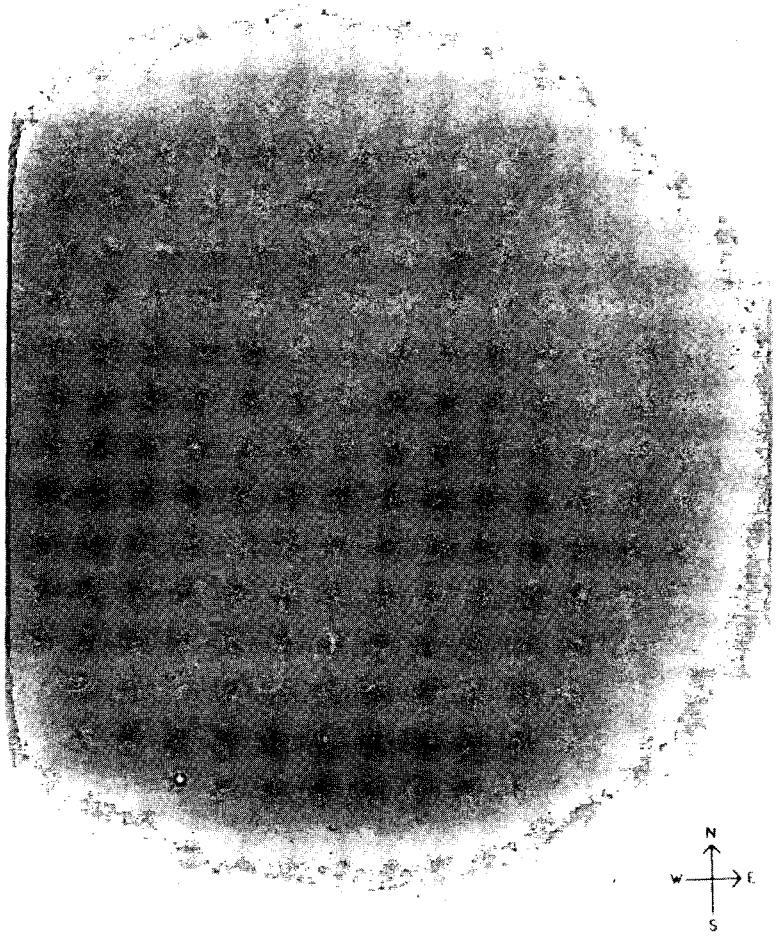

Figure 2.6 -- Reducing Lens Vignetting

0.5 second exposure of twilight sky. Display range 5000 A/D counts (white) to 10,000 A/D counts (black).

Table 2.1 -- Transmission of Various Optical Elements

Element	Transmission	Cumulative Transmission
Star at top of atmosphere	--	1.0
Atmospheric absorption	0.85(a)	0.85 at zenith
Telescope		
Aperture (obstruction)	0.93	0.79
Primary reflectivity (b)	0.90(c)	0.71
Secondary reflectivity (b)	0.90(c)	0.64
Optics		
Field lens	0.92(d)	0.59
Bronica lens (on axis)	0.60(c)	0.35
Nikon lens	0.85(c)	0.30
Dewar window	0.92(d)	0.27
(IR filter)	0.92(d)	0.25
(Field flattener)	0.92(d)	0.23

Notes:

(a) Measured by observation at various zenith angles

(b) Primary washed, secondary realuminized

(c) Measured using He-Ne laser (633 nm) and calibrated photodiode

(d) Assuming 4% reflection per surface

2.3. Computer System

2.3.1. Processor

The computer system used for image processing and system control for the supernova search system is a PDP-11/44 with floating point processor hardware and (currently) 2.5 megabytes of main memory. The 11/44 is a compact, high speed 16 bit minicomputer which provides convenient real-time control and interfacing via the DEC UNIBUS. At the time the search system was designed, it was the smallest DEC computer providing sufficient memory space (4 megabyte maximum) and processor speed for the search -- although at the time we anticipated using no more than one megabyte of memory. In view of the problems we have encountered in using a 16 bit computer, we would today make every effort to purchace a VAX 11/750 or similar 32 bit machine for any comparable project.

2.3.2. Peripherals

As currently configured, the PDP-11 system has nine serial ports available (including the system terminal port), all in use -- one for the system terminal/printer, five for user terminals, two for modems used to communicate with Leuschner observatory, and one for control and monitoring of the videotape data receiver. A 160 megabyte Fujitsu sealed (Winchester) disk drive provides mass storage, with two RL02 ten megabyte removable cartridge drives for backup, testing, and auxiliary storage. Graphics display is provided by a Grinnell GR-270 display system, with 512 x 512 x 8 bit display memory. Graphics hard copy is provided by a Tektronix dry silver photo printer, model 4634, connected to the video signal from the Grinnell. We have ordered a nine track magnetic tape drive (Kennedy model 9100 with Emulex controller), but have not yet received it.

2.3.3. CCD Interface

CCD data, either direct from the CCD or recovered from videotape, is read into the computer system via a direct memory access (DMA) interface. The main interface is a standard DR-11B single board DMA interface from MDB, Inc. 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 during times when the computer bus is not available for DMA.

2.4. Telescope Control System

2.4.1. Leuschner 30" Reflector

The Leuschner 30" telescope has been equipped with a fully automatic control system over the last two years, largely through the efforts of Richard Treffers. The telescope position is sensed by absolute position 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 256 kilobyte memory and dual floppy disk drives; however, the telescope has operated automatically with a 24 kilobyte Commodore PET computer with one disk drive. Both IBM and PET computers are programmed in BASIC. The IBM computer uses a compiled BASIC which is fast enough to handle all control functions, while the PET requires an assembly language routine to sample the telescope status and control the drive motors.

The control computer operates the telescope slew and guide motors via simple D/A converter interfaces; a separate synchronous motor drive is mechanically summed with the right ascention guide drive for tracking. A simple three-step procedure is used to point the telescope: full speed slew, 1/4 speed slew, and guide to within specified limits of the requested pointing. Minor modifications (adding damping

resistors across the slew motors, and changing set points in software) have reduced the typical pointing time from over a minute to approximately 30 seconds. We expect to reduce this time to under 10 seconds for a five degree motion by replacing the guide and tracking motors on each axis with single electronically controlled torque motors.

Pointing repeatability is approximately ten arc seconds. Refraction, flexure, and polar axis errors result in a pointing accuracy on the sky of approximately two arc minutes. We are now measuring coefficients for an existing pointing correction program, and we expect to obtain 10 - 15 arc second accuracy on the sky shortly.

The Leuschner control computer communicates with the PDP-11/44 computer via an RS-232 serial port and dial up modem at up to 1200 baud. A second 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.

2.4.2. MIRA Telescope Control

The MIRA telescope control system is shown schematically in figure 2.7. Incremental shaft encoders are coupled via rollers to both axes, giving a readout resolution of 0.1 arc second. The RA drive is a roller drive, and both axes are designed for high speed pointing. The anticipated pointing performance is 5-10 arc sec accuracy in three seconds, for a two degree motion. Note that since we will eventually be measuring the telescope pointing error on each CCD field in real time, we can correct any cumulative pointing errors. However, until we begin real time processing, the absolute encoders used at Leuschner provide more reliable pointing than MIRA's incremental encoders. MIRA's control computer is a 16 bit HEX minicomputer using

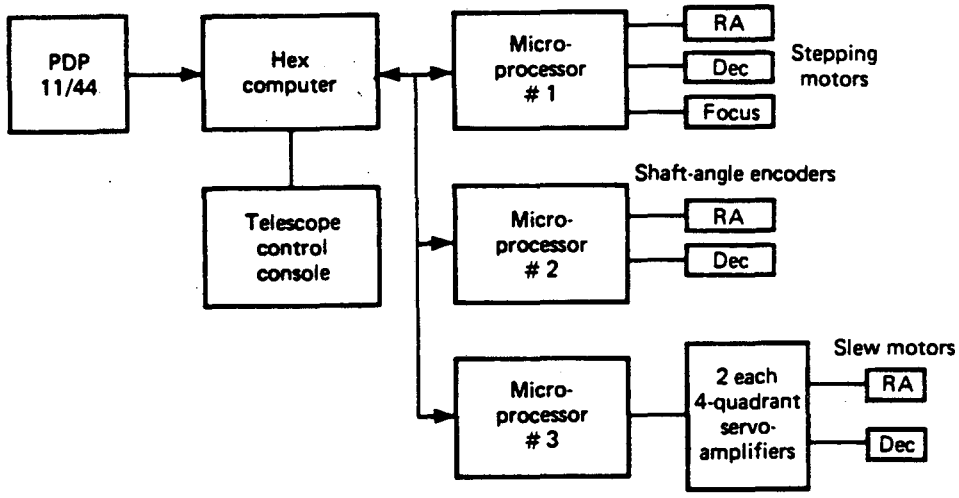

Figure 2.7 — MIRA telescope control system

bit slice technology; it communicates with individual Z-80 control microcomputers ("trolls") via RS-232 serial links. This control system is not yet fully operational.

2.5. Special Hardware

2.5.1. Videotape Storage Unit

The output of the CCD electronics package is a serial data stream, with a data rate of 20,000 14 bit words per second for the current eight second readout time. Until the computer is installed at the telescope site, we must store this information for later processing. Even after we can read out directly into the computer, it will be desirable to store the raw CCD data for archival purposes.

Because of the high data rate and very large volume of data (320 megabytes for 1000 CCD fields) the only practical storage medium is magnetic tape. Conventional computer tape drives (1600 bpi) typically store only 40 megabytes per reel (2400 feet) of tape; they are also large, relatively expensive, and complex to control. We therefore chose to build an interface to a consumer videotape recorder. The theoretical capacity of such a recorder is very large -- several million bits per second for three hours, or several billion bytes per tape. As constructed, our recorder 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 megabytes per three hour tape. The recorder used is a Sony portable Betamax, model SL-2000. Beta format is preferred over VHS because the Beta recorder retains tape loading during fast forward and rewind operation; one can thus shuttle quickly between frames. VHS tape is unloaded from the tape head drum during rewind, a process requiring several seconds for each operation.

The major problem with videotape storage is a relatively high burst error rate; the signal to noise ratio of videotape is only some 40 dB, and dropouts (visible in ordinary use as horizontal streaks on the video screen) are common. Our videotape interface records each pixel twice to reduce errors; the single pixel error rate with this redundancy is a few times 10^{-5} (between zero and about 20 pixels received incorrectly or completely lost in each 160,000 pixel image). No error correction is done; if a pixel is not received correctly in two attempts it is discarded. This level of error is acceptable for the supernova search as long as no pixels are skipped completely. Skipped pixels cause the entire remainder of the CCD field to be shifted horizontally; vertical lines acquire a "jog", and star coordinates are shifted.

To prevent this, a two bit count is recorded with each pixel. A modulo 4 counter keeps track of the number of pixels recovered from the videotape; if this

counter disagrees with the recorded count, dummy pixels are sent to the computer to fill the gap. A typical image without this "jog corrector" circuit is shown in figure 2.8a; the jagged white line near the right (top) edge of the image is the column of "cold" (low value) pixels at the left edge of the normal CCD image, shifted out of position by the missing pixels. Note that the CCD scans from the lower left to the upper right, so the displacement increases as more pixels are lost. The same image with the jog corrector on is shown in figure 2.8b.

We keep track of CCD pictures by writing ASCII data onto 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 control computer operates the videotape recorder via a simple interface which allows the computer to "push buttons" on the remote control by sending characters to a serial port.

At this writing, the videotape interface has been operating for some two years, with only occasional problems due primarily to dirty videotape recorder heads. The total amount of data stored is approaching 100 gigabytes, or the equivalent of many hundreds of reels of computer tape. Schematics and further details on the videotape system are in the videotape system manual, attached to this thesis as appendix 2.

2.5.2. Scan Converter

A simple scan converter (video frame store) is used to display a portion of the CCD output directly on a video monitor at the telescope site. The scan converter contains 16K bytes of memory, sufficient to store a 128 x 128 x 8 bit image. The left-most 256 columns of the CCD are divided into eight 128 x 128 pixel regions, and the scan converter can store any single region, or every eighth pixel (every second column by every fourth row) of the entire 256 columns. Any eight of the 14 bits produced by

Figure 2.8a -- Retrieved image with jog corrector off. Note progressive shift of white line (actually located at the edge of the CCD) as pixels are lost.

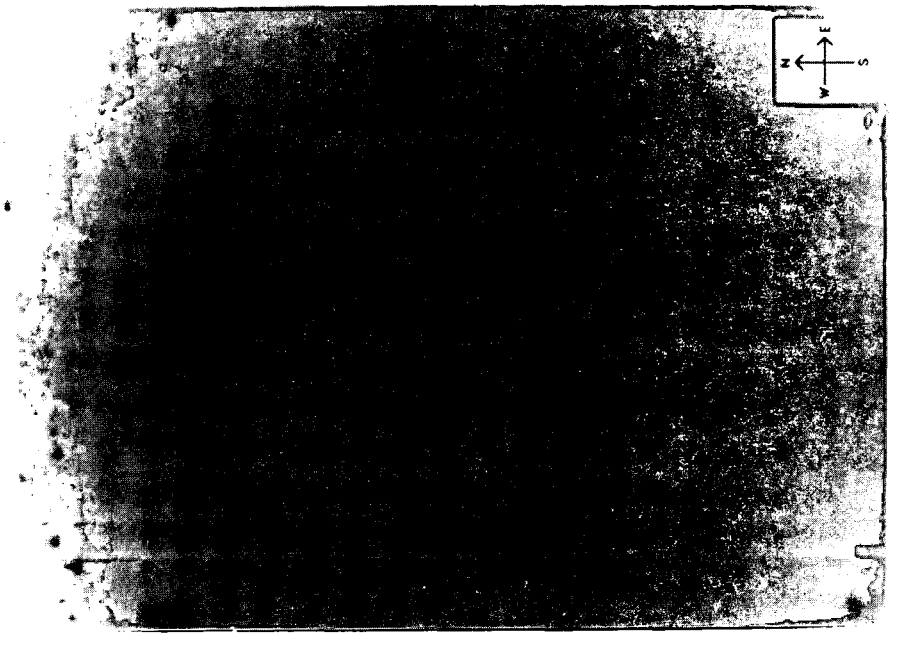

Figure 2.8b -- Same image, retrieved with jog corrector on. Scattered white specks are single dummy pixels, inserted by jog corrector to keep image aligned.

Figure 2.8 -- Operation of Jog Corrector in videotape recorder system

the CCD A/D converter can be stored.

Once an image is stored, it is displayed on a local monitor. Any five of the eight stored bits can be displayed, providing a contrast control. Binary switches can be set to subtract off any value from the top four bits, allowing crude offset and background subtraction.

The entire scan converter occupies two 3" x 8" circuit boards inside the video-tape interface transmitter box. It has proved invaluable for testing and operating the CCD. A larger scan converter with 256 Kbytes of memory has been partly designed, but has not so far proved necessary.

2.5.3. Fiber Optics

In early bench tests of the CCD detector, we observed irregular spikes in the CCD output, with a frequency of approximately 20 KHz. These were traced to radiation from the switching power supplies of the PDP-11/44 computer, conducted via the CCD interface cable to the CCD electronics package. To eliminate this interference, we constructed a fiber optics link, providing 300k bit per second transmission of CCD data to the computer, plus a low speed link multiplexing up to eight signals in each direction over two fibers.

The first version of this link used Fairchild OPTO demonstration kits, with plastic-coated glass core fiber and AMP OPTIMATE connectors. These kits provide 5 MHz bandwidth over a maximum distance of approximately 100 meters. These worked satisfactorily for about a year, but gradually became intermittent, apparently due to connector deterioration. A replacement system, using Hewlett Packard plastic fiber cable and snap-in connectors, has proved more convenient and has so far worked without difficulty. The relatively high attenuation of the plastic fiber (200 dB/km)

limits the fiber length to approximately 15 meters, but this has not been a problem to date.

In addition to reducing noise pickup, the fiber optics link provides extra protection against destruction of electronics due to lightning; conventional cables acting as antennas can pick up destructive voltages even from nearby strikes.

3. Design of the Berkeley System -- Software

3.1. Computer System

3.1.1. Operating System

The supernova search PDP-11/44 runs under the RSX-11M operating system, version 3.2. This is a real time operating system with extensive scheduling and interrupt facilities, and is well suited to a dedicated computer which will eventually be controlling a telescope and processing images in real time. The major limitation of RSX-11 is that individual programs (tasks) are limited to a 16 bit (64 kilobyte) address space. Large data regions can be accessed through memory mapping registers, which allow access to the full 22 bit (4 Mbyte) address space of the processor. However, programs are limited to 64 kilobytes for code and variables. We work around this restriction by sequentially spawning multiple tasks, each performing a single function; the cost is a substantial overhead in code and memory space and a modest overhead in execution time. Overlays, which allow different program subsections to share the same task address space, also allow large programs to be run, but are difficult to construct efficiently for FORTRAN programs, and have been found to lead to subtle run-time errors.

3.1.2. Languages

The bulk of the supernova search software is written in DEC FORTRAN IV-Plus. This is a standard language, lacking many modern features such as structured control flow, but providing full access to the features of RSX-11M.

We considered several other languages and operating systems, including UNIX/C, PASCAL, and FORTH, for our purposes. UNIX is not designed for real time control purposes, and we were advised at the time we selected our operating

system that the standard Berkeley UNIX would require substantial additions to handle our application. Available C and PASCAL packages were not guaranteed to properly use the facilities of RSX-11. FORTH, a language developed originally for telescope control, would have been a reasonable choice, but none of the members of the supernova search project were sufficiently familiar with it to consider using it for a large project.

A small amount of code has been written in MACRO-11, the PDP-11 assembly language, for device drivers for the CCD and Grinnell graphics display interfaces. Additional assembly code may be written eventually to speed up inner loops in the supernova image processing, in the interest of full real time processing.

3.2. Image Cycle

3.2.1. Introduction

The core of the real time image processing is the image cycle. The image cycle routines accept one CCD image of a galaxy as input, and produce a list of coordinates of "supernova candidates" -- CCD pixels brighter than can be accounted for by noise or known sources of false alarms -- as output. At the highest level, the image cycle routines select galaxies to be examined (from an input list), and determine which supernova candidates are sufficiently well confirmed to warrant alerting a human operator.

The major functions of the image cycle are:

1. Acquisition of new image
2. Retrieval of reference information
3. Matching and subtraction of old and new images
4. Identification of supernova candidates
5. Logging of new information, and
6. Confirmation of candidates and selection of next field

The following sections describe qualitatively the algorithms used to accomplish these

functions. Because we are still developing the software, no complete description can be given. A relatively complete set of software documentation will be available as The Berkeley Automated Supernova Search Software Manuals at some future time; a section of the Software Manuals describing some of the image cycle tasks is attached as appendix 3.

3.2.2. Acquisition of Images

Because images are currently recorded on videotape and processed "off line", acquisition of images consists simply of setting the videotape recorder to "play" and reading in the next image. One task (currently CCDMEM) sets up the DMA transfers for the image data, and a separate task reads log information, including the identity of the sky region, from the audio channel of the tape. "Sky region" is the term for a particular CCD-image-sized area of the sky, usually containing a galaxy; reference information is filed by sky region identity, or SRID.

Information for the telescope control computer is transmitted via serial link; for each sky region the coordinates, SRID, and exposure time are transmitted. Currently this is done by an independent program, working from a list generated by the user for a particular night. Eventually, the telescope control program will be a part of the image cycle, so that images can be acquired and processed directly from the telescope.

3.2.3. Retrieval of Reference Information

Reference information about each sky region is stored in a series of files, identified as

(SRID).DES Description: Telescope pointing information, galaxy catalog references, classification, history of supernovae, operators notes.

(SRID).OBC Object catalog: CCD Coordinates and brightness of up to 10 fiducial stars, up to 16 "special objects" such as variable stars, and up to 128 "threshold objects" -- all bright objects in the field not otherwise identified.

(SRID).RSP Reference subimage parameters: CCD coordinates of galaxy and galaxy core; size, brightness, and other parameters of core; coordinates of up to 16 "subimage problem objects" such as variable stars or previously discovered supernovae within the galaxy subimage; and up to 128 "hot pixels" -- sites of false alarms around the galaxy core or foreground stars.

(SRID).RSA Reference subimage array: The actual image of the galaxy, derived from one or more reference CCD exposures.

(SRID).OLG Observation Log: Detailed record of the four most recent observations of a galaxy, including a list of all supernova candidates detected.

(SRID).CLG Chronological Log: Summary record of all observations of a galaxy, and of all manual changes to the data files.

This approach leads to a large number of small disk files, thus making somewhat inefficient use of disk space and access time, but allows all the work of locating information to be given to the operating system.

3.2.4. Matching and Subtraction of Images

CCD coordinates and brightnesses of 5 to 10 bright foreground stars ("fiducial stars") are used to determine both the offset (telescope pointing error) between the current and reference images, and the required corrections for exposure, atmospheric

transmission, and seeing. At this time, no corrections are made for optical distortions (mainly pincushion), field rotation, or magnification changes; over the width of a galaxy (typically less than 100 pixels), these effects shift the image by much less than a pixel.

Because the telescope pointing during the search is expected to be accurate to about 10 pixels, we do not need to scan the entire 160,000 pixels of the CCD image to locate fiducial stars. A small area around the nominal position of each star (as stored in (SRID).OBC) is scanned for pixels above a threshold; the star is assumed to be centered on the brightest such pixel. A centroiding routine determines the exact center, with a typical accuracy of 1/8 to 1/4 of a pixel. We calculate a mean offset, with missing or poorly aligned stars discarded.

We have found experimentally that, given the vignetting of our optical system and the substantial variations in CCD quantum efficiency, it is necessary to "clean" each star subimage to obtain consistent brightnesses. In addition, especially for our simple summing and centroiding routine, the level of the sky background around each star must be known to one or two A/D counts to get accurate results -- the total brightness of a typical reference star is only a few hundred counts, and an error of 2 counts summed over 49 pixels (7x7 integration) is a 20% error in brightness.

The cleaning, or field flattening, process consists of subtracting a standard dark field from the subimage, and dividing it by a quantum efficiency field (actually an exposure of twilight sky, minus a dark field, normalized to a convenient mean brightness). Both the dark field and the quantum efficiency field are stored in memory, so that small areas can be efficiently cleaned. We do find that this simple flattening process tends to leave smooth variations of up to 5 percent in the background level over the CCD (figure 3.1). We therefore calculate a local background for each fiducial star

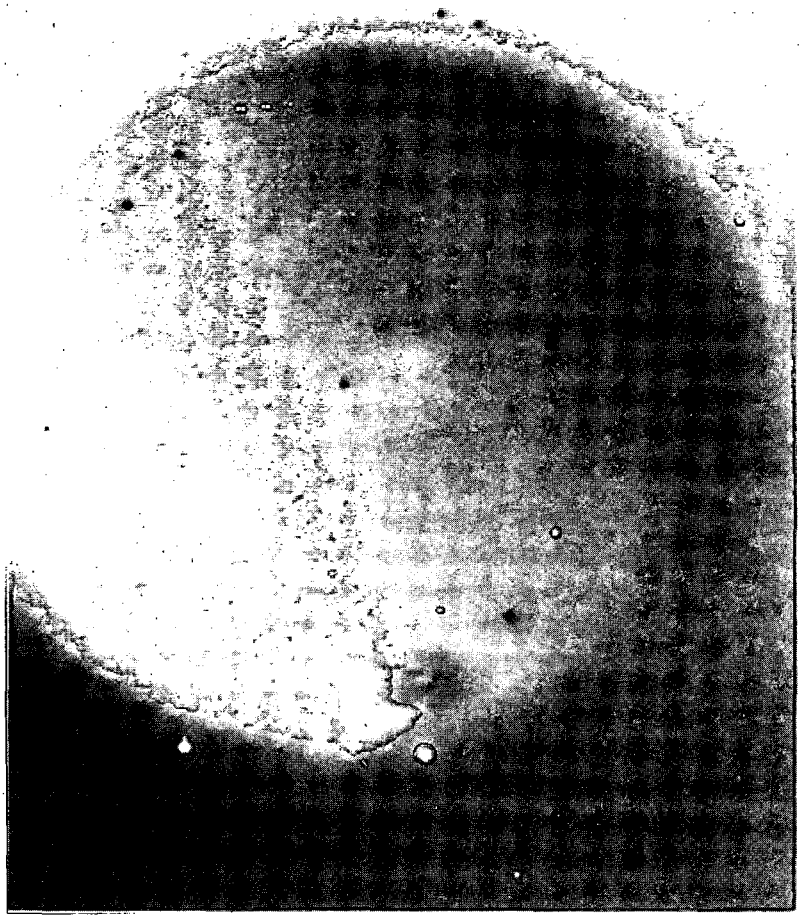

Figure 3.1

Background variations after field flattening: Ratio of two flat fields (twilight sky exposures) from different nights. Display range is 95% (white) to 105% (black); the original images averaged approximately 7000 A/D counts per pixel. Bright and dark spots are stars in the two images.

before finding the brightness and centroid.

Given the current offset, the software can locate the portion of the CCD containing the galaxy and copy just this portion to a separate region in memory; all further processing is done on this subimage. The subimage is flat fielded in our standard fashion, by subtracting off a dark frame and dividing by a quantum efficiency frame. Again, we find a local sky background level and subtract it off. Finally, the current subimage is multiplied by a constant to correct for atmospheric transmission and exposure variations.

At this point, a reference subimage is retrieved from disk and "jogged" (shifted by a fraction of a pixel) to match the current subimage. Also, the subimage with better seeing (normally the reference subimage) is blurred by convolution with a gaussian to further match the images. Currently, both functions are performed by one routine, which does a convolution between the reference image and a 6 x 6 array containing an appropriate gaussian blurring function, off center by the appropriate fractional pixel shift. The fractional shift is made to the nearest 1/16 pixel, though in practice errors of 1/8 or even 1/4 pixel cause only small changes in the results of a subtraction.

This routine is quite slow, taking up to 20 seconds for a 100 x 100 pixel subimage. Also, each pixel is treated as a point value, and no account is taken of the "shape" of the image, which results in a substantial loss of information. In particular, shifting an image by fractional pixels produces substantial blurring even when none is desired, because 4 input pixels are simply averaged together to obtain each output pixel. Various more sophisticated algorithms are possible, from simple fits to neighboring pixels to Fast Fourier Transform (FFT) processing, which use more of the available information to produce more accurate results. These require much more

processing time, however, and are thus slower still -- too slow for real time processing of an entire subimage.

Based on our current experience, sophisticated algorithms will be needed only in the immediate neighborhood of galaxy cores, and that far from the core (and from bright foreground stars) even the current algorithm may be unnecessarily complex. Future versions of the blurring routine will probably employ FFT routines on 8 x 8 or 16 x 16 pixel patches surrounding galaxy cores, and may do only a 2 x 2 convolution -- or perhaps nothing at all -- outside of the core. Foreground stars occupy a small fraction of the galaxy subimage, and may be easier to simply discard than to fit with a complex algorithm.

Once the current and reference subimages are matched, a simple routine subtracts the reference subimage, as processed by the blurring and jogging routine, from the current image, and stores the resulting difference subimage in a separate region in memory.

3.2.5. Identifying Supernova Candidates

The first step in identifying supernovae is to simply list all pixels which are significantly above the background of the difference subimage. "Significantly" is defined as a set number of standard deviations, typically 4, above mean background, with both the mean and the standard deviation of the background calculated anew for each difference subimage (note that later stages in the selection process will require a higher threshold; this is simply a first cut). In an ideal system, this would be sufficient; a bright pixel would imply a bright point in the sky. In practice, such a pixel will usually be a false alarm. Table 3.1 lists some sources of false alarms, with estimated frequencies.

Table 3.1 -- Expected False Alarm Rates

Source of alarm	Rate	Identified by
Statistical fluctuations	10^{-6} pixels 10^{-2} galaxy subimages	Second look
Cosmic rays	10^{-1} CCD frames 10^{-2} galaxy subimages	Second look
Asteroids	10^{-2} galaxy subimages in the ecliptic	Third look (30 minute delay)
Subtraction errors ("Hot pixels")	Several per subimage	Software, Test exposures*, Second look
CCD faults ("Bad" pixels or columns)	Few on current CCD	Software
Data transmission errors	Several per subimage	Software Using VTR storage
Variable stars		
Short period	?	Test exposures*
Long period or low amplitude	?	Software
flare stars, novae, etc.	Rare	Light curves**

* Test exposures are taken immediately after the reference exposure, and for the first weeks of searching, to identify hot pixels and short period variables.

** Flare stars, etc., which can only be distinguished from supernovae by light curve or spectrum, will occasionally be identified as possible supernovae; the goal of early identification takes priority over complete elimination of false alarms.

Most false alarms that occur far from the galaxy core and from foreground stars are due to rare events -- cosmic rays, statistical fluctuations, etc. -- which do not recur. These do not appear in a second image, and are thus eliminated by the confirmation process (see 3.2.7). A few of these can be detected directly by software; for example, single, extremely bright pixels are probably cosmic rays. However, the major source of false alarms, which must be handled by the image cycle software at this level, is poor subtraction of the steep brightness gradients around galaxy cores and stars. (To be precise, errors are due to gradient variations, or second derivatives, in brightness, since even our simple interpolation routines are exact for uniform gradients. In practice, most brightness gradients are not uniform over more than a few pixels, so we simply refer to "gradient errors".) Such subtraction errors can be seen in the sample image processing cycle in the next section (figures 4.1 to 4.7)

Subtraction errors can only be eliminated by oversampling at least the reference image. If one records CCD pixels smaller than 1/2 cycle of the highest spatial frequency present, and aligns images with perfect accuracy, an exact subtraction, even of abrupt gradient changes, is possible. For search images, however, this lowers the signal to noise ratio for supernovae, as the supernova light is spread over several pixels, and in any case requires unreasonable amounts of processing from our present computer.

At present, we cannot accomodate different pixel sizes for reference and search images: the telescope optics are not easily interchanged, and the existing software cannot correct for changes in magnification. With a single fixed pixel size, optimized for best signal to noise, one can reduce subtraction errors with a sophisticated image processing algorithm, such as an FFT, but not eliminate them entirely. Even with high resolution reference images, fluctuations in seeing (particularly when the reference image seeing is relatively poor) will leave some errors in subtraction.

Our solution to this problem is to acknowledge that abrupt gradients in our images will always produce some "hot pixels", but because these gradients are localized, the errors will always occur in the same small regions. We therefore simply generate a list of all the possible hot pixels, with a maximum brightness, and disregard any pixel on that list which does not exceed that brightness. In the interest of simplicity, the source of this list is a series of test images (currently 10) of the galaxy in question; these are processed through the image cycle in the usual way, except that all bright pixels are collected. Any location which appears bright in 2 or more of these 10 exposures is included on the "Subimage threshold" pixel list, and pixels at the corresponding locations in search images are disregarded unless they exceed a suitable (high) threshold.

This process typically generates 10-50 hot pixels, out of 10,000 pixels in a 100 x 100 subimage, so the area lost is very small. Most of these are in the neighborhood of foreground stars, and thus at very poor points for studying supernovae in any case. The only important loss is in the cores of galaxies -- supernovae can still be detected in cores, especially in nearby galaxies, but only at an increased threshold. We consider this an acceptable trade for simple software. In the future, we can use more sophisticated image processing, as described above, plus models for the shape of galaxy cores and possibly high resolution reference images of the core (obtained either with different reducing optics or by combining many exposures) to achieve accurate subtractions in galaxy cores at the expense of increased computer and programmer effort.

Once a list of bright pixels is generated, the list is checked in several ways. The reference hot pixel list described above is used to eliminate as many candidate pixels as possible. In addition, neighboring pixels in the difference image are checked; if adjacent pixels are above background, even by two or three standard deviations, the

bright pixel is a supernova candidate; if adjacent pixels are at the background level, the pixel must exceed a higher brightness threshold (typically five standard deviations) to be retained. Other "supporting evidence" is the presence of the same pixel location (allowing for offset) on the list of bright pixels from the previous observation; other negative evidence is the presence of even a very faint foreground star at that location in the reference image. With all these factors included, one or more pixels may be listed as supernova candidates. The exact weighting of the different factors is still being adjusted; the goal is to raise our sensitivity until we find one supernova candidate, on the average, in every 10th field observed, and thus need to recheck about one field in ten.

3.2.6. Logging of New Information

The main output of the image cycle routines is a list, freqently empty, of supernova candidates. However, much additional information is also logged, including the coordinates and brightnesses of all fiducial stars (to aid in correcting the telecope pointing, and to provide a record of atmospheric properties), and the list of all pixels exceeding the minimum threshold (to update the hot pixel list, and for consideration in the next observation). Also, "special objects", such as variable stars or supernovae already located, may be specified by the observer; the brightness and position of each special object is recorded. The total amount of stored information is a few thousand bytes for each observation.

The full log file for each sky region is retained for the latest four observations; older observations are purged to save disk space. This allows time for non-real-time programs to analyse the stored information and condense it, or transfer it to magnetic tape or other off line storage, as desired.

A chronological log file is maintained for each sky region in addition to the main observation log. This contains a one-line record of each observation, and of any modification of the reference files. The chronological log provides an "audit trail", so that any activity relating to that sky region can be traced by time and date.

Whenever a supernova candidate is identified, the entire current subimage is also saved on disk. This allows the operator to double check the image cycle results, and to analyze the causes of false alarms.

3.2.7. Confirmation of Candidates; Selection of Next Field

Normally, the controlling program of the image cycle simply runs down a list of sky regions provided by the operator, acquiring images of each galaxy in turn. If a possible supernova is identified, however, the name of that sky region is placed in a separate queue. The region is rechecked as soon as possible -- when operating with the computer at the telescope, immediately. Currently, when the videotape recorder is in use, we take two pictures of each galaxy in rapid succession so that the computer can recheck the region immediately.

If no bright pixels are found at the matching coordinates on the second exposure, the candidate is discarded as a statistical fluke, a cosmic ray, or other transient effect. If a bright pixel is found at the appropriate spot, the candidate is obviously of interest, and a printout is generated. However, the supernova is not considered confirmed until a third image, delayed by about an hour, is taken. This delay is sufficient to rule out asteroids, which move at a rate of approximately 1 arc second every two minutes, or some 10 pixels per hour. (Note that we need to exercise a bit of care not to pick up asteroids right at the inflection points of their apparent motion; fortunately, the area of sky where, at any given time, asteroids are turning around and backing up is very small.) For the videotape-based search, the delay for this

third exposure is normally 24 hours, which is thus the worst-case delay for confirmation.

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 is variable foreground stars, including ordinary novae, which increase in brightness by at least two magnitudes. This is the difference between the detection threshold for a reference exposure (typically three minutes) and for a normal exposure (currently 30 seconds). We cannot now eliminate such variables, but we do not expect to find many of them. When we are not trying to collect very early time spectral information, we can simply follow the light curve of each possible supernova for long enough to confirm our find. For nearby supernovae, we may be able to rule out some candidates by color difference measurements, but in the interest of very fast response, we will have to risk an occasional public false alarm.

3.3. Support software

In addition to the image cycle programs, we have written several tens of thousands of lines of FORTRAN code for support programs and subroutines. These may be divided roughly into three sets: Low level utility routines, high level (user) utility routines, and new region generation software. The following sections give short overviews of this software. Because many of the utility routines are continually being upgraded (and new routines added) a more complete description in this paper is not practical.

3.3.1. Low level utilities

Because of the 16 bit addressing limit of RSX-11M tasks, Roger Williams has written a series of subroutines which allow the creation and manipulation of large

blocks of memory, called "regions" (as distinguished from "sky regions" -- an unfortunate but now established repetition). These region access routines are used for essentially all manipulation of images and subimages larger than a few hundred pixels; actual processing is done in local buffers, and these buffers are copied from and to the regions as necessary. We originally expected to be able to operate with one or two large (160K word) regions and a few smaller regions for subimages; in practice, we have added memory until we can keep as many as 5 complete images in memory at once.

Because regions are independent of individual tasks or users, we have also used the region access routines to pass data among tasks, in place of either global FORTRAN common blocks or disk files.

Other low level routines which have been heavily used include the software interfaces to the CCD and the graphics display, a series of disk access routines which format supernova search data for storage, and particularly a set of spawning routines. These last provide for the convenient spawning, or calling and running, of independent RSX-11 tasks from within a program. Although the main purpose of multitasking is to allow many programs to run concurrently, we have used multiple spawned tasks primarily to escape from the 64K byte task address limitation.

In retrospect, many of the functions performed by these utility tasks, at substantial cost in time and programming effort, would have been superfluous on a 32 bit computer with a large direct address space. However, the limitations of our current machine have been indirectly beneficial, in forcing a modular appproach to programming, and in giving us great control over the exact sequence and timing of operations -- control which will eventually allow us to run multiple tasks in parallel for maximum speed.

3.3.2. High level utilities

Most of the development of the supernova search hardware and software has proceeded without the use, or availability, of a complete image cycle program. Image processing during development and debugging has been done by a series of separate, user-callable programs which perform simple functions. Many of these are simply components of the image cycle, equipped with a calling program which can be run from a terminal -- for example, to subtract one subimage from another. Others are written specifically for debugging, although they have found wide use -- for instance, the program GNDISP (reGioN DISplay with Print), which displays the numeric values of pixels in a specified patch of a specified memory region, was written simply to check the operation of other image processing programs. However, with the addition of some simple statistical functions, calculated on the displayed values, this program has been useful for testing the optical system transmission, estimating seeing, and inspecting videotape data recording errors.

3.3.3. New regions

A series of routines closely related to the image cycle routines perform the function of generating new regions. The new region programs accept as input one or more high quality (long exposure) images of a galaxy, and generate the set of data files, including a reference subimage file, needed by the image cycle to process that galaxy. In addition, the new region program uses the image cycle routines to process a set of, typically, ten sample images in order to generate the subimage threshold pixel list. This is the list used by the image cycle to identify false alarms due to foreground stars and galaxy cores.

4. Performance of the Berkeley Supernova Search

4.1. Demonstration of the Image Cycle

Figures 4.1 - 4.7, following page 59, show the operation of the image cycle. Figure 4.1 shows three CCD images -- a three minute reference exposure, a 30 second "test" exposure, and a 30 second "search" exposure. All are "raw" images, as read in from videotape, except that the "search" exposure contains a 16th magnitude supernova -- a bogus one, unfortunately, added via software. The brightest pixel in the supernova contains 200 A-D counts, and the total brightness is 625. The vignetting due to the reducing optics is visible in all three images.

Figure 4.2 shows the same three exposures after field flattening with a standard dark frame and quantum efficiency/vignetting frame. Note that such full frame field flattening is normally only done on the reference exposure, for the purpose of selecting fiducial stars.

Figure 4.3a is a printout of the list of fiducial stars generated from the reference frame -- the stars are selected by the operator with a joystick, and measured by the computer, although we have routines which automatically select the ten brightest usable stars in the field. Figure 4.3b shows these stars as retrieved from the test frame -- a small (16x16) box around each star is cleaned and saved in a separate memory region. These star subimages are processed by the offset finding routine to obtain the data printed out in 4.3c. Note that the offset is 0.59 rows and -0.44 columns (these exposures are all from the same night, so the offsets are small), with errors of about 0.1 pixel. The calculated sky transmission (brightness of fiducial stars in the test frame vs. the expected value calculated from the reference frame) is very close to 1.0. The "SKY" column gives the local sky background measured by the star finding routine; the variations from star to star demonstrate the need to use such a

local background measurement with our crude field flattening.

Figure 4.4 shows the subtraction process. 4.4a and b are the raw and cleaned current (test) subimages, and 4.4c is the reference subimage, scaled in brightness to match the standard (30 second) exposure and jogged by the necessary fractional pixels. Note that the background noise level in the reference much lower than in the current image -- the dynamic range is the same (100 A-D counts) in all four subimages. Finally, 4.4d is the difference between the current and reference subimages. The jogging process broadens the reference picture, as does the longer exposure (due to telescope tracking errors) and to some extent nonlinearities and charge leakage in the CCD. The net effect is that sharp gradients around foreground stars are not well subtracted -- here stars appear as central peaks surrounded by negative (light, on these pictures) rings. Even these apparently bad subtractions leave peaks only about 20% of the amplitudes of the original stars, but some further correction is required. The symmetry of the foreground star remnants indicates that the subimages were well aligned. Figure 4.5 shows the need for such alignment; it includes the reference subimage and difference subimage if no fractional pixel jog is done, leaving about a 1/2 pixel error on each axis.

One way to discriminate against such subtraction errors is to look for nearby "antisupernovae" -- the average of a 7 x 7 box around each foreground star is very near the background level -- but this requires averaging many pixels, which introduces additional noise. The supernova search uses a different technique to eliminate false alarms due to gradient variations, illustrated in figures 4.6 and 4.7.

Figure 4.6a shows the test difference image from figure 4.4. Figure 4.6b shows the same subimage, but with the display range adjusted so that only pixels brighter than four standard deviations (20 A/D counts) above the local background are visible.

These are "hot pixels" -- known sources of false alarms. A list of such pixels is generated and stored with the reference data for each galaxy. Normally, the list is compiled from ten or more test images, but for these figures, we use just the pixels in 4.6b.

In figure 4.6c, each pixel in 4.6b has been expanded into a 3 x 3 block of pixels, forming a mask. In this demonstration, we will simply ignore the pixels covered by this mask. A more sophisticated processing algorithm would take into account the brightness of each hot pixel, and the exact location including fractional pixel shifts, to adjust the detection threshold in each pixel just enough to eliminate false alarms. However, as shown in figure 4.6d, very little of the subimage is blocked by even this crude mask, so only a very few supernovae would be lost.

Figure 4.7a shows the search exposure subimage, ready for subtraction; the added "supernova" is visible but not conspicuous. Figure 4.7b shows the difference subimage. In this figure, the supernova is conspicuous to the human eye, as it is the largest bright object remaining, and is not surrounded by a ring of "antisupernova". However, it is still difficult to arrive at a simple algorithm which will allow the computer to recognize the supernova, except by averaging over many pixels and losing signal to noise ratio.

In figure 4.7c, however, we have used the mask from figure 4.6 to block out expected false alarms. The supernova stands out alone above the background. It is even more conspicuous in figure 4.7d, where the display range has been adjusted to show only pixels more than five standard deviations above background. At this stage, a single remaining bright pixel would be noticeable, even to a computer.

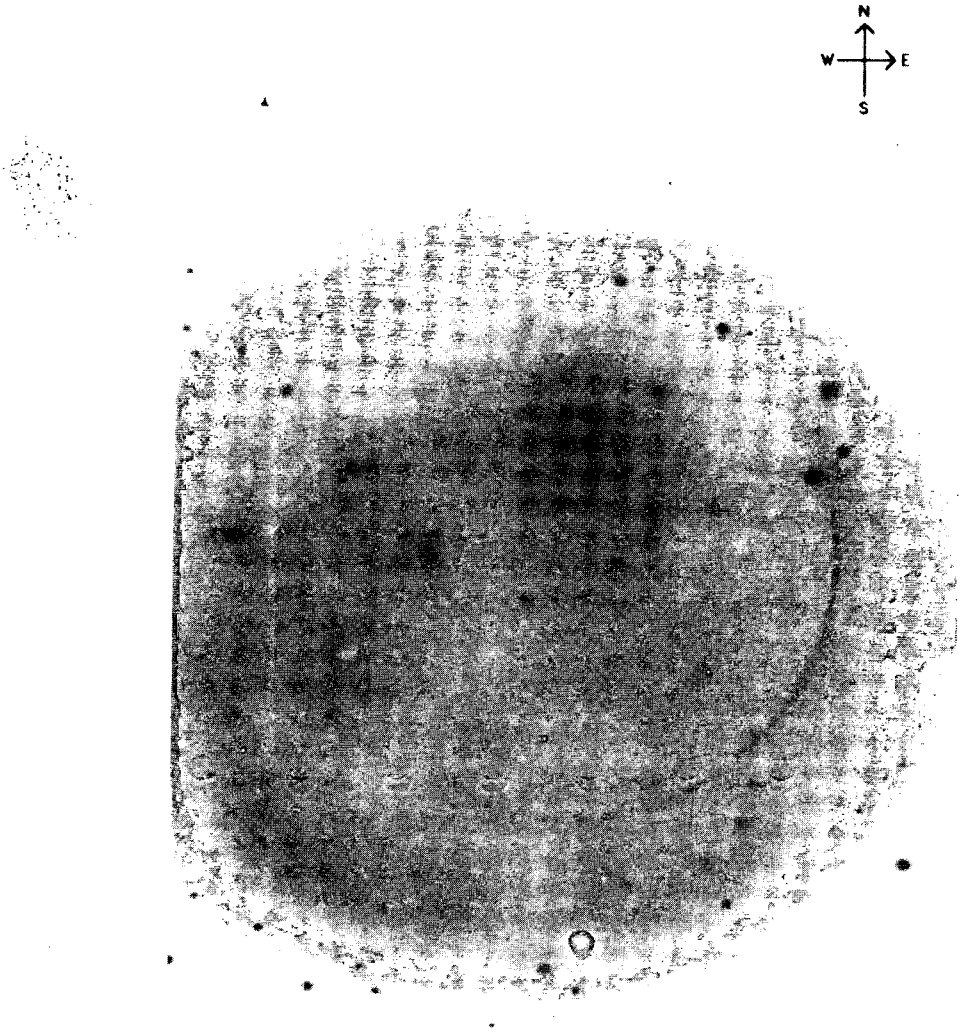

Figure 4.1a

Reference exposure of NGC0628, a large spiral galaxy. Exposure time 3 minutes. Display range 1250-1750 A/D counts. Note CCD defects: Dimple (white spot in lower center), flare (black triangular mark at upper left, due to pixel with high dark current leakage), teacup rings (dark rings of high quantum efficiency around center), and cold columns (white vertical lines at left). Overall shading is due to vignetting in the reducing optics. Note: the "top" of the CCD is always the north, narrow end, even for figures oriented sideways on the page. The CCD is read out from left to right (west to east) within rows, and by rows from bottom to top (south to north).

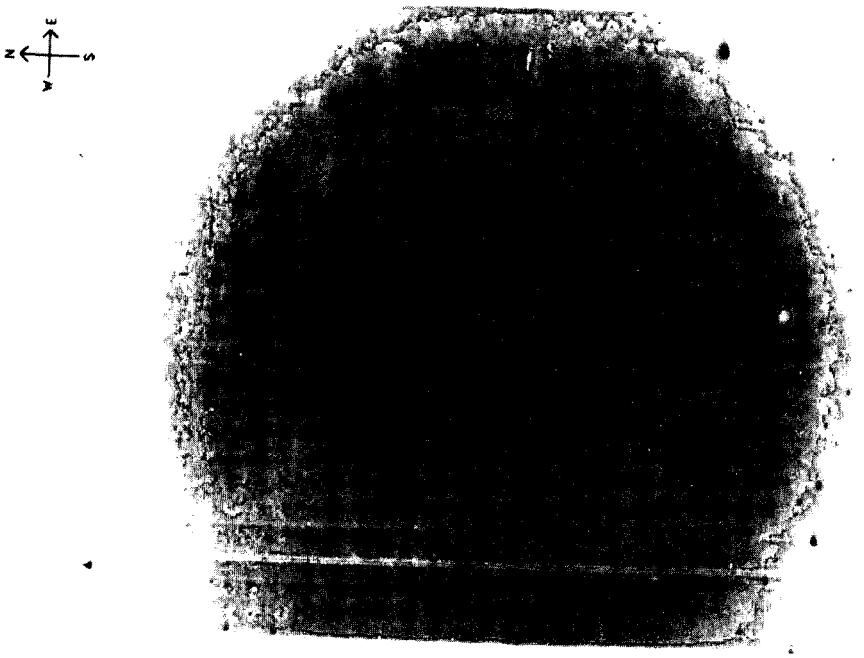

Figure 4.1b

Test exposure of NGC0628. 30 second exposure, display range 550 - 700. Note prominent cold columns and faint 120 Hz ripple (bars running roughly east-west).

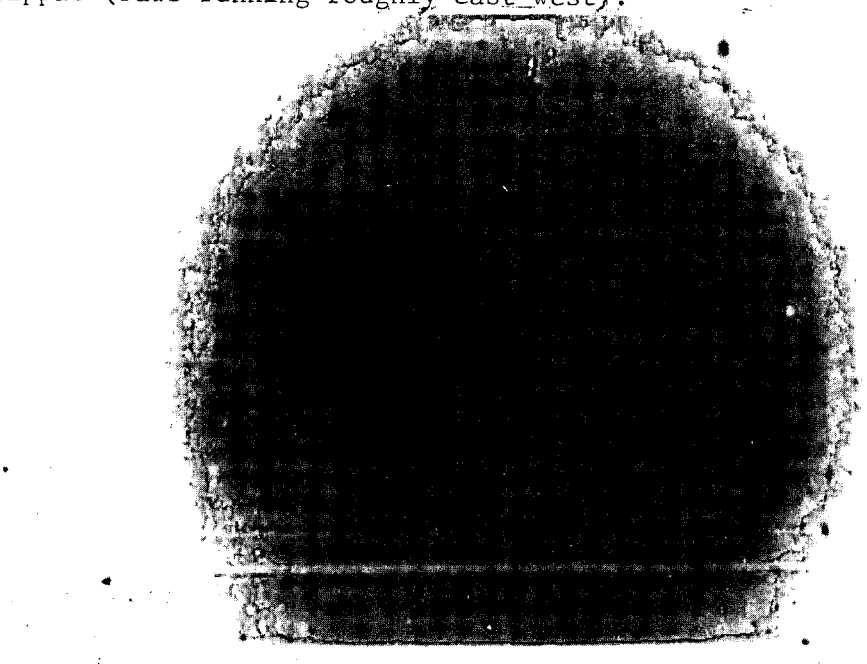

Figure 4.1c

"Search" exposure of NGC0628. 30 second exposure, as in 4.1b, with approx. 15th magnitude star added east of galaxy core. Note: all exposures were taken the same night.

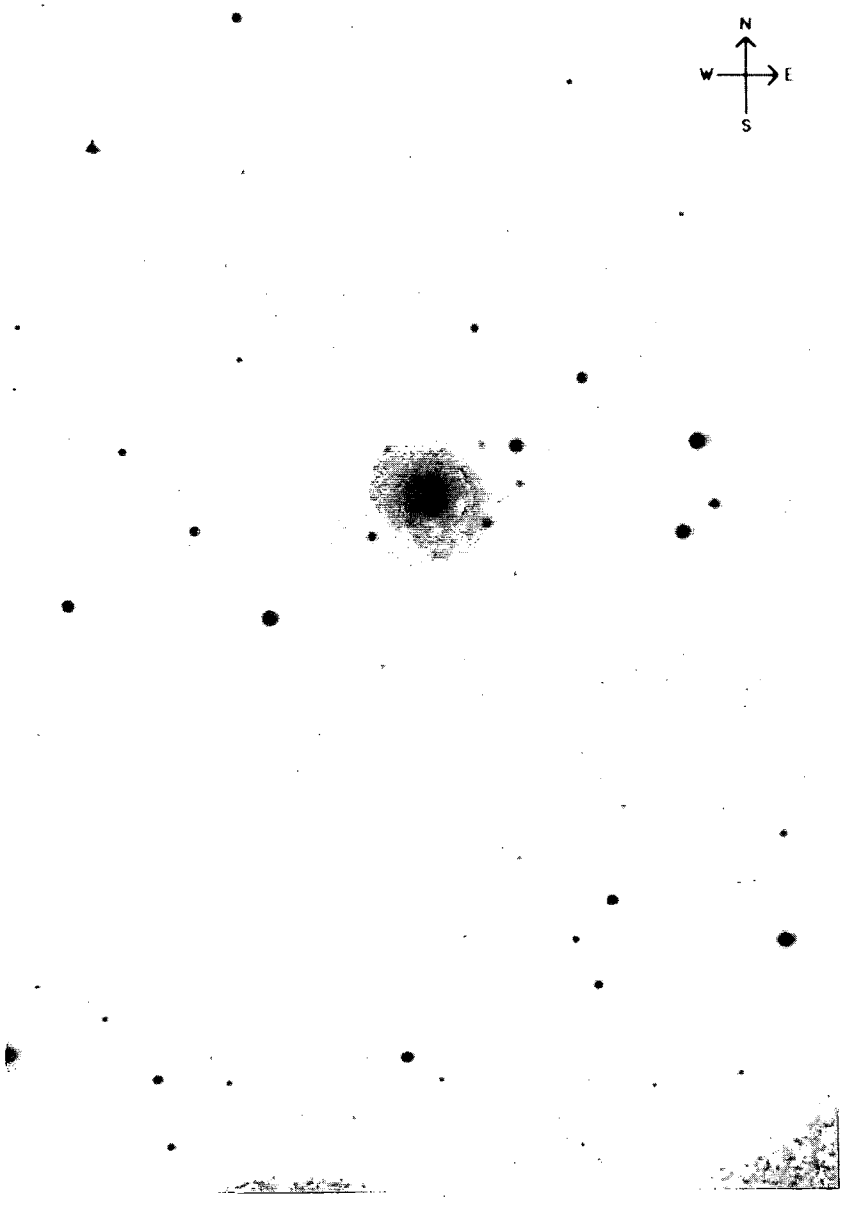

Figure 4.2a

Reference exposure of NGC0628 after field flattening by subtraction of a dark frame and multiplication by a "quantum efficiency" frame (actually the inverse of a normalized exposure of the sky at twilight, which corrects for both quantum efficiency and vignetting. Flares along bottom edge and bottom right corner may be due to electro-luminescense in the output register. Display range 1200 - 1500.

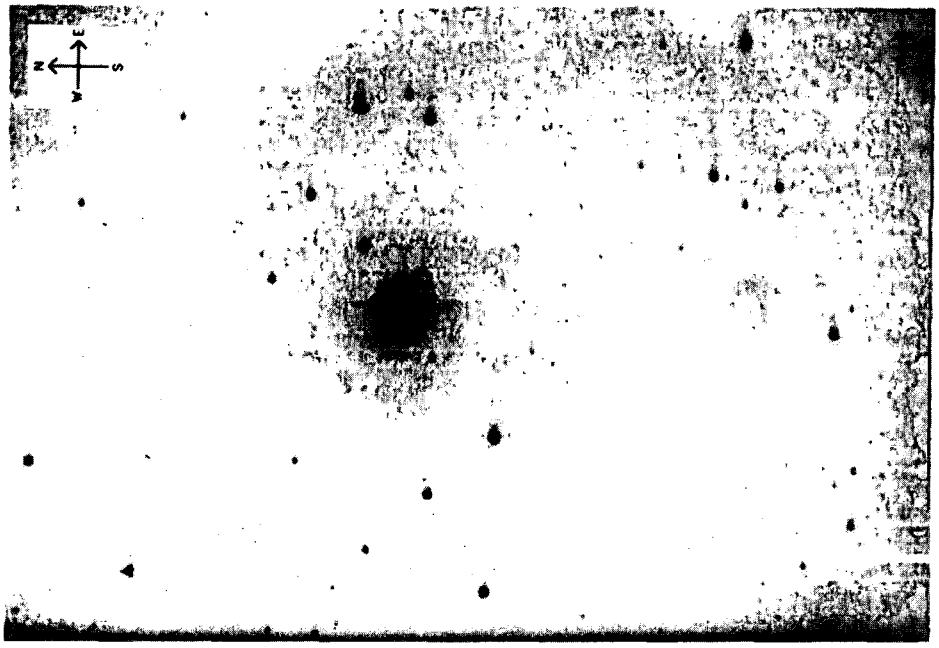

Figure 4.2b

Test exposure of NGC0628 after field flattening. Display range 570 - 670. Note relatively high noise level compared to 4.2a.

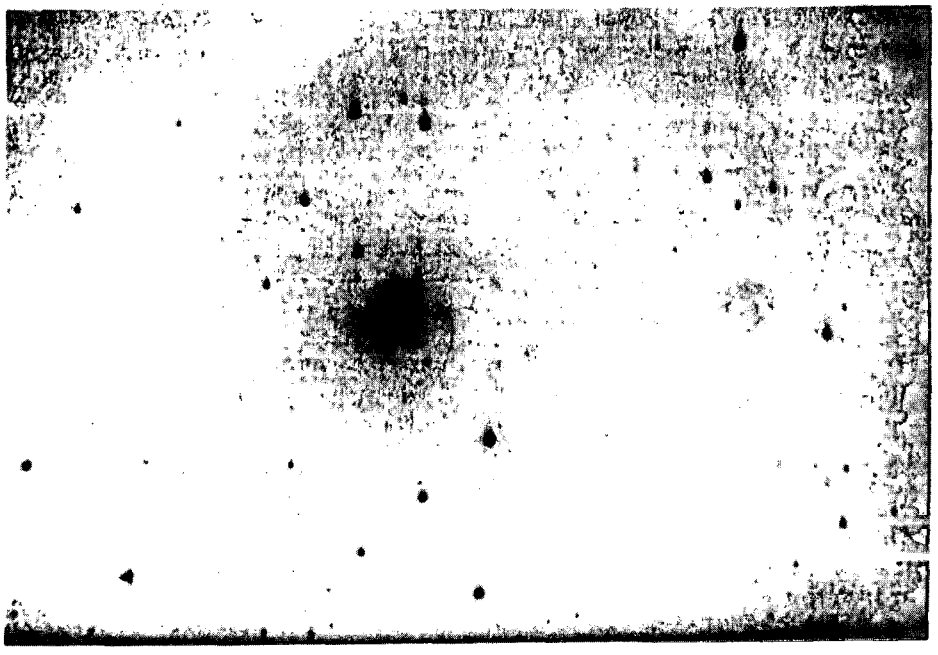

Figure 4.2c

Search exposure of NGC0628 after field flattening. Added "supernova" is just visible on the east side of the galaxy core.

```
FIDUCIAL STARS IN n0628
Image coord.s:      1   512      1   320
OFFSET:     0.00 rows,     0.00 columns
SENSITIVITY:    6.000 / exp:180.000 sky:   1.000
BACKGROUND:    1270 / sky:    870 / sigma:     30

#     ROW      COLUMN     BRIGHTNESS  FLAG
     69.75    154.75          4471      1
    134.25    233.19          2799      1
    246.88    103.88          8521      1
    251.88     27.69          5216      1
    281.25     75.25          2614      1
    314.88    197.69          2280      1
    280.69    261.06          5774      1
    316.69    266.88          9093      1
    341.81    222.88          2564      1
    487.94     91.25          3840      1
Select input source:  J for Grinnell Joystick,
                      K for keyboard,
                      QT to quit
```

Figure 4.3A -- Fiducial stars in reference image, selected by system operator but measured automatically. Brightness is the total signal above local background in a 7 x 7 pixel box

```
ZSKYTR:     0.9949122          (Sky transmission; ratio of current to ref. star
ZSKYTV:     5.3845730E-02       brightness, and standard dev. of sky transmission)
ZCSENS:     0.9949123          (Current Sensitivity, nominally 1.0)
sky back:     174                        (Mean sky background)
BLUR:        1036.869  -539.8750         (not used)
OFFSET:     0.5898800 -0.4382077         (coordinate offset in pixels)
OFF. VAR:   8.9996502E-02  0.1077006     (Std. dev. of coordinate offsets)
NSTARO:       10                         (Number of stars located)

      1      2     10      0
 type star data? (Y or N): y
DATBUF  1:  SRID: N0628
  #    row      col      brt      cur/ref     offset              sky
  1   70.376  154.196  4381.000    0.980     0.626    -0.554     175
  2  134.970  232.822  2705.000    0.966     0.720    -0.366     174
  3  247.435  103.428  8768.000    1.029     0.560    -0.447     175
  4  252.410   27.067  5345.000    1.025     0.535    -0.620     168
  5  281.999   74.903  2705.000    1.035     0.749    -0.347     170
  6  315.456  197.194  2474.000    1.085     0.581    -0.493     179
  7  281.210  260.736  5789.000    1.003     0.522    -0.326     176
  8  317.159  266.390  8842.000    0.972     0.472    -0.485     181
  9  342.378  222.431  2245.000    0.876     0.566    -0.444     179
 10  488.646   91.007  3678.000    0.958     0.709    -0.243     161
```

Figure 4.3C -- Fiducial star data recovered from the test exposure.

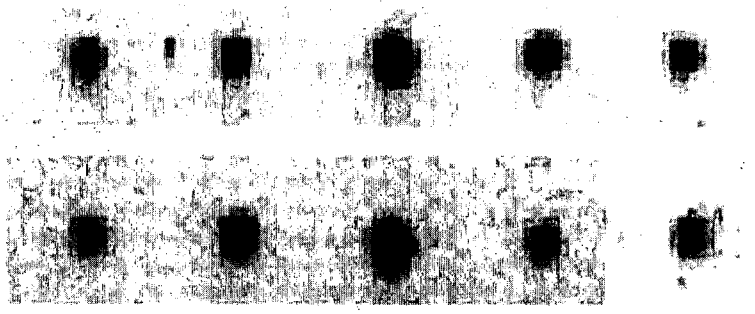

Figure 4.3b

16 x 16 subimages of ten fiducial stars retrieved by the offset finding routine. A neighboring star is visible in the second image from the left, but it is far enough from the main star to have no effect on the coordinate finding routines.

Figure 4.4

Test subimage subtraction

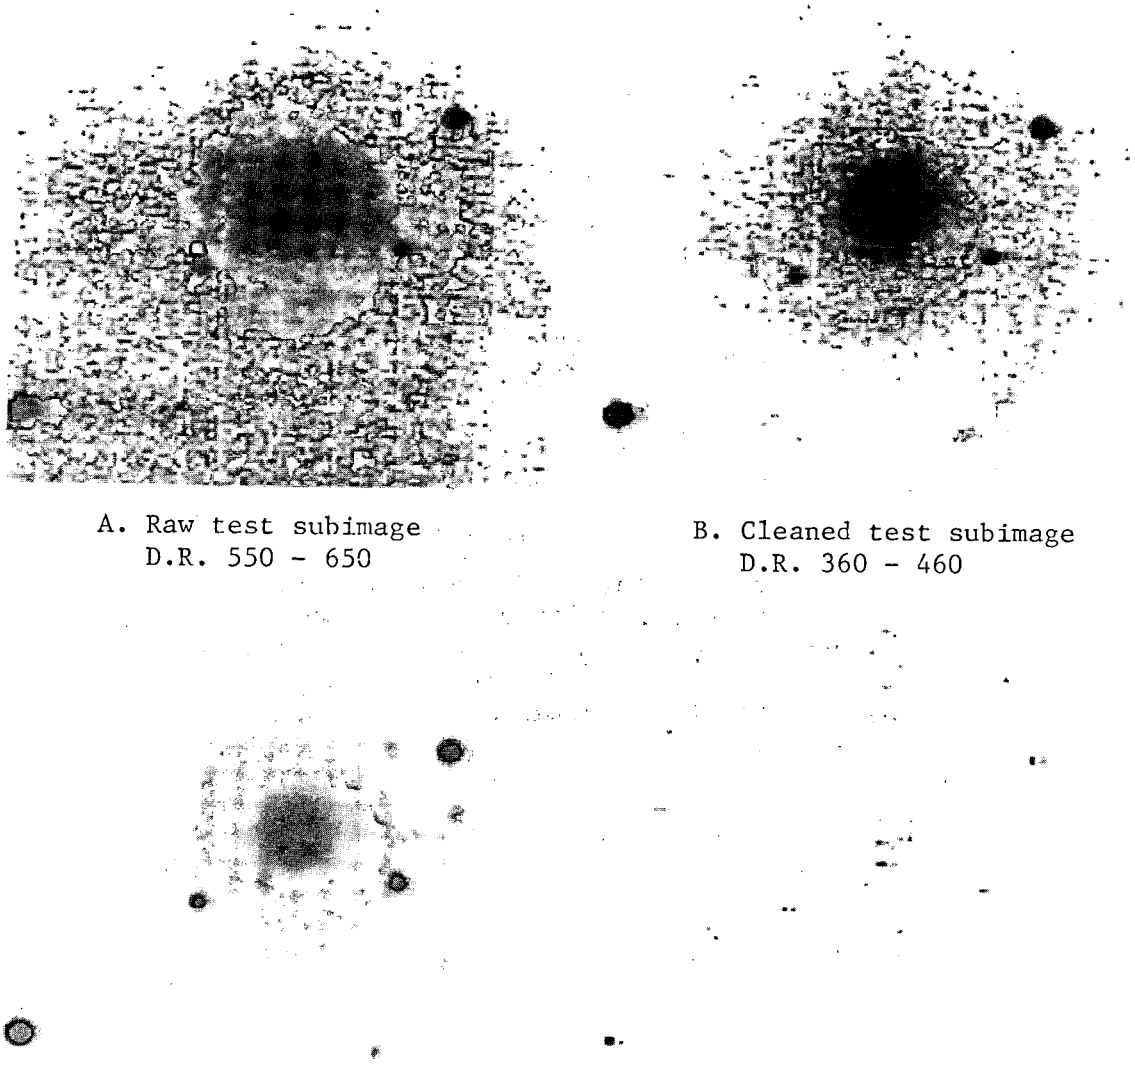

A. Raw test subimage
D.R. 550 - 650

B. Cleaned test subimage
D.R. 360 - 460

C. Reference subimage
D.R. 360 - 460

D. Difference subimage
D.R. 360 - 460

Notes:

B. Somewhat less grainy than A due to cancelling of pixel to pixel quantum efficiency variations.

C. Much less grainy than A or B -- 6 times longer exposure; subimage scaled by 6x to match B. Sky background noise is higher, so noise in scaled image is only reduced by approx. 3x. Image is jogged by 10/16 pixels on each axis, not blurred at all.

D. D = B - C + 400. See Fig. 4.6 for further notes.

Figure 4.5

Effect of fractional pixel misalignment on subtraction

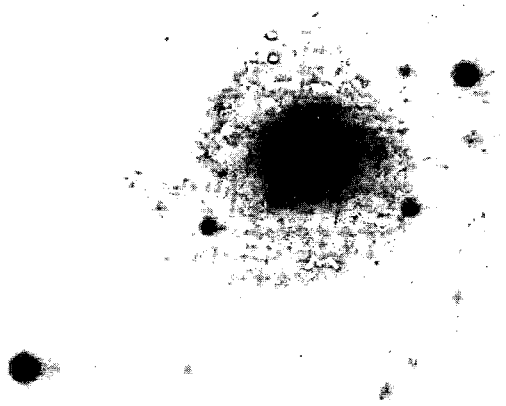

A. Reference subimage, not jogged
D.R. 360 - 460

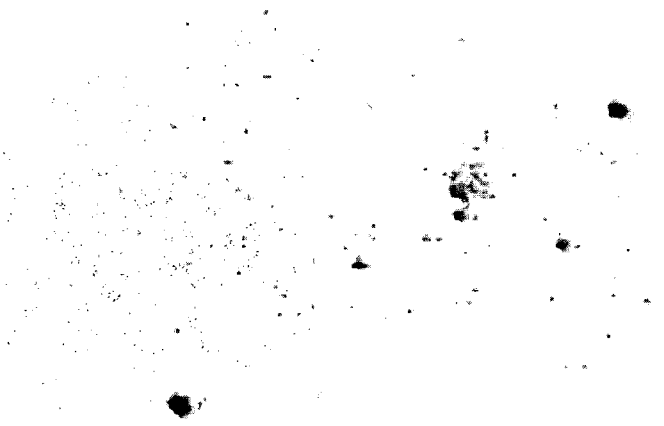

B. Difference subimage (Fig. 4.4B - fig. 4.5A), D.R. 360 - 460

Note asymmetric star/antistar objects, and relatively large signal in galaxy core compared to figure 4.4d. Misalignment is approximately 0.6 pixels (1.5 arc seconds) on each axis.

Figure 4.6

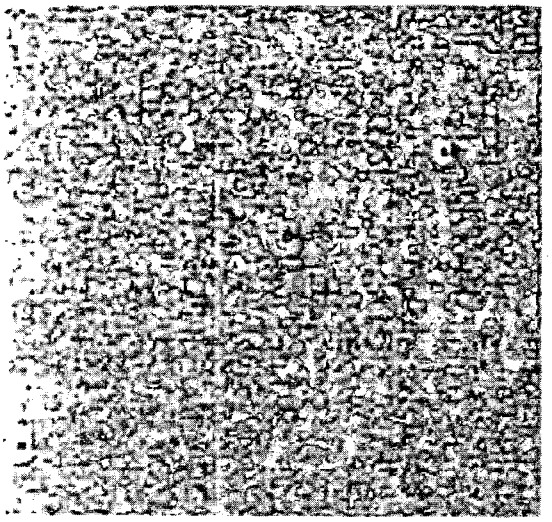

A. Test difference subimage (same as fig. 4.4d) D.R. 340 - 440

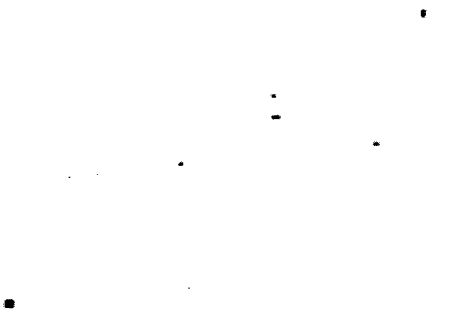

B. Same as A, display range set to show hot pixels. D.R. 420 - 430

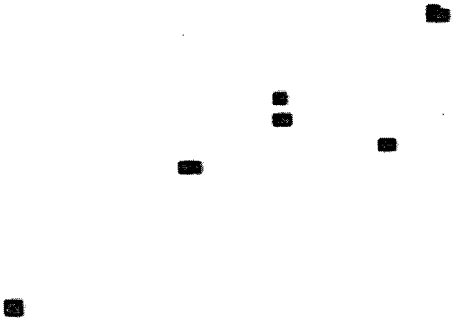

C. Hot pixel mask, generated by 3x3 convolution of B.

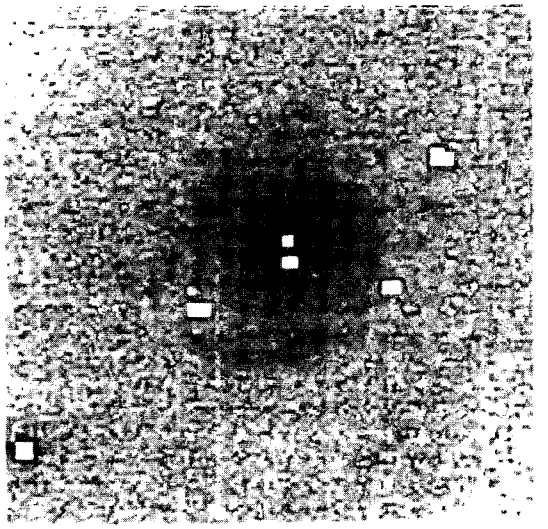

D. Test subimage masked by C -- masked pixels set to 0.

See next page for notes.

Notes for figures 4.4 and 4.6

All subimages are 120 x 120 pixels, with a local background level of approximately 400 A/D counts. In figure 4.4b, the brightest pixel in the galaxy core is 204 counts above background, and the integrated brightness of the galaxy (7 x 7 pixel sum) is 4950 counts. The bright star in the lower left corner contains 8915 counts (7 x 7 sum), and the brightest pixel is 1716 counts above background.

In 4.4d (also 4.6a), the brightest pixel in the galaxy core is 64 counts above background, and the integrated brightness is 93 counts. The brightest pixel in the foreground star is 234 counts, and the integrated brightness is only 57 counts. The background noise level in this subimage is approximately 5 counts r.m.s. in each pixel. The r.m.s. error in a 7 x 7 pixel sum (including the contribution from error in finding the local background by averaging 36 pixels) is approximately 46 counts; a one-count error in finding the local background produces a 49 count error in the sum.

It is possible to eliminate most of the subtraction errors in 4.4d by blurring the test subimage approximately 1/2 pixel to compensate for the blurring introduced in jogging the reference subimage. The reference blurring is unusually bad in this case because the jogging required was close to 1/2 pixel on each axis. Note, though, that it is better to blur the reference image too much, leaving a few bright pixels surrounded by low values, than too litte, leaving many bright pixels around dim central spots.

Figure 4.6b displays the pixels from this test image which would be included in a reference list of hot pixels: all pixels more than 4 standard deviations above local background. There are 12 such pixels in this subimage.

Figure 4.6c is generated by setting each pixel in the reference list derived from 4.6b to 100 counts above a flat background, and expanding each single pixel into a 3 x 3 block of pixels. 4.6d is generated by setting the same expanded list of pixels to zero in the test subimage (same subimage as figure 4.4b). The total number of pixels masked is 82, or approximately 0.6% of the subimage.

Figure 4.7

Finding a Supernova!

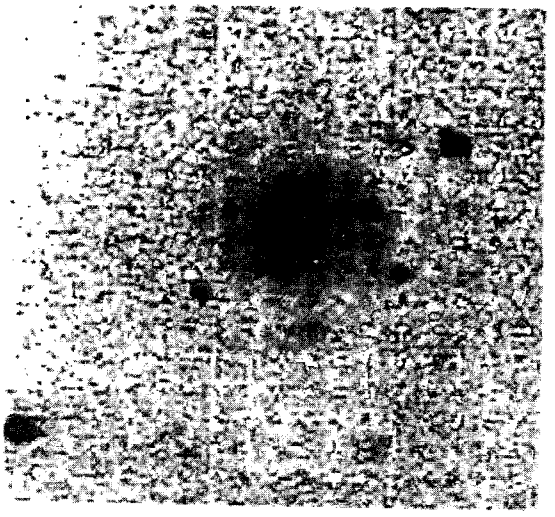

A. Search subimage with 15th mag."supernova"added D. R. 360 - 460

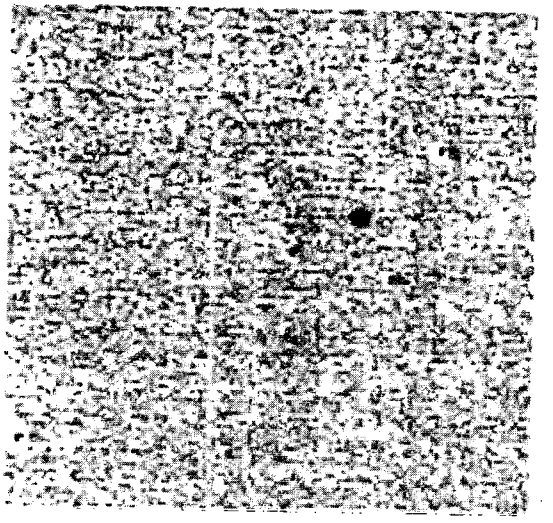

B. Search difference subimage D. R. 360 - 460

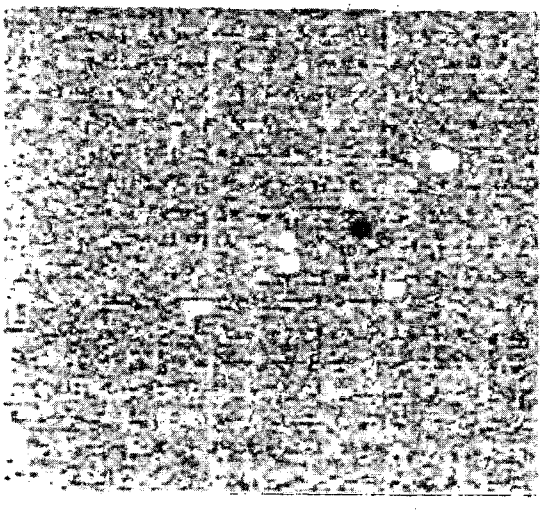

C. Masked difference subimage, using fig. 4.6c. D. R. 360 - 460

D. Masked difference subimage, set to show pixels above 5 std. dev. D.R. 425 - 435

The fake supernova is added by manually editing the original CCD image. The brightest pixel in the supernova is 200 counts above background, and the total number of counts added is 625. The background noise level in the difference subimage is between 4 and 5 counts, r.m.s. The search and test images are offset from each other by approximately 1/2 pixel, with no rotation or change in magnification.

XBB840-8736

4.2. Search Sensitivity

4.2.1 Signal to Noise Calculation

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

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

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 (arc sec^2)
N_{CCD}	=	CCD rms noise, photoelectrons/pixel

Note that the rms noise in the supernova signal does not need to be included in the denominator.

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 88±10 for our current CCD. A new CCD has been delivered by RCA, and is expected to have less than 60 electrons RMS noise; for calculations, we have

assumed that the noise level of the new chip will be 40 electrons.

f is a random variable, depending on the seeing and on the position of the supernova image on the CCD. In the limit of perfect seeing, f goes to 1.0. For typical seeing at the Leuschner site, we find a mean value for foreground stars (brightest pixel/total brightness) of 0.18. This is a rather low value, partly because our current pixels are relatively small (2.4 arc sec square). Figure 4.8 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.

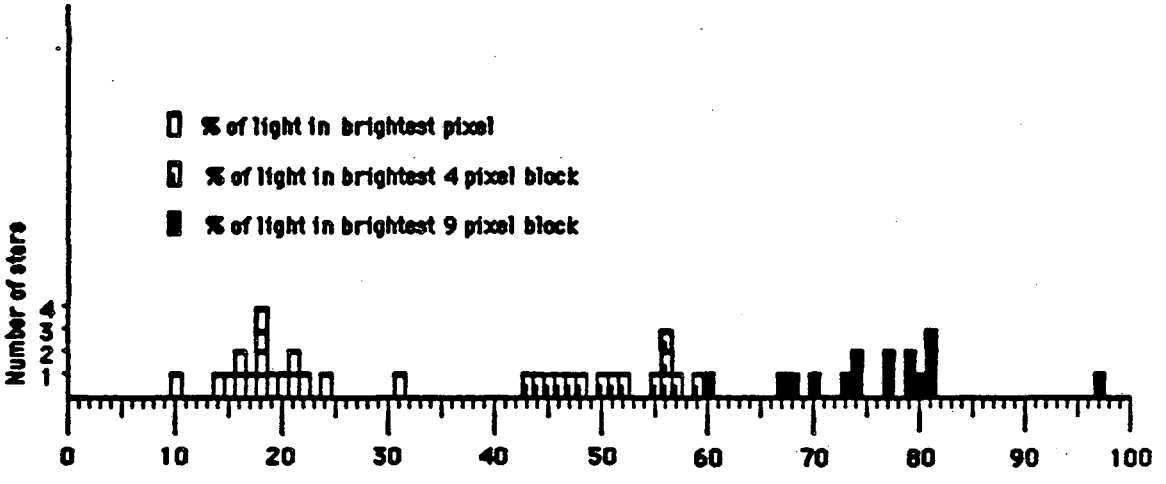

Figure 4.8 -- Fraction f of light in 1, 4, and 9 pixels for 17 stars observed at Leuschner.

4.2.2 Current and Future Threshold for Detection

Figure 4.9 plots the number of photoelectrons detected (or A/D counts, assuming a conversion factor of 28 photoelectrons per count) versus t_{exp} for various combinations of telescope and CCD performance. The signal from various magnitudes of

Figure 4.9 -- Sensitivity and Detection Thresholds vs. Exposure Time, for Various Combinations of CCD and Telescope.

Diagonal lines represent the light collected in one pixel from stars or supernovae of various apparent visual magnitudes. Curved lines represent the system noise level (S/N 1), and the detection threshold of five standard deviations (S/N 5). f (fraction of supernova light in one pixel) is 0.25 except in figure 4.9e; unidentified factor of 2 loss is included in calculations.

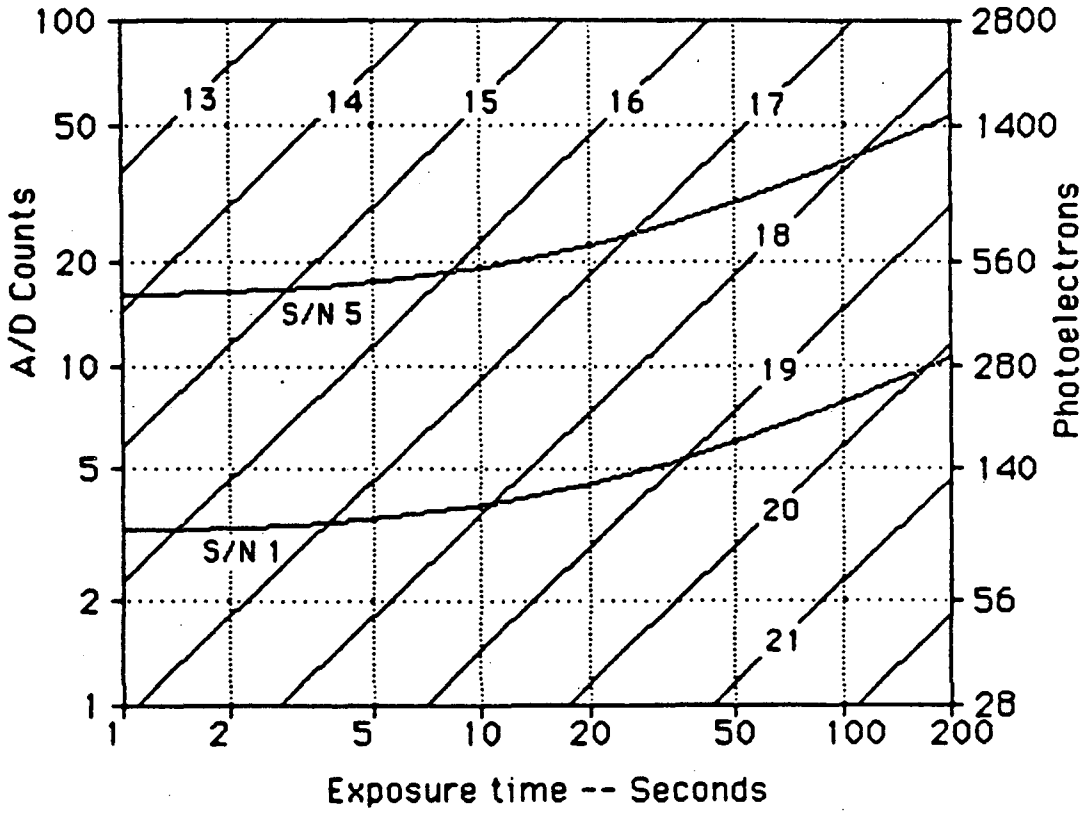

Figure 4.9a -- Current CCD (88 photoelectrons r.m.s. noise, quantum efficiency 30%) at Leuschner telescope (30" aperture, sky background 400 photoelectrons per second per pixel).

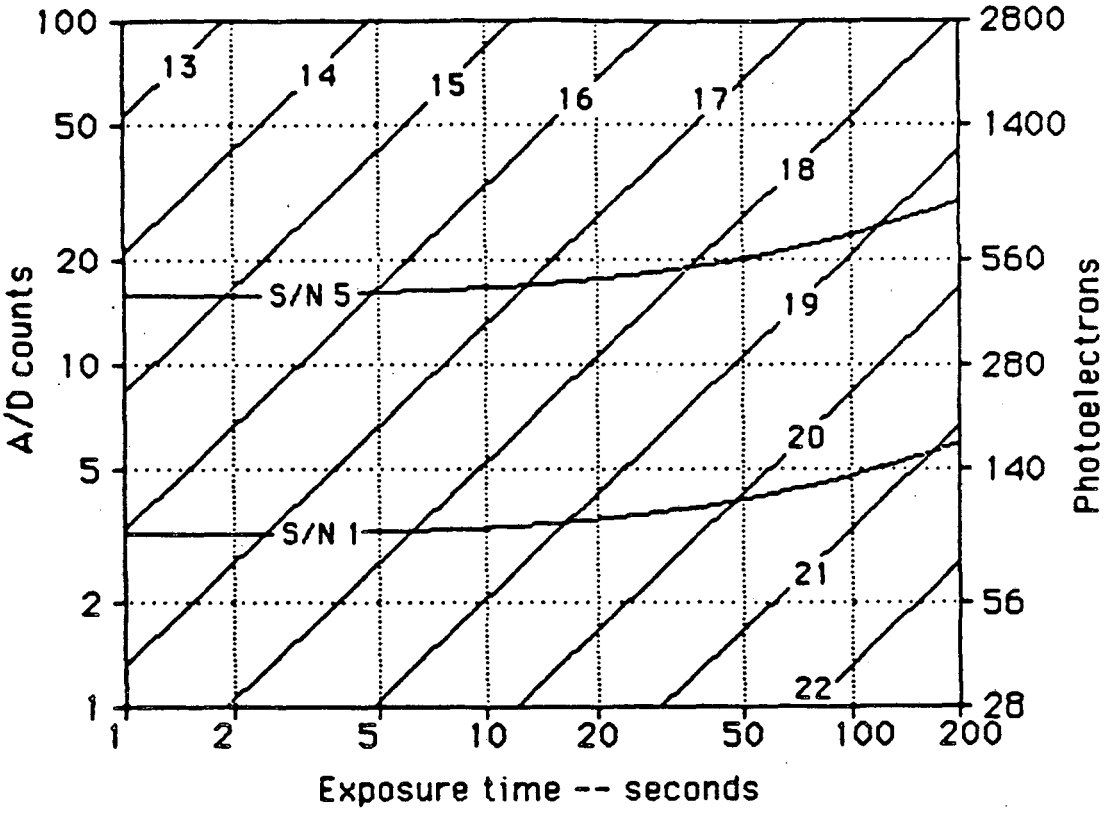

Figure 4.9b -- Current CCD (88 photoelectrons noise, quantum efficiency 30%) at MIRA telescope (36" aperture, sky background 100 p.e./sec/pixel).

supernova is plotted, along with the CCD noise level (signal to noise level of 1) and detection threshold (s/n of 5). f is assumed to be 0.25, except in figure 4.9e, where a larger pixel size and improved seeing at the MIRA telescope are assumed to give f = 0.5.

The conversion from n_{ss} to limiting magnitude shown in these plots depends on the spectrum of the supernova, the transmission of the entire optical system including the atmosphere, and the quantum efficiency of the CCD. We still have a factor of approximately 2 between the expected luminosity at the top of the atmosphere and

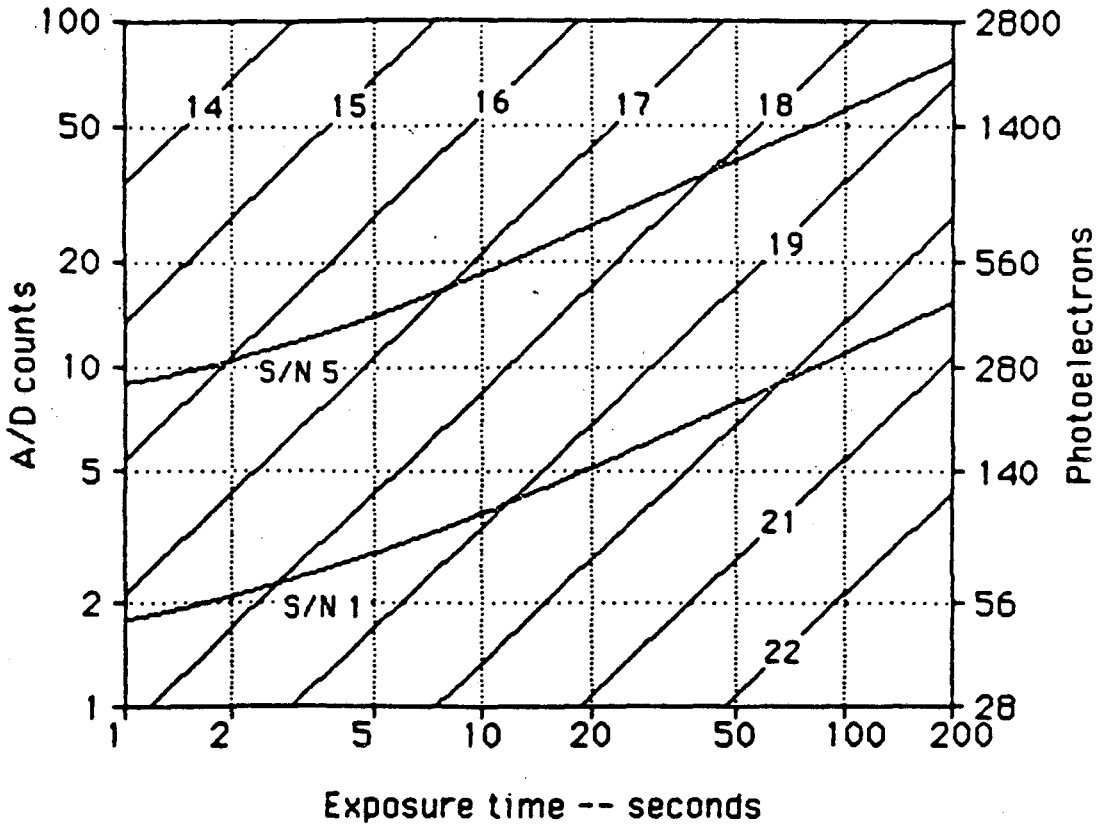

Figure 4.9c -- New CCD (Assumed noise 40 photoelectrons r.m.s., assumed quantum efficiency 70%) at Leuschner telescope (sky background 933 p.e./sec/pixel, due to higher quantum efficiency).

the observed signal at the CCD output. The most probable explanation for this difference is a variation in CCD quantum efficiency with wavelength, but we have not yet measured the quantum efficiency at other than 633 nm. Including this loss, and neglecting differences due to the extremely blue color of early time supernovae, the conversion factor is

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

This values is based on observations of standard stars of known visual magnitude using our current CCD at the Leuschner telescope (figure 4.9a). The conversion

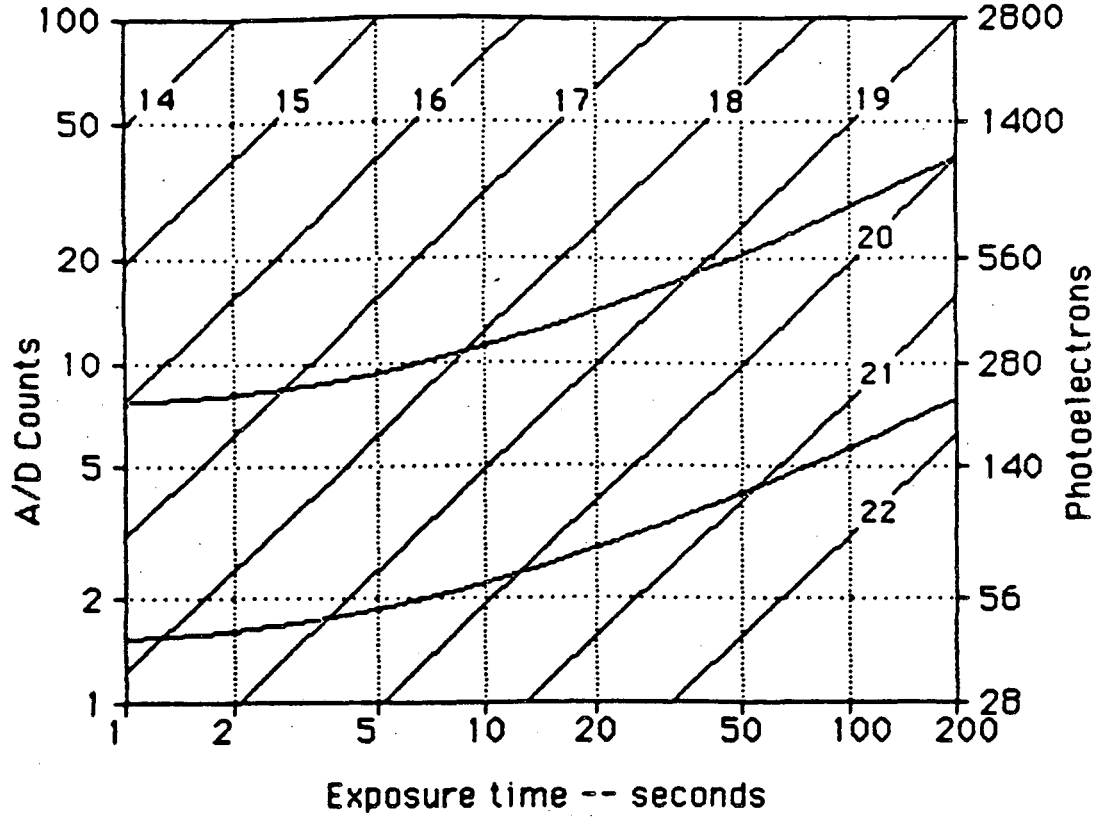

Figure 4.9d – New CCD (40 photoelectrons noise, 70% quantum efficiency) at MIRA telescope (sky background 233 p.e./sec/pixel).

factor for the other plots is derived from this by correcting for changes in telescope aperture, mean CCD quantum efficiency, and (in figure 4.9e) transmission of the optics and fraction f of light in one pixel. Improving the optics could gain us another factor of 1.5 in sensitivity – mostly through locating a suitable single camera lens and eliminating the Bronica/Nikon pair (50% total transmission).

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

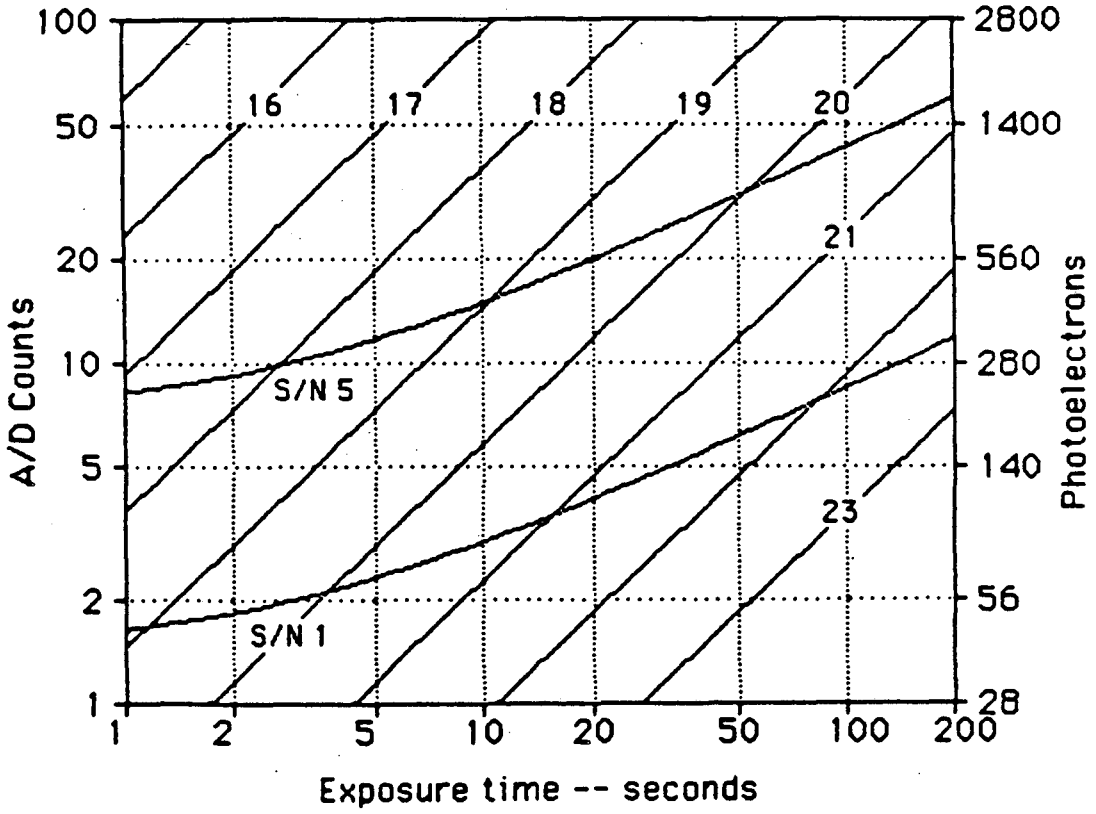

Figure 4.9e – Design Goal: New CCD (assumed noise 40 photoelectrons, assumed quantum efficiency 70%) at MIRA telescope with improved optics (75% transmission, f/2) and 3 arc second pixels, giving f (fraction of supernova light in one pixel) = 0.5, and sky background 538 p.e./sec/pixel.

rises more slowly. With our current CCD at Leuschner, we can reach 17th magnitude in approximately 30 seconds. An improved CCD would let us reach 17th magnitude in about 8 seconds, but with the telescope move time still 30 seconds the gain would be small – and reaching 18th magnitude would require nearly 50 second exposures. At MIRA, we would do only slightly better with the current CCD (15 seconds to 17th magnitude), but a lower noise CCD would be much more valuable, reaching 18th magnitude in under 10 seconds. Any of these values could be improved, however, by

improving the optics and increasing the pixel size (or getting better seeing) to raise the fraction f of light in one pixel.

4.3. Search Speed

4.3.1 Telescope Performance

The Leuchner telescope currently requires approximately 30 seconds to point over a typical intergalactic angle of two degrees. This includes sufficient time for the telescope to settle. Taking two CCD pictures, separated by an 8 second delay for readout (the second readout occurs while the telescope is moving), the total time spent per galaxy is approximately

$$t = t_{move} + t_{readout} + 2 \times t_{exp} = 38 \text{ sec} + 2 \times t_{exp}$$

If we are able to achieve a low enough false alarm rate to use only a single CCD picture, the time becomes

$$t = t_{move} + t_{exp} = 30 \text{ sec} + t_{exp}$$

Finally, the MIRA telescope can slew between galaxies in under 3 seconds. Figure 4.10 indicates that the settling time of the telescope is negligable; a star is trailed across the CCD using both Right ascention and declination stepping motors, then one motor is abruptly stopped. No overshoot or ringing is observable. We expect to achieve comparable performance at the Leuschner telescope shortly, using the upgraded drive system. In this case,

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

Figure 4.11 plots the number of galaxies observable per (6-hour) night versus t_{exp} for these three cases, plus an intermediate case assuming a 6 second move time.

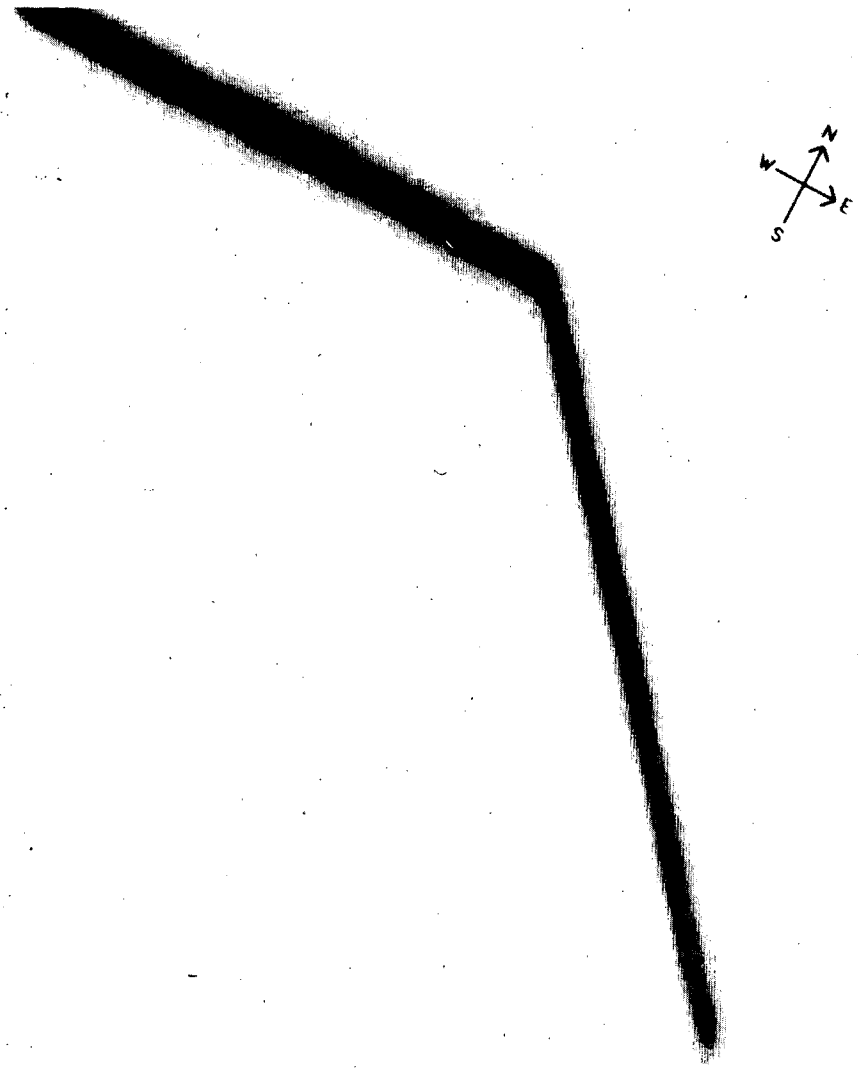

Figure 4.10

Telescope Skid Marks

Illustration of rapid settling time of MIRA telescope: Telescope moves northwest with both RA and DEC stepper motor drives operating at approx. 150 arc sec/sec (50 pixels/sec). DEC drive stopped abruptly; no visible vibration on a time scale longer than a few pixels (100 msec.) with an amplitude greater than approx. 1 pixel (3 arc sec).

XBB840-8737

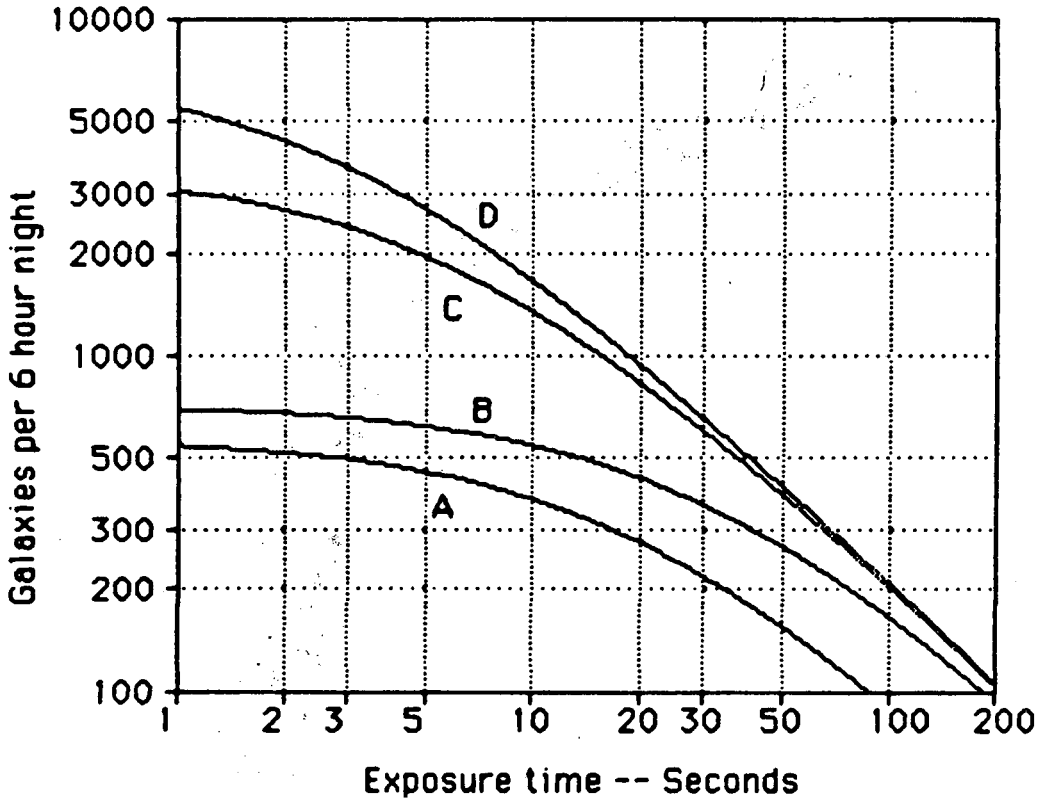

Figure 4.11 -- Observable galaxies per night vs. exposure time

A. 2 exposures per galaxy, 30 second move time + 8 second readout time.

B. 1 exposure, 30 second move time

C. 1 exposure, 30 second move time

D. 1 exposure, 3 second move time (videotape recorder buffered)

4.3.2 Other Considerations

4.3.2.1 CCD Readout

We have read out our CCD in times as short as 3 seconds (20 microseconds/pixel), with no conspicuous ill effects, and have been informed that readouts as fast as 1 microsecond/pixel can be achieved without increasing the readout noise [32]. The videotape recording system is currently limited to a max-

imum readout rate of about 25,000 pixels per second (6.4 second minimum readout time); however, by using a suitable buffer memory, the CCD data can be stored and transferred to videotape during the following exposure as well as during the telescope move time, so this is not a major limitation.

4.3.2.2 Computer Time

The total computer time required for one image cycle is currently about 60 seconds for a 100 x 100 pixel galaxy subimage. Nearly half of this time is spent in the blurring and jogging subroutines. By revising the algorithm used in this routine, and by simple optimizing of other routines (for example, minimizing disk accesses by buffering reference files in memory) we expect to reduce this time to under 30 seconds. However, achieving a mean processing time per galaxy of under 8 seconds (including CCD readout) will certainly require careful hand-coding of several routines. Alternatively, several hardware-based routes to faster processing are now available at modest cost, ranging from 68000-based coprocessor boards and array processors to custom-built dual-port memories (eliminating DMA overhead) to, as an extreme example, upgrading our system to a VAX-11/750, with at least a doubling of processor performance.

4.4. 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 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. An additional confusion is introduced by the ability of the Berkley search to select individual galaxies to search, and thus to concentrate on relatively massive galaxies, or on specific types of galaxies which have been found to produce large numbers of supernovae.

A relatively subtle point is that, in trying to detect supernovae at very early times, there is no inherent advantage in scanning one set of galaxies frequently. If one can examine n galaxies per night, one option is to scan the same n galaxies each night, obtaining, say, n/100 supernovae per year with an average delay (between a supernova crossing the detection threshold and the actual detection) of 12 hours. The number detected within a time t of crossing threshold is then (n/100)(t/24 hours) per year. However, one can equally well scan 2n galaxies every 2 nights, detecting twice as many supernovae per year. The average delay is now 24 hours, but the number detected within time t is (2n/100)(t/48 hours) – the same number. Unfortunately, one soon runs out of nearby galaxies to search. The maximum range for detection of type I supernovae at a fraction f_{max} of maximum light is roughly

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

For $M_v = -19.4$

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

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

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

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

where d_{gal}, the density of galaxies, is roughly 0.02 per Mpc^3 [33], and f_{sky}, the fraction of the sky accessible from our site at one time, 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). This is somewhat above the number that can be conveniently searched in 2 to 3 nights, so we select galaxies by brightness or (where available) red shift (z) value.

In practice, we have generated our lists of galaxies from standard catalogs (primarily the Nilson catalog [34]) available on magnetic tape. 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. A cut in apparent integrated visual magnitude is made to select the brightest (and thus nearest and/or largest) galaxies, with the cutoff adjusted to yield a suitable total number of galaxies.

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 also expect the time spent per galaxy to decrease; for the foreseeable future, we will essentially run out of accessible galaxies (at f_{max}=0.2) in a 3 to 4 day cycle. In order to increase the chances of very early detections, we will eventually scan a selected set of nearby galaxies every night.

5. Current Status, Summary, and Conclusions

5.1. Current Status

The Berkeley supernova search is now collecting reference galaxy pictures at the Leuschner telescope as a secondary task, while gathering data for the solar companion search. Processing of this data, and completion of the fully automated version of the image cycle software, are planned following the completion of this paper.

Once the first data set for the solar companion search is complete, searching for supernovae will begin at a rate of approximately 300 galaxies per night, searching 1200 galaxies on a 4 night cycle. Two 30-second exposures per galaxy will be used, giving a supernova detection limit of 16th magnitude. As soon as sufficient false alarm statistics are available to warrant switching to single 30 second exposures, the search rate will be increased to approximately 500 galaxies per night. With 2000 galaxies searched, we expect to find one to two supernovae per month of operation.

As various improvements are made, including replacing the CCD, replacing the current optics, rebuilding the Leuschner telescope drives, and possibly moving to the new MIRA observatory, the search rate will be increased and the detection threshold lowered. Full real time processing (computer at the telescope site) is some time away; full sensitivity will depend on acquiring a better CCD chip and on many small improvements to the system.

5.2. Summary and Conclusions

The Berkeley automated supernova search is an operating system for automated astronomy. It is now producing reference data for both the supernova search and a search for a solar companion. 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.

6. Possibilities for Future Work

The Berkeley supernova search system is capable of doing much more than finding supernovae. It can collect additional information on the supernovae found, or further analyse the data collected in the search. In addition, with relatively minor changes, it can search for other types of objects, or investigate any time-varying astronomical phenomenon of interest.

6.1. Additional Data Collection.

The simplest additional measurement possible is the tracking of supernova light curves. Since the foreground fiducial stars provide a reference for photometry, the only effort required is to maintain a list of supernovae, and take extended exposures of the appropriate galaxies to improve the signal to noise ratios. Photometric accuracies of 10% should be obtainable down to 1-2 magnitudes fainter than the detection threshold. We expect to follow light curves routinely in the future.

By adding color filters in front of the CCD, we can do color photometry on supernovae. Currently, filters must be inserted manually, but an automatic filter wheel can be added easily. Absolute calibration of sensitivity will require frequent measurement of reference stars, but relative measurements can still be made by using fiducial stars as standards. Low dispersion spectroscopy could be done by inserting a grating into the optical path, but the system sensitivity would be very low. Spectroscopy can be done more easily at other telescopes, and is not a priority for the supernova search.

6.2. Additional Data Analysis

The supernova search fields include large numbers of foreground stars, as well as the target galaxy. Both the galaxy and the foreground stars are of only peripheral

interest for the search itself, but can be studied independently. The data collected are appropriate for studying time variability on timescales from a few days to a few months. Variable stars close to galaxies must be catalogued in any case, to minimise false alarms; it is straightforward to record the brightnesses of other interesting objects in each CCD field.

It is also possible that certain galaxy cores are variable at a low level [35]. By averaging large numbers of CCD pictures, we can measure variations in galaxy core brightness to high accuracy.

6.3. Other Automated Searches

The Berkeley automated supernova search is one of the first systems designed for automated astronomy – completely automated processing of large numbers of images in real time. This technology can be used to search for other things besides supernovae. We mention here two particularly interesting possibilities: a solar companion star search and a search for extrasolar planets.

6.3.1. A Search for a Solar Companion Star

Davis, Hut, and Muller [36] have recently proposed that the Sun posesses a companion star in a 28 million year orbit; the current distance to this star is roughly 2.4 light years. If this companion is heavier than approximately 0.08 solar masses, it should be visible as a dim red star. It is possible to separate a very nearby red star from more distant (and more numerous) red giants by looking for parallax as the earth moves around the sun. The expected parallax shift is approximately 3 arc seconds peak to peak. This motion is large enough to be detected easily on a CCD image.

The Berkeley supernova search has already been pressed into service searching for a solar companion. Because of the small area covered by the CCD, it is necessary to select target stars for study; we have generated a list of roughly 5000 faint red stars from the Dearborne Observatory catalog [37]. Exposures centered on the target star typically include 5 or more fiducial stars; the position of each star can be found with an accuracy of better than .05 pixels, so parallaxes of less than 1 arc second can be detected. Searching the target list is now underway, and we expected to scan the entire accessible portion of the list in 10 to 15 nights -- after which we must wait 1 to 2 months before scanning again.

By reading out the CCD continuously while scanning across the sky at a matching rate, one can scan a continuous strip of sky very efficiently. This drift-scan technique [38] can be used to cover the entire sky in a search for a solar companion (or anything else), but CCD's and Schmidt plates yield comparable performance in such a full sky search. Also, the volume of data to be processed is very large. However, such a full sky search may yet prove worthwhile.

6.3.2. A Search for Extra-Solar Planets

One way of locating extra-solar planets is by searching for transits -- temporary reductions in the apparent luminosities of stars [39]. The reduction in luminosity is approximately 1% for a Jupiter-sized planet passing in front of a sunlike star. Schemes based on this observation have been rejected as impractical because of the large number of stars that would have to be surveyed -- the probability that we lie exactly in the orbital plane of a distant planet is quite low. However, a CCD can provide photometric measurements on many stars simultaneously, and a system such as ours (or a suitable drift-scan system) can examine many star fields per night. If sufficient photometric accuracy ($<1\%$ error) can be achieved, a planet search system

capable of monitoring tens or hundreds of thousands of stars may be feasible. Data on false alarms and photometric accuracy from the Berkeley search will be valuable in designing such a system.

REFERENCES

[1] J. B. Oke and L. Searle, ANN. REV. ASTRON. ASTROP. 12, 315 (1974).

[2] R. Barbon, F. Ciatti, and L. Rosino, *On the Light Curve of Type I Supernovae,* in SUPERNOVAE AND SUPERNOVA REMNANTS, C. B. Cosmovici, ed. (D. Reidel, Dordrecht, 1974).

[3] G. A. Tammann, *Supernova Statistic and Related Problems,* in SUPERNOVAE: A SURVEY OF CURRENT RESEARCH, M. J. Rees and R. J. Stoneham, eds., (D. Reidel, Dordrecht, 1982).

[4] R. Barbon, F. Ciatti, and L. Rosino, *Recent Observations of Supernovae at Asiago,* in SUPERNOVAE AND SUPERNOVA REMNANTS, C. B. Cosmovici, ed. (D. Reidel, Dordrecht, 1974).

[5] J. C. Wheeler, *Theoretical Models of Type I Supernovae,* in SUPERNOVAE: A SURVEY OF CURRENT RESEARCH, M. J. Rees and R. J. Stoneham, eds., (D. Reidel, Dordrecht, 1982).

[6] D. Branch, *Type I Supernovae -- Observational Constraints,* in SUPERNOVAE: A SURVEY OF CURRENT RESEARCH, M. J. Rees and R. J. Stoneham, eds., (D. Reidel, Dordrecht, 1982).

[7] T. Axelrod, *Late Time Optical Spectra from the ^{56}Ni Model for Type I Supernovae,* in PROC. TEXAS WORKSHOP ON TYPE I SUPERNOVAE, ed. J. C. Wheeler (U. of Texas Press, Austin, 1980).

[8] S. E. Woosley and T.A. Weaver, *Theoretical Models for Supernovae,* in SUPERNOVAE: A SURVEY OF CURRENT RESEARCH, M. J. Rees and R. J. Stoneham, eds., (D. Reidel, Dordrecht, 1982).

[9] G.E. Brown, *The Fate of Massive Stars: Collapse, Bounce and Shock Formation,* in SUPERNOVAE: A SURVEY OF CURRENT RESEARCH, M. J. Rees and R. J. Stoneham, eds., (D. Reidel, Dordrecht, 1982), and following papers.

[10] G. A. Tammann, op. cit.

[11] F. Zwicky, *Some Results of the International Search for Supernovae,* in SUPERNOVAE AND THEIR REMNANTS, Brancazio & Cameron, eds., (GordonAnd Breach, 1969).

[12] D. H. Clark and F. R. Stevenson, *The Historical Supernovae,* in SUPERNOVAE: A SURVEY OF CURRENT RESEARCH, M. J. Rees and R. J. Stoneham, eds., (D. Reidel, Dordrecht, 1982).

[13] Ibid

[14] G. Muller and E. Hartwig, GESCHICHTE UND LITERATUR DES LICHTWECHSELS II, p. 417 (Leipzig: Peschel & Trepte, 1920).

[15] F. Zwicky, op. cit.

[16] W. L. W. Sargent, L. Searle, and C. T. Kowal, *The Palomar Supernova Search,* in SUPERNOVAE AND SUPERNOVA REMNANTS, C. B. Cosmovici, Ed., (Reidel, Dordrecht, 1974).

[17] L. Rosino, *Observations of Supernovae at the Astrophysical Observatory of Asiago,* in SUPERNOVAE, D. Schramm, ed. (Reidel, Dordrecht, 1977).

[18] J. Maza, in PROC. TEXAS WORKSHOP ON TYPE I SUPERNOVAE, ed. J.C. Wheeler, p. 7 (U. Texas, Austin 1980).

[19] G. A. Tammann, op. cit.

[20] J. R. Dunlap, J. A. Hynek, and W. T. Powers, in ADVANCES IN ELECTRONICS AND ELECTRON PHYSICS, J. B. McGee, D. McMullan, and E. Kany, eds., 33B, 789 (1972).

[21] 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).

[22] 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).

[23] S.A. Colgate, E.P. Moore, and R. Carlson, PUBL. ASTRON. SOC. PACIFIC, 87, 565 (1975).

[24] T. A. Weaver, T. S. Axelrod, and S. E. Woosley, *Type I Supernova Models vs. Observations,* in PROCEEDINGS OF THE TEXAS WORKSHOP ON TYPE I SUPERNOVAE, J. C. Wheeler, ed., (U. Texas, Austin, 1980).

[25] W. D. Arnett, *The Nature of Supernovae as Determined from Their Light Curves,* in SUPERNOVAE: A SURVEY OF CURRENT RESEARCH, M. J. Rees and R. J. Stoneham, eds., (D. Reidel, Dordrecht, 1982).

[26] R. Barbon, F. Ciatti, and L. Rosino, *On the Light Curve of Type I Supernovae,* in SUPERNOVAE AND SUPERNOVA REMNANTS, C. B. Cosmovici, ed. (D. Reidel, Dordrecht, 1974).

[27] J.D.E. Beynon and D.R. Lamb, CHARGE COUPLED DEVICES AND THEIR APPLICATIONS, (McGraw Hill, 1980).

[28] 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).

[29] J. Culver, *A CCD Camera for an Automated Supernova Search,* Senior Honors Thesis, University of California, Berkeley (1982).

[30] 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).

[31] J. R. Janesick, Jet Propulsion Laboratory, private communication.

[32] R. Aikens, private comm.

[33] C. W. Allen, ASTROPHYSICAL QUANTITIES, 3rd Edition (Athlone, London, 1973), p. 289.

[34] P. Nilsen, UPPSALA GENERAL CATALOG OF GALAXIES, (Uppsala, 1973).

[35] M. Gorenstein, private comm.

[36] M. Davis, P. Hut, and R.A. Muller, *Extinction of Species by Periodic Comet Showers,* NATURE, 308, #5961, 715-717, 19 April 1984.

[37] T. A. Nagy, *Documentation for the Machine Readable Version of the Dearborne Observatory Catalog of Faint Red Stars,* SASC Report R-SAW-8/79-01 (Systems and Applied Sciences Corp., Riverdale, MD 1979).

[38] 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).

[39] W.J. Borucki and A. L. Summers, *The Photometric Method of Detecting Other Planetary Systems,* ICARUS 58, 121-134 (1984).

APPENDIX I

A CCD CAMERA FOR AN AUTOMATED SUPERNOVA SEARCH

Senior Honors Thesis by
John Culver

Uncopyrighted material
Included for reference only
Used by permission of author

A CCD CAMERA FOR AN AUTOMATED SUPERNOVA SEARCH

In this paper, I shall describe an automated optical imaging system using a Charge-Coupled Device imaging array. I will first briefly describe the application to which this camera will be put: an automated search for supernova explosions in external galaxies. Next, I will review the physics of the charge coupled effect, and describe the actual imaging device, comparing it to other imaging devices currently available. Finally, I will focus on the design of the support equipment necessary to produce a working imaging system that will interface with a computer image processing system.

AN INTRODUCTION TO THE SUPERNOVA SEARCH

In a typical galaxy, a supernova occurs, on the average, every 100 years. At the end of its life, a star of 4 to 20 solar masses can explode so violently that it can be as bright as an entire galaxy. Many such supernovae have been detected in the past, in fact as far back as 185 B.C. Many have been studied with modern spectroscopic instruments, but practically none have been caught before they reached their maximum brightness.

This is significant because most theories predict that heavy element synthesis occurs in the early stages of supernovae. Spectrographic measurements on supernovae in their early stages can tell us much about the synthesis of heavy elements. In addition, the availability of data on large numbers of supernovae is useful for the study of supernova statistics, and the use of supernovae as standard candles in distance measurements.

The supernova search is essentially an updated version of Stirling Colgate's systematic search for supernovae. A 36 inch diameter telescope is used to scan 2,500 galaxies each night. A picture is taken of each galaxy, and compared to the picture from a previous night. If this comparison indicates that a significant increase in brightness of a single picture element has occured, the observers are signaled that a supernova has been found. The expected detection rate is one supernova per 3 nights of running the experiment.

The scanning function of the telescope and the image comparison is controled by a PDP 11-44 computer. All operations are performed in real time. While the telescope moves to the next galaxy and takes its picture, the previous image is read into the computer and compared to the standard image. This calls for a fast imaging device and

efficient use of the computer.

The CCD imaging system provides this capability. The CCD is an efficient low light-level imager, has very low noise, good transfer efficiency, and wide dynamic range. The output is indexed for convienent correction of device flaws, and is ideally suited for digitizing and correlated double sampling. This digitizing capability makes it possible to input the image data directly to the computer memory through a direct memory access board, saving considerable time, and freeing the computer to work on other things.

While CCD's have been available since the mid-1970's, the technology for producing a low noise, very efficient, television-type imaging array has only recently been developed. CCD technology is blossoming, so the time is right for an experiment of this sort.

AN INTRODUCTION TO THE CHARGE-COUPLED DEVICE

A. The Physics of CCDs

1. The MOS Capacitor

The CCD imaging array is essentially an array of metal oxide semiconductor (MOS) capacitors. The capacitors are used to form potential wells in which photo-generated electrons are trapped. By modulating the voltage on the capacitors, the charge packets can be made to shift across the array. The packets are then detected at one end of the device, and used to generate an image. We will see that the depth of the potential well is set initially by the voltage on the capacitor, and then varies linearly with the amount of charge stored there. This leads to a simple hydraulic model (the "bucket brigade") for the charge storage and transfer which will prove very useful in visualizing what is happening in the CCD.

The analysis in this section follows that given in Part I of Beynon and Lamb (see REFERENCES) Consider the MOS capacitor shown in Figure 1. A layer of silicon oxide is grown over a lightly doped p-type substrate. An electrode is added over the oxide layer. If we apply a positive voltage to this "gate" electrode, the holes (the majority cariers in the substrate) are repelled from the interface creating a depletion region underneath the electrode. Further increase in the gate voltage beyond a threshold

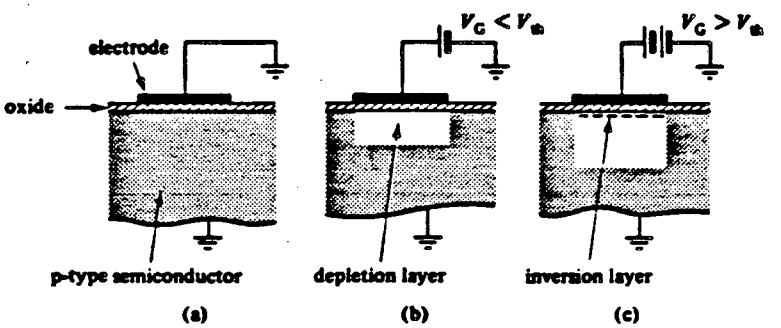

Single CCD electrode showing the creation of depletion and inversion layers under the influence of an increasingly positive electrode voltage.

Figure 1

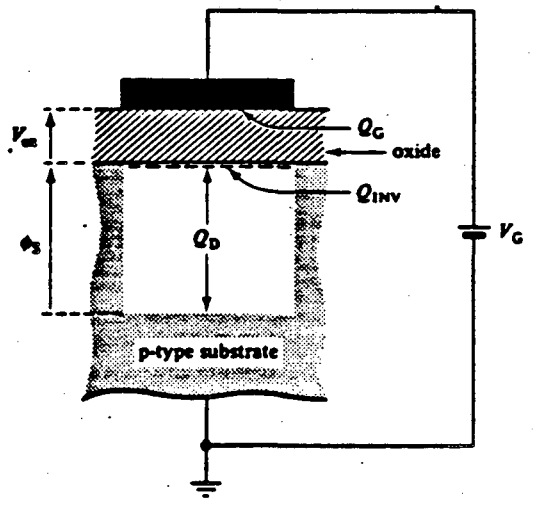

Potentials and charges associated with a single CCD electrode. For the p-type substrate shown, V_G is positive whilst Q_{INV} and Q_D are negative; for an n-type substrate the signs would be reversed.

Figure 2

voltage V_{th} drives the device into deep depletion (Figure 1(c)). This produces a surface potential so great that minority carriers (electrons) are attracted to the interface. This is called an inversion layer. There are two main ways to form an inversion layer. Thermal energy in the substrate can excite an electron from the valence band into the conduction band. The electron half of this thermally generated electron-hole pair is attracted to the interface. In time, enough of these thermal electrons are attracted to the interface to put the system into equilibrium. The time required for this to occur is called the thermal relaxation time of the CCD and is on the order of several minutes for a device at liquid nitrogen temperatures.

For imaging applications, electrons are made available for the inversion layer by incident photons exciting electron-hole pairs. Since the characteristic times for light integration and device read-out are much less than the thermal relaxation time, the inversion layer will hold a unique signal generated by the photoelectrons. If the image were held for longer than the thermal relaxation time, thermal electrons would reduce the contrast in the image.

By driving the CCD into deep depletion, the inversion layer can hold the charge packets generated by the photoelectrons. These electrons will be attracted to the high surface potential at the substrate-oxide interface, since this is the region of lowest potential energy for negative charges. This accumulation of negative charge decreases the surface potential, and shrinks the depletion layer. We can derive a relationship between the surface potential and the charge accumulated in the inversion layer.

Referring to Figure 2 we see that

$$V_G = V_{OX} + \phi_s \quad .$$

Since the system has no net charge, we require that the charge densities sum to zero:

$$Q_G + Q_{INV} + Q_D = 0 \quad .$$

Now we use the depletion approximation. Let N_A be the acceptor atom number density, and ϵ_s be the permittivity of silicon. Let the distance variable x extend from the bottom of the oxide layer into the substrate. Solving Poisson's equation

$$\nabla^2 \phi(x) = - e \frac{N_A}{\epsilon_s}$$

with the boundary condition $\phi(x_D) = 0$ gives us

$$\phi(x) = - e \frac{N_A}{2\epsilon_s} (x^2 - x_D^2)$$

$$\phi_s = \phi(0) = e \frac{N_A}{2\epsilon_s} x_D^2$$

$$x_D = \left(\frac{2\phi_s \epsilon_s}{eN_A} \right)^{1/2} .$$

The charge density created in the substrate by driving it into depletion to a distance x_D is given by

$$Q_D = - eN_A x_D$$

$$= - eN_A \left(\frac{2\epsilon_s \phi_s}{eN_A} \right)^{1/2}$$

$$= - (2eN_A \epsilon_s \phi_s)^{1/2}$$

Let C_{OX} be the capacitance per unit area of the oxide layer. Then,

$$V_{OX} = \frac{Q_G}{C_{OX}}$$

$$V_G = \frac{Q_G}{C_{OX}} + \phi_s$$

$$= \frac{-1}{C_{OX}} (Q_{INV} + Q_D) + \phi_s$$

$$= \frac{- Q_{INV} + (2eN_A \epsilon_s \phi_s)^{1/2}}{C_{OX}} + \phi_s$$

Rearranging the last equation yields

$$\phi_s = V_G + V_O + \frac{Q_{INV}}{C_{OX}} - \left[2 \left(V_G + \frac{Q_{INV}}{C_{OX}} \right) V_O + V_O^2 \right]^{1/2}$$

where $$V_O = \frac{e\epsilon_s N_A}{C_{OX}^2} .$$

Here are some typical values for these constants:

$$N_A = 10^{15} \text{ cm}^{-3}$$

$$\epsilon_s = 1.14 \times 10^{-12} \text{ f/cm}$$

$$C_{OX} = 1.14 \times 10^{-8} \mu\text{f}$$

$$V_O = 0.14 \text{ V}$$

$$V_G = 10 \text{ V}$$

The small value for V_O suggests an approximation:

$$\phi_s = V_G + \frac{Q_{INV}}{C_{OX}} \quad .$$

This is the fundamental equation for the surface-channel CCD. It suggests a simple hyraulic model. Just as the depth of a bucket of water (as measured from the rim to the surface of the water) varies linearly with how much water is poured in, the depth of the potential well varies linearly with the number of electrons in the inversion layer. These electrons "fill up the well".

We can now understand how the CCD array can be used as an imaging device. Imagine a photon incident on the substrate. If it has sufficient energy , it can excite an electron out of the valence band into the conduction band. This is a photo-generated electron-hole pair. The absorption spectrum should therefore be a continuum of strong absorption at high energies bounded by a steep drop-off in absorption below the band-gap energy.

The absorption spectrum for silicon at room temperature is shown in Figure 3. This shows that for the visible part of the spectrum (1.8 to 3.0 eV) silicon is a useful detector. Note that the absorption in this region is relatively independent of the photon energy. This indicates that the number of electron-hole pairs produced is also insensitive to photon energy. Thus the number of electrons is proportional to the number of photons and hence to the incident intensity of radiation. These photo-electrons are atracted to the local minimum in potential energy, so they essentially "fall" into the nearest well. An array of these wells will then contain charge packets each of which has a size proportional to the local intensity of light. Thus the information stored in the charge packets is identical to that which would be stored on a piece of photographic film. The array contains an image.

2. Charge Transfer

We would like to read out the charge in each well in some indexed fashion so that we may reconstruct the image. This is accomplished by adding two more wells for each

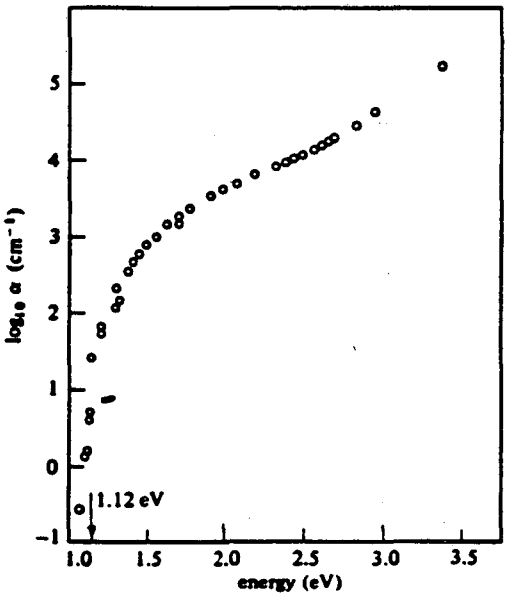

Absorption spectrum of silicon at 300 K (Dash and Newman (1955)).

Figure 3

one that stores the original packet. These extra wells can be used to transfer the charge in a "bucket brigade" across the CCD. Figure 4 shows a section of a linear array of electrodes. 4(a) shows the potential well containing the photo-electrons, and three surrounding wells. All the gate bias voltages are set greater than the threshold voltage, but only one is biased into deep depletion. This situation is maintained during light integration. At the end of the integrating period we read out the stored charge packets by first driving one of the adjacent wells into deep depletion. The charge packet distributes itself evenly across both wells as shown in 4(b-c). The original well can now be collapsed, producing a spatial shift in the charge packet one electrode to the right, as in 4(d-e).

In this way, the the charge packets can be propagated in a controlled manner across the array simply by applying the correct sequence of gate voltages. The three well structure enables us to connect every third electrode to the same clocking waveform, producing a device that will transfer charge unambiguously in one direction without having to have individual access to each electrode. Three is the minimum number of electrodes for the type of device shown in Figure 4, which is called a <u>three phase CCD</u>. The waveforms for clocking this device are shown in 4(f).

The output end of the CCD must contain something that can detect the charge packet,and reset to prepare for the next. Figure 5(a) shows a schematic of the on-chip detecting electronics. The output diffusion is set initially to a reference voltage RD by turning the reset transistor on with a positive pulse to the reset gate RG. After the reset transistor has been turned off, the charge packet is dumped on the output gate OG which is held at a fixed bias. This produces a voltage variation at the gate of the amplifier transistor which is operated in the source-follower configuration. Figure 5(b) shows the clocking waveforms and the output voltage; $\Delta V \propto Q_{INV}$.

Now we can see how a CCD image is read out. Figure 6 depicts the image transfer for a <u>frame transfer device</u> similar to the one used in the Supernova Search. Photoelectrons are stored in the image section potential wells. When light integration is complete, the image is transferred vertically by one set of three-phase vertical clocks into the image storage section. From there, each horizontal line of the image is transferred to a horizontal readout register by another set of three-phase vertical clocks. Each of these lines is then transferred horizontally to the output port by the horizontal clocks. In this way, each charge packet is indexed to its place in the array by how many vertical and horizontal transfers have taken place. The output voltage is

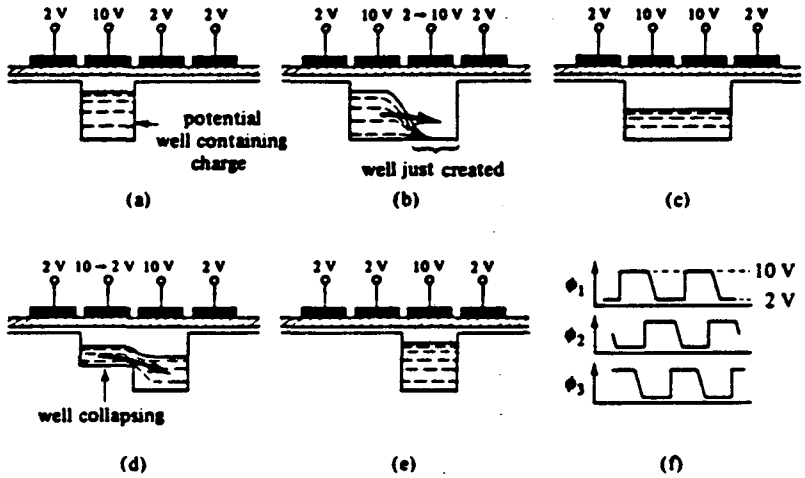

(a)–(e) Movement of potential well and associated charge packet by clocking of electrode voltages: (f) clocking waveforms for a three-phase CCD.

Figure 4

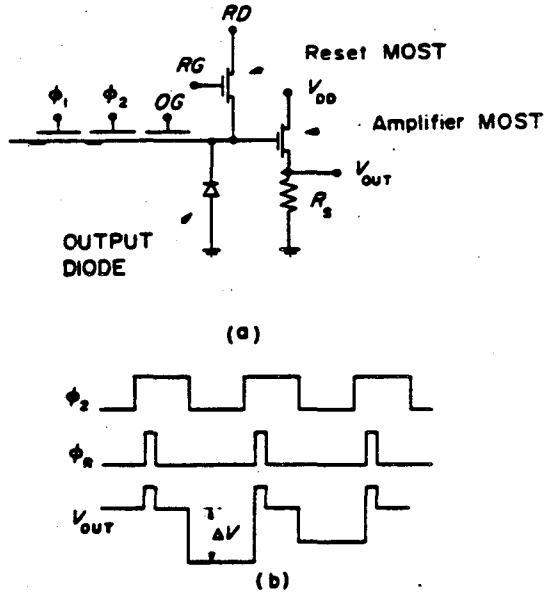

(a) Circuit diagram of a gated charge integrator; (b) clocks to operate the circuit and resulting output wave form

Figure 5

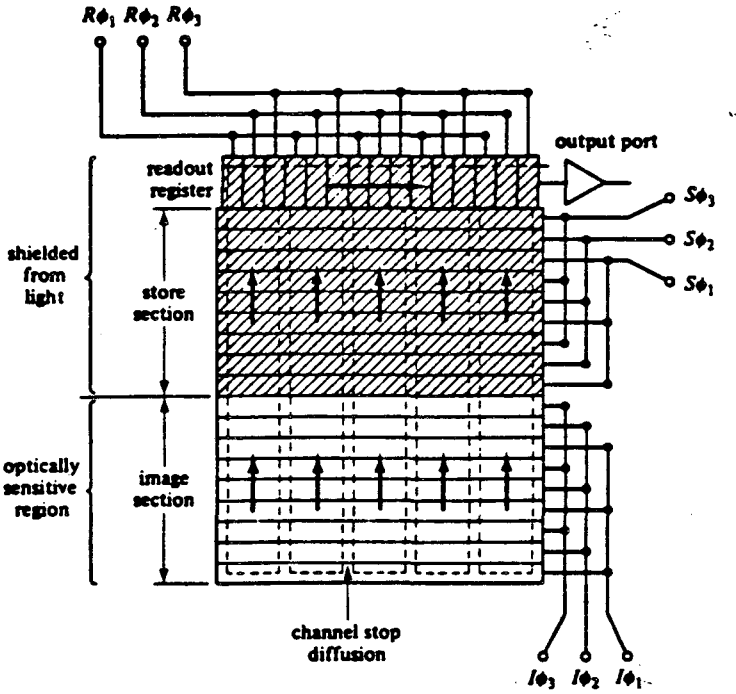

Frame transfer array schematic.

Figure 6

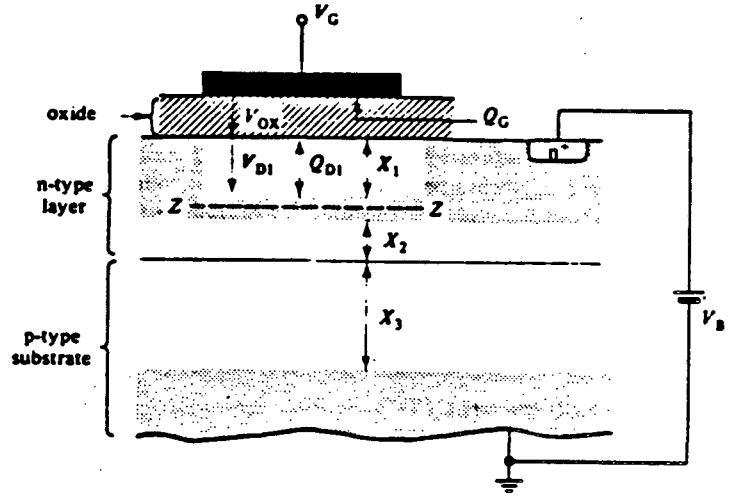

Section through part of a BCCD showing formation of field-induced and junction depletion regions.

Figure 7

amplified further, and the noise is reduced through correlated double sampling. The signal is digitized and read into the computer.

3. Refinements and Optimization

A few refinements to the surface channel CCD should be mentioned at this time. The CCD used for the Supernova Search is an RCA 53612-X0. It is a thinned-back, buried channel CCD. A buried channel CCD is similar to the surface channel device, but the charges are not stored just under the oxide layer. Instead, they are stored at a potential minimum which is burried in the substrate. Figure 7 shows a diagrammatic representation of the burried channel device. The main difference is that the burried channel device has a small n-type layer deposited over the p-type substrate. The doping level in the n-type silicon is made greater than that in the substrate. Following the same type of analysis used for the surface channel device, one can show that the potential minimum will lie somewhere within the n-type region. This device will retain the important characteristics of the surface channel device (like the linearity of the well-depth/charge relation) but avoid some of the disadvantages (like fast surface states limiting of transfer speed). They have greater transfer efficiency and lower intrinsic noise as well. The disadvantages are relatively minor. Fabrication of the device is more complex. It has a smaller signal handling capability, and is more sensitive to thermal generation of electrons. These are not significant for the Supernova Search, as the device is cooled and the light-levels are low.

An ordinary "thick" CCD must make the electrodes transparent to visible radiation. The photons must traverse the electrode and oxide regions before producing electron-hole pairs. The thinned-back device avoids this problem by making the substrate as thin as possible and having the light impinge on that side. This technique significantly improves the quantum efficiency of the device, as shown in Figure 8.

B. Manufacturers

CCD's are now available from several manufacturers. Most of these are also conducting research and development of CCD's. Here is a partial list:

Bell Labs
Fairchild

General Electric
Intel
RCA

CCD CAMERA DESIGN

A. Description of CCD Camera for Supernova Search

A block diagram for the supernova detection system hardware is shown in Figure 9. Light from a galaxy is focused onto the CCD array. Signals from the CCD are amplified within the dewar and sent to the CCD driving electronics. The electronics provide the clocking and signal processing for the camera. The image is then sent to the computer for further processing in software, and comparison to previous images of the galaxy. The computer also controls the aiming of the telescope.

In this section I would like to discuss electronics in detail, but first I would like to give an overview of what the electronics is supposed to do. Figure 10 shows a functional diagram of the electronics. Correlated Double Sampling is a reset-noise reduction techique which effectively eliminates the reset noise from the signal The signal is then converted to serial digital information by an A/D converter and sent to the computer.

When the computer is ready to accept the data, it sends a signal to the timing boards. These boards produce the basic clock phases for CCD readout of an image. They also provide clocking for the A/D converter and transmission of data to the computer. The clock waveforms are buffered and shaped by analogue circuitry to provide the CCD with phases that produce efficient charge transfer. These are then sent to the CCD.

There is some additional circuitry for temperature regulation and protection of the CCD from voltage and current spikes, but this circuitry is not too important for our discussion. We will now look at the electronics in more detail.

B. The Electronics

1. Output Clocking and Timing

The schematics for the Digital Timing Boards are shown in Figures 11(a) and

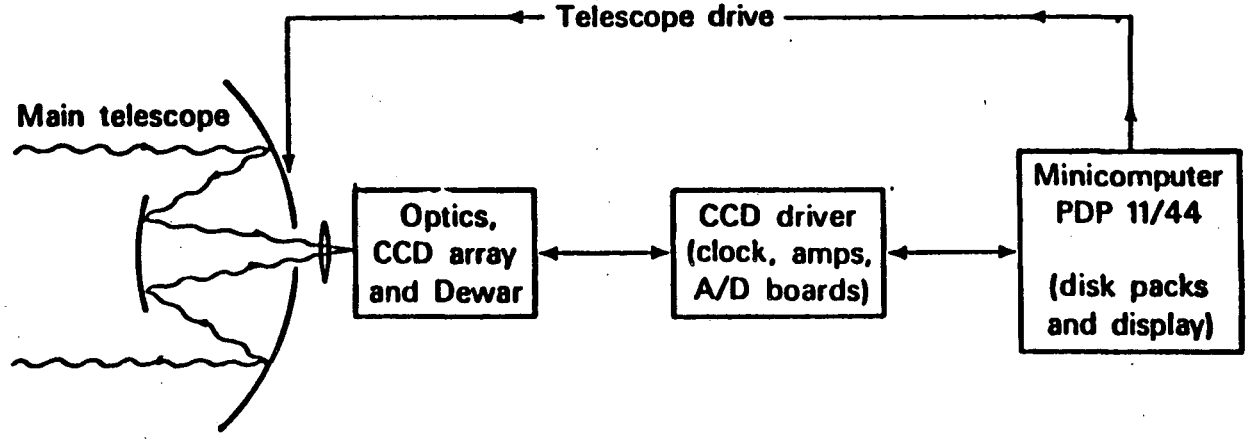
Telescope drive
Main telescope
Optics,
CCD array
and Dewar
CCD driver
(clock, amps,
A/D boards)
Minicomputer
PDP 11/44
(disk packs
and display)

Figure 9

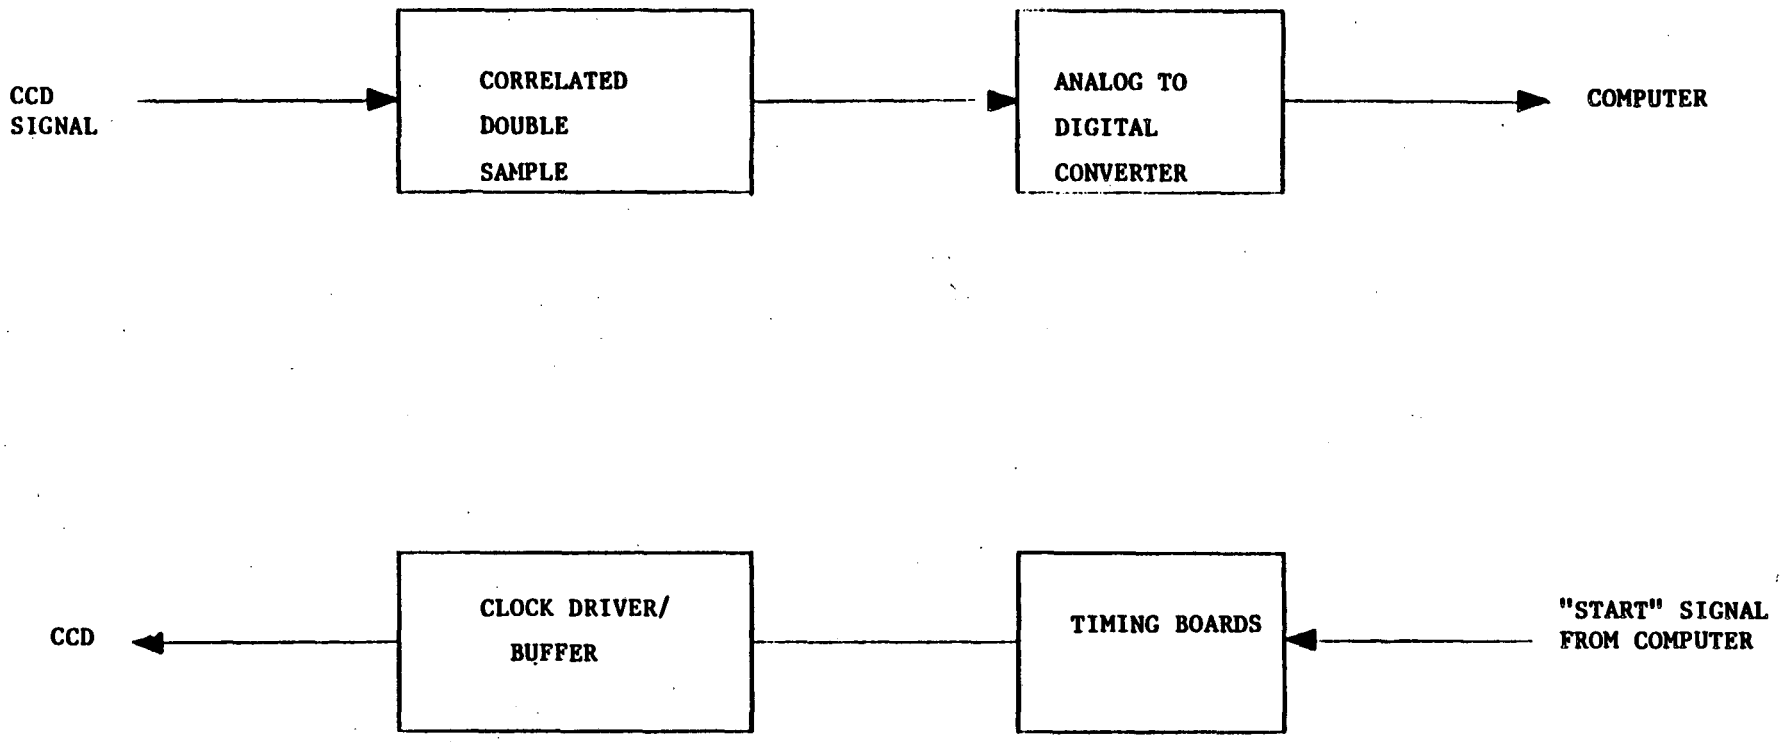
CCD
SIGNAL
CORRELATED
DOUBLE
SAMPLE
ANALOG TO
DIGITAL
CONVERTER
COMPUTER
CCD
CLOCK DRIVER/
BUFFER
TIMING BOARDS
"START" SIGNAL
FROM COMPUTER

Figure 10

11(b). The corresponding timing diagram is shown in Figure 12. The CCD readout sequence is begun by a START signal from the computer to Timing Board 1 (Figure 11(a)). There is a delay so that the output is referenced to the same place in the 60Hz line frequency so that noise from nearby power lines is always the same and can be compensated for. A set of counters is then cleared. These counters count the number of horizontal transfers by counting the number of reset pulses. After 320 horizontal transfers a vertical transfer is enabled. A second set of counters counts the number of vertical transfers. After 520 vertical transfers the image has been read out so the system resets and stops. This board also contains the master clock for the second board, the A/D converter, and data transfer.

The primary function of Timing Board 2 (Figure 11(b)) is to generate clock pulses for the vertical and horizontal phases. The mechanism for generating the pulses is to divide the master clock into 100 separate pulses. Two counters accomplish this. In this way, a pulse can be produced which is indexed relative to the overall timing scheme. The pulse width can be made any multiple of 0.5 microseconds, up to 50 microseconds total.

As an example of these features, consider the reset pulse and the horizontal phases. A horizontal transfer takes one timing period, so we would like the reset pulse to occur at the beginning of the interval. To this end, clock pulse 00 is used to set the S/R latch for the reset pulse high. The pulse should only be a few microseconds wide so clock pulse 04 resets the latch. This produces a 2 microsecond pulse at the beginning of every timing cycle.

From the timing diagram (Figure 12) we can see that the same proceedure is used for the horizontal phases, and indeed, for all the clock functions. The numbers on the edges of the pulses indicate where in the timing cycle the given latch is set or reset. Most of these clock functions are sent directly to the Correlated Double Sampling Board (CDS Board) and will be discussed in more detail below. The horizontal and vertical clocks, however, are sent to the Clock Driver Board.

The general schematic for the driver and output buffer is shown in Figure 13. The Clock Driver Board consists of nine sets of these driver/buffer combinations. Three sets are for the horizontal phases, and two groups of three sets are used for the two sets of vertical phases. One set is used to clock the image area, while the other is used to clock the storage area (see Figure 6). Since the storage area is identical to the image area, the Supernova Search uses both areas for imaging. This implies that the two sets of vertical

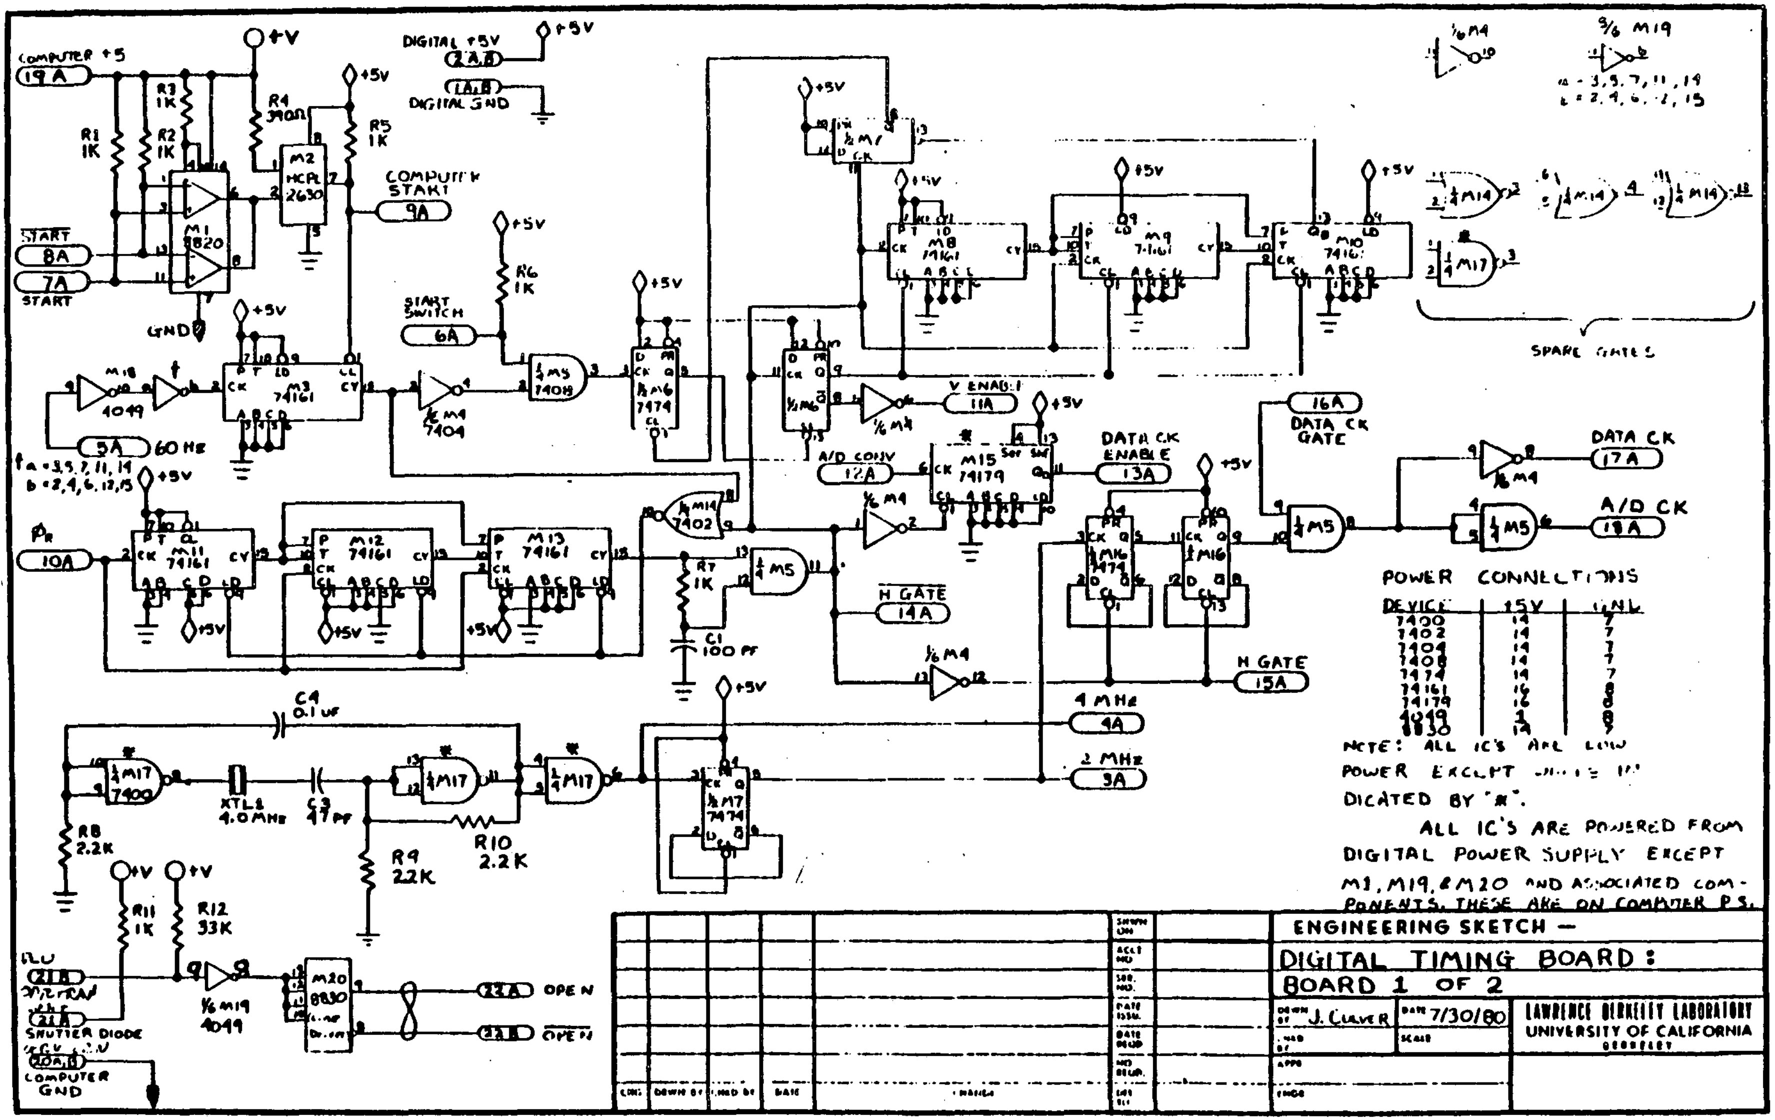
ENGINEERING SKETCH —
DIGITAL TIMING BOARD:
BOARD 1 OF 2
J. CULVER
7/30/80
LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
POWER CONNECTIONS
NOTE: ALL IC'S ARE POWERED FROM DIGITAL POWER SUPPLY EXCEPT M1, M19, & M20 AND ASSOCIATED COMPONENTS. THESE ARE ON COMPUTER P.S.
SPARE GATES
COMPUTER +5
COMPUTER START
START SWITCH
DIGITAL +5V
DIGITAL GND
V ENABLE
A/D CONV
DATA CK ENABLE
DATA CK GATE
DATA CK
A/D CK
H GATE
4 MHz
2 MHz
XTL1 4.0 MHz
COMPUTER GND
SHUTTER DIODE

Figure 11(a)

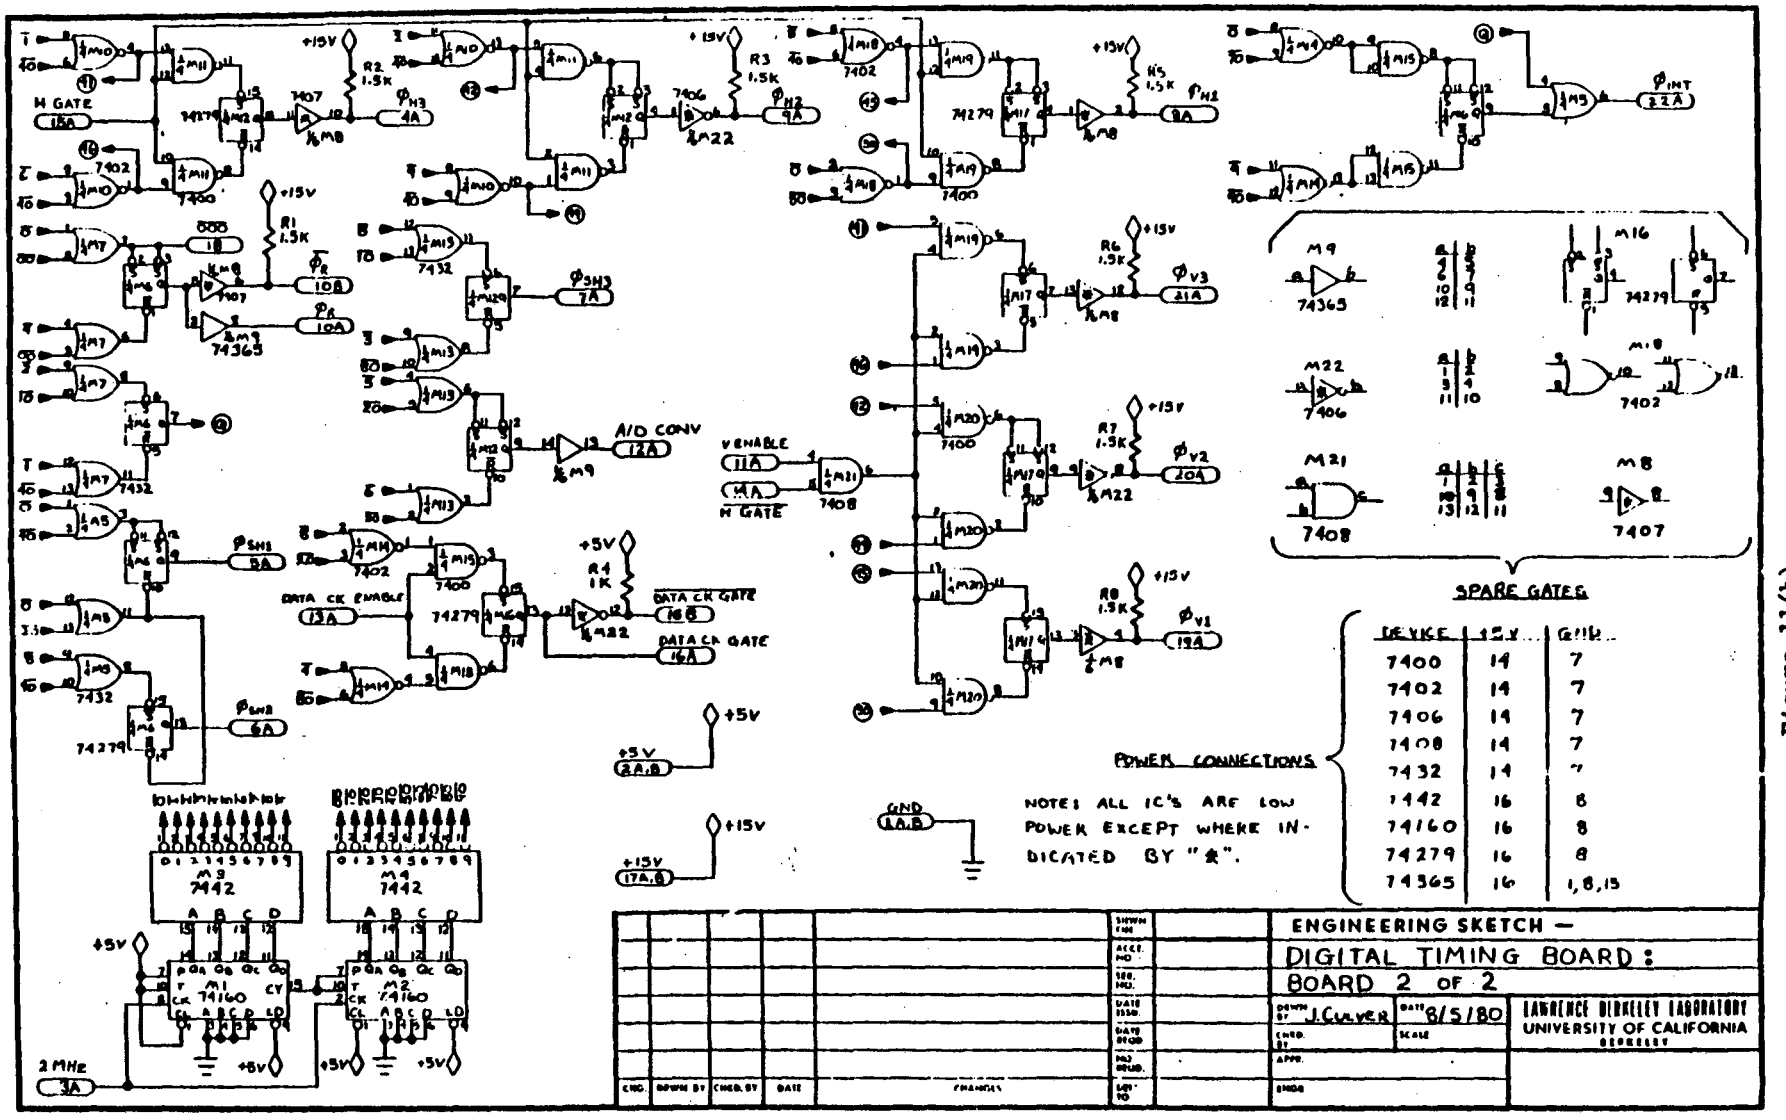
H GATE
15A
A/D CONV
12A
VENABLE
11A
H GATE
14A
DATA CK ENABLE
13A
DATA CK GATE
16B
DATA CK GATE
16A
2 MHz
3A
+15V
+5V
GND
1A,B
+5V
2A,B
+15V
17A,B
POWER CONNECTIONS
NOTE: ALL IC'S ARE LOW POWER EXCEPT WHERE INDICATED BY "*".
SPARE GATES
DEVICE | +5V | GND
7400 | 14 | 7
7402 | 14 | 7
7406 | 14 | 7
7408 | 14 | 7
7432 | 14 | 7
7442 | 16 | 8
74160 | 16 | 8
74279 | 16 | 8
74365 | 16 | 1,8,15
ENGINEERING SKETCH —
DIGITAL TIMING BOARD:
BOARD 2 OF 2
DRWN BY J.CULVER
DATE 8/5/80
SCALE
CHKD. BY
APPR.
LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
BERKELEY

Figure 11(b)

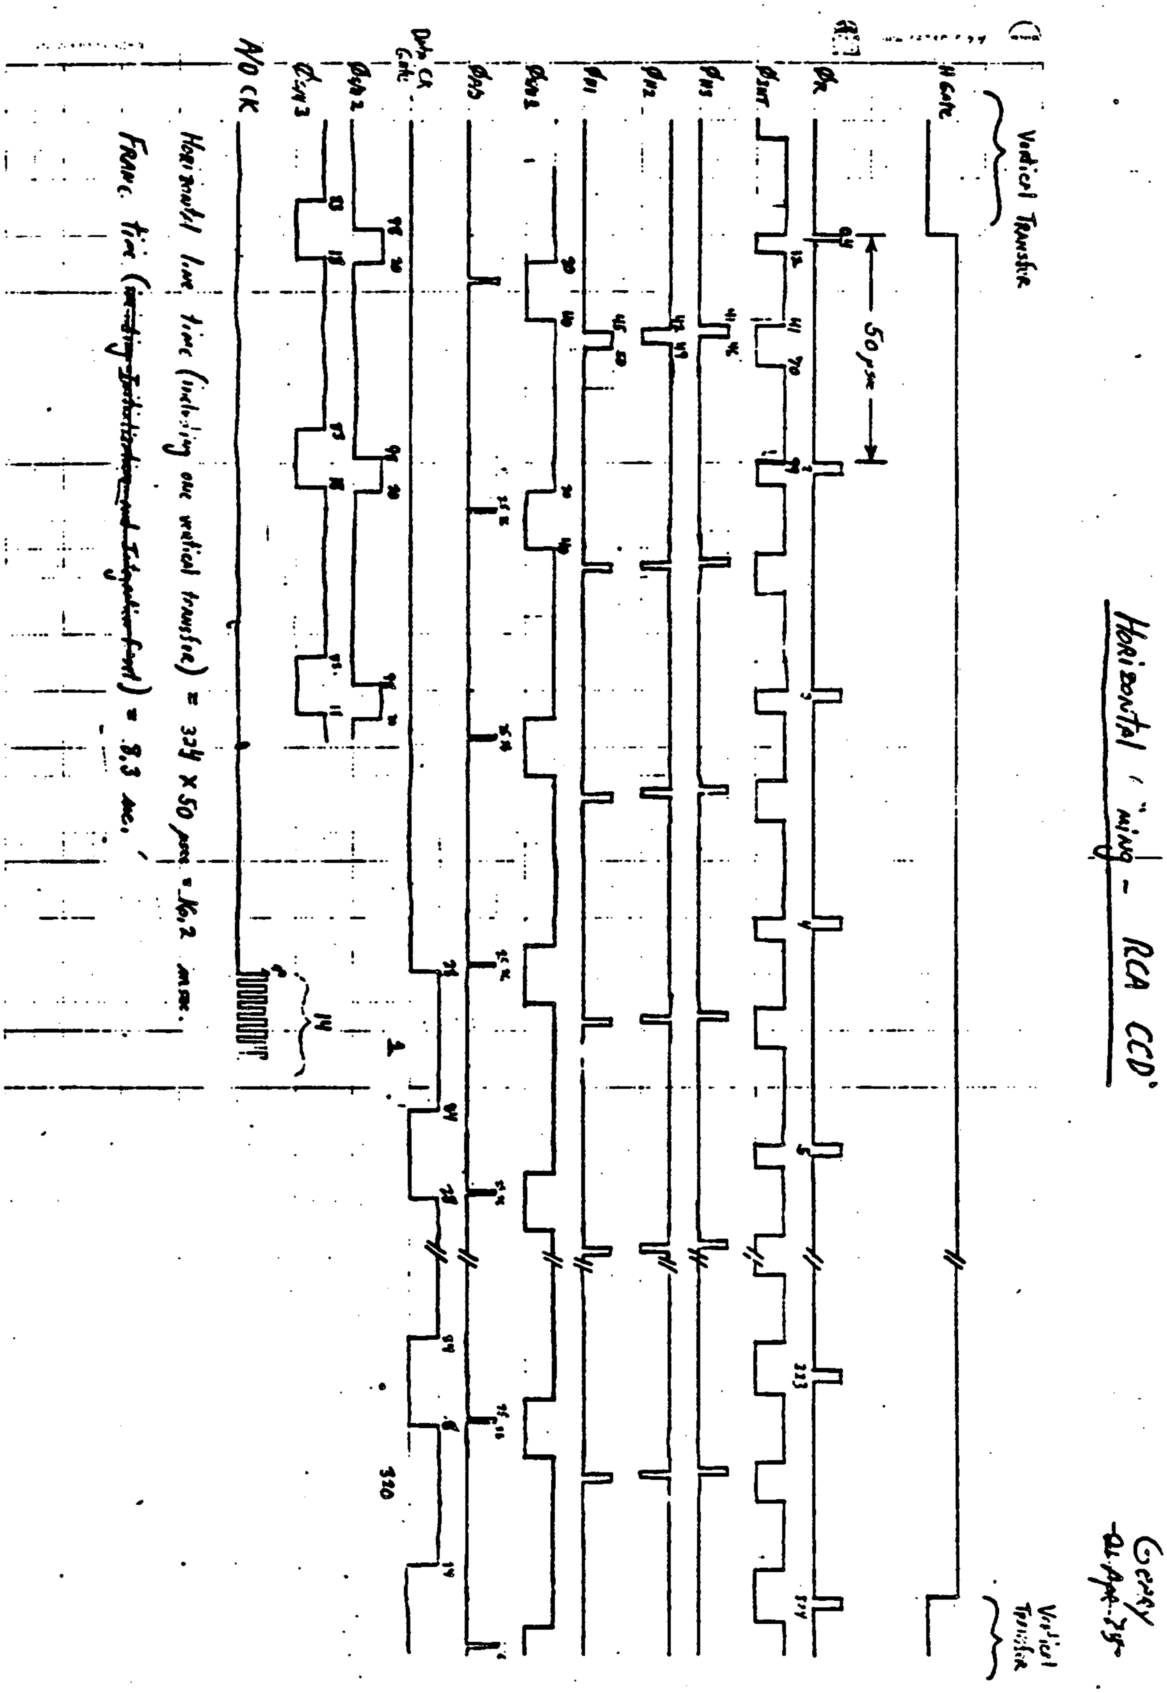

Figure 12

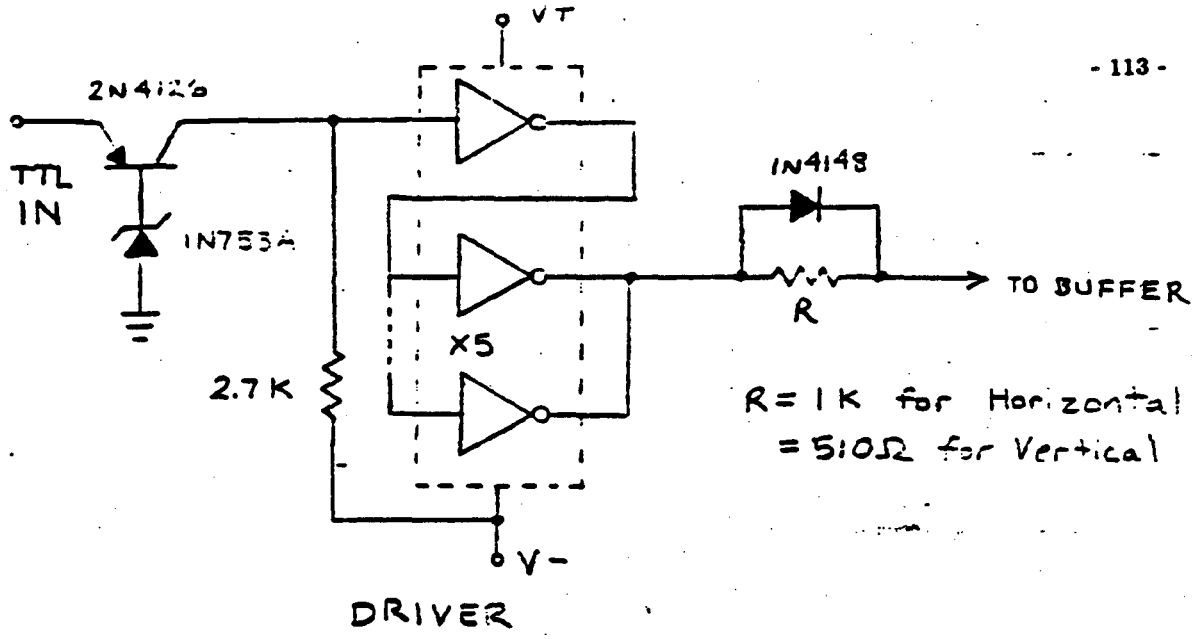
VT
2N4126
TTL
IN
1N753A
1N4148
TO BUFFER
R
X5
2.7K
R = 1K for Horizontal
= 510Ω for Vertical
V-
DRIVER

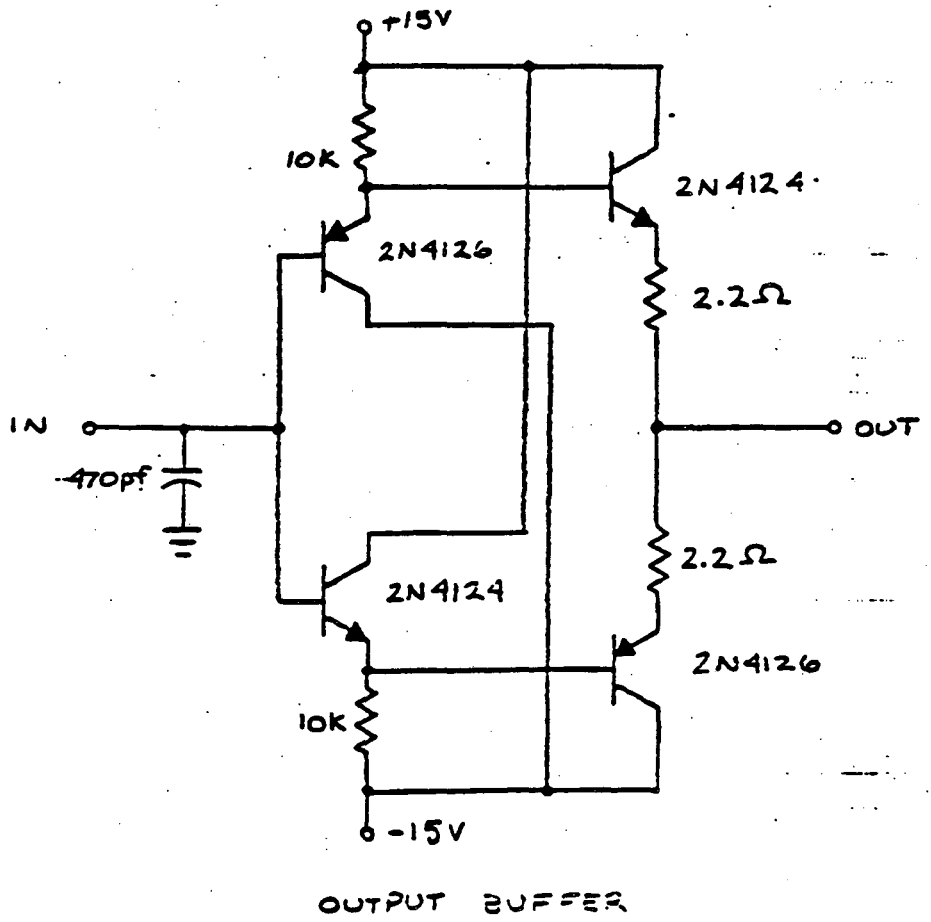
+15V
10K
2N4124
2N4126
2.2Ω
IN
OUT
470pf
2.2Ω
2N4124
2N4126
10K
-15V
OUTPUT BUFFER

Figure 13

clocks are the same, they simply require seperate driver/buffer circuits.

The driver circuit works in the following way. The transistor and zener diode serve as a level-shifter for the TTL input. The CMOS inverter (4049) is used to obtain fast switching and to clamp the output to the voltages specified by V+ and V-. This means that a logic low will produce an output at V-, and a logic high will output a V+. This characteristic of CMOS devices is used to tailor the amplitude of the waveform for each set of phases by varying the voltages from the Clock Driver Power Supply Board (these are simple, low current, op-amp regulated power supplies, and are not relevent to the discussion at hand).

The input to the buffer circuit is shaped by the resistor-capacitor-diode arrangement. The diode provides a fast rise-time, while the RC circuit produces an exponential fall-off. This gradual fall-off transfers the charge packets more efficiently The output buffer is a "push-pull temperature-compensated cascaded-emitter follower stage". It has unity gain, and is designed to drive the capacitive load of the CCD.

2. Signal Processing*

Correlated double sampling (White et al. (1974)) is an output sensing technique whereby the effects of reset noise can be eliminated. Referring back to Figure 5 we see that the reset proceedure is to set the output gate to a reference potential once each clock period. This produces reset noise which is still present on the output gate when the charge packet is transferred to it. The noise before the charge transfer and the noise present in the signal after charge transfer are correlated so that by sampling the noise, it can be subtracted from the signal + noise which is sampled after the charge packet is transferred. What is left is pure signal. An equivalent circuit for this reset proceedure is a capacitor being charged by a resistor as depicted in Figure 14. R is the resistance of the reset transistor and C_{FD} is the capacitance of the sense node. It can be shown (Barbe (1975)) that the mean-square deviation of the voltage across C_{FD} is given by

$$\overline{v^2(t)} = \frac{k\,T}{C_{FD}}\left[1 - \exp\left(\frac{-2t}{RC_{FD}}\right)\right].$$

When the transistor is on, $R_{on}C_{FD} \approx 10^{-9}$ seconds

* Following Beynon and Lamb, section 2.10.3.

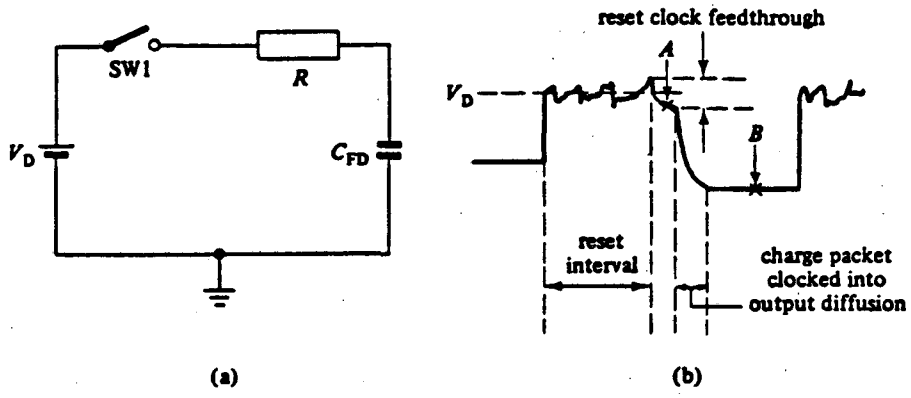

Elimination of resetting noise by correlated double sampling: (a) equivalent circuit; (b) output voltage waveform of reset amplifier circuit, showing sampling points *A* and *B* (after Barbe and Baker (1975)).

Figure 14

When the transistor is off, $R_{off}C_{FD} \approx 10^{-1}$ seconds

So, in terms of noise, the circuit will have reached equilibrium within the 2 microsecond reset time. When the reset transistor is turned off, there will be a drop in the output voltage due to clock feed-through, then a longer delay on the order of a tenth of a second. This gives us enough time to sample the noise (point A), transfer the charge, and sample the noise + signal (point B) before there is significant variation in the noise level. The noise is the same at B as at A (i.e. it is correlated) as long as the time between A and B is much less than $R_{off}C_{FD}$. Subtracting V_A from V_B produces a signal free from reset noise. This is the Correlated Double Sampling technique.

In the CDS Board (Figure 15) the signal from the CCD is amplified, integrated, and double-sampled. To understand the operation in detail, refer back to the timing diagram (Figure 12). The gate is reset at the beginning of the timing cycle. The transistor is turned off, and 4 microseconds later the noise voltage is integrated for 14 microseconds. This integrated voltage value is then latched into a sample-and-hold (S/H1) for later subtraction. The integrator is reset during charge transfer which takes 5 microseconds. The signal + noise is integrated for another 14 microseconds, at which time the voltage is latched into S/H2. The time required for this entire process is less than the characteristic noise correlation time so that the noise is correlated. The two signals are then subtracted and sent to the A/D converter. The third sample-and-hold retains the signal voltage while the conversion takes place, allowing the integrator to proceed with the next charge packet. Finally, the output of the A/D converter is sent to the computer along with a reference clock. DATA CLOCK GATE indicates when data is available ,and V ENABLE indicates when vertical transfer is taking place. This is used to index the data unambiguously in the computer memory.

The CCD driving electronics serve as the link between the CCD and the computer. The circuits described above provide the necessary go-between to create an integrated system for astrophysical imaging

CONCLUSIONS

CCD technology is quite recent, but it has already established itself as a useful tool for solid-state imaging. Current technology is providing the means to use the CCD in the more specialized imaging applications like the low noise, high efficiency, low

in the more specialized imaging applications like the low noise, high efficiency, low light-levels one finds in astrophysics. This, coupled with the availability of inexpensive and powerful small computers may herald a new era in automated astronomical observing.

REFERENCES

J. D. E. Beynon and D. R. Lamb, ed., Charge-Coupled Devices and their Applications, McGraw-Hill Book Company (UK) Limited, 1980

M. J. Howes and D. V. Morgan, ed., Charge-Coupled Devices and Systems, John Wiley & Sons, 1979

R. A. Muller et al., A Search for Supernova Explosions, research proposal to NASA, 1981

C. Pennypacker, Electronics Memos, 1980

KEY TO FIGURES

Figures 1-4, 6, 7, 14: Beynon and Lamb (1980)

Figures 8, 9: Muller (1981)

Figure 12: J. Geary, Harvard (1980)

Figures 10, 11, 13, 15: J Culver, based on circuits from J. Geary

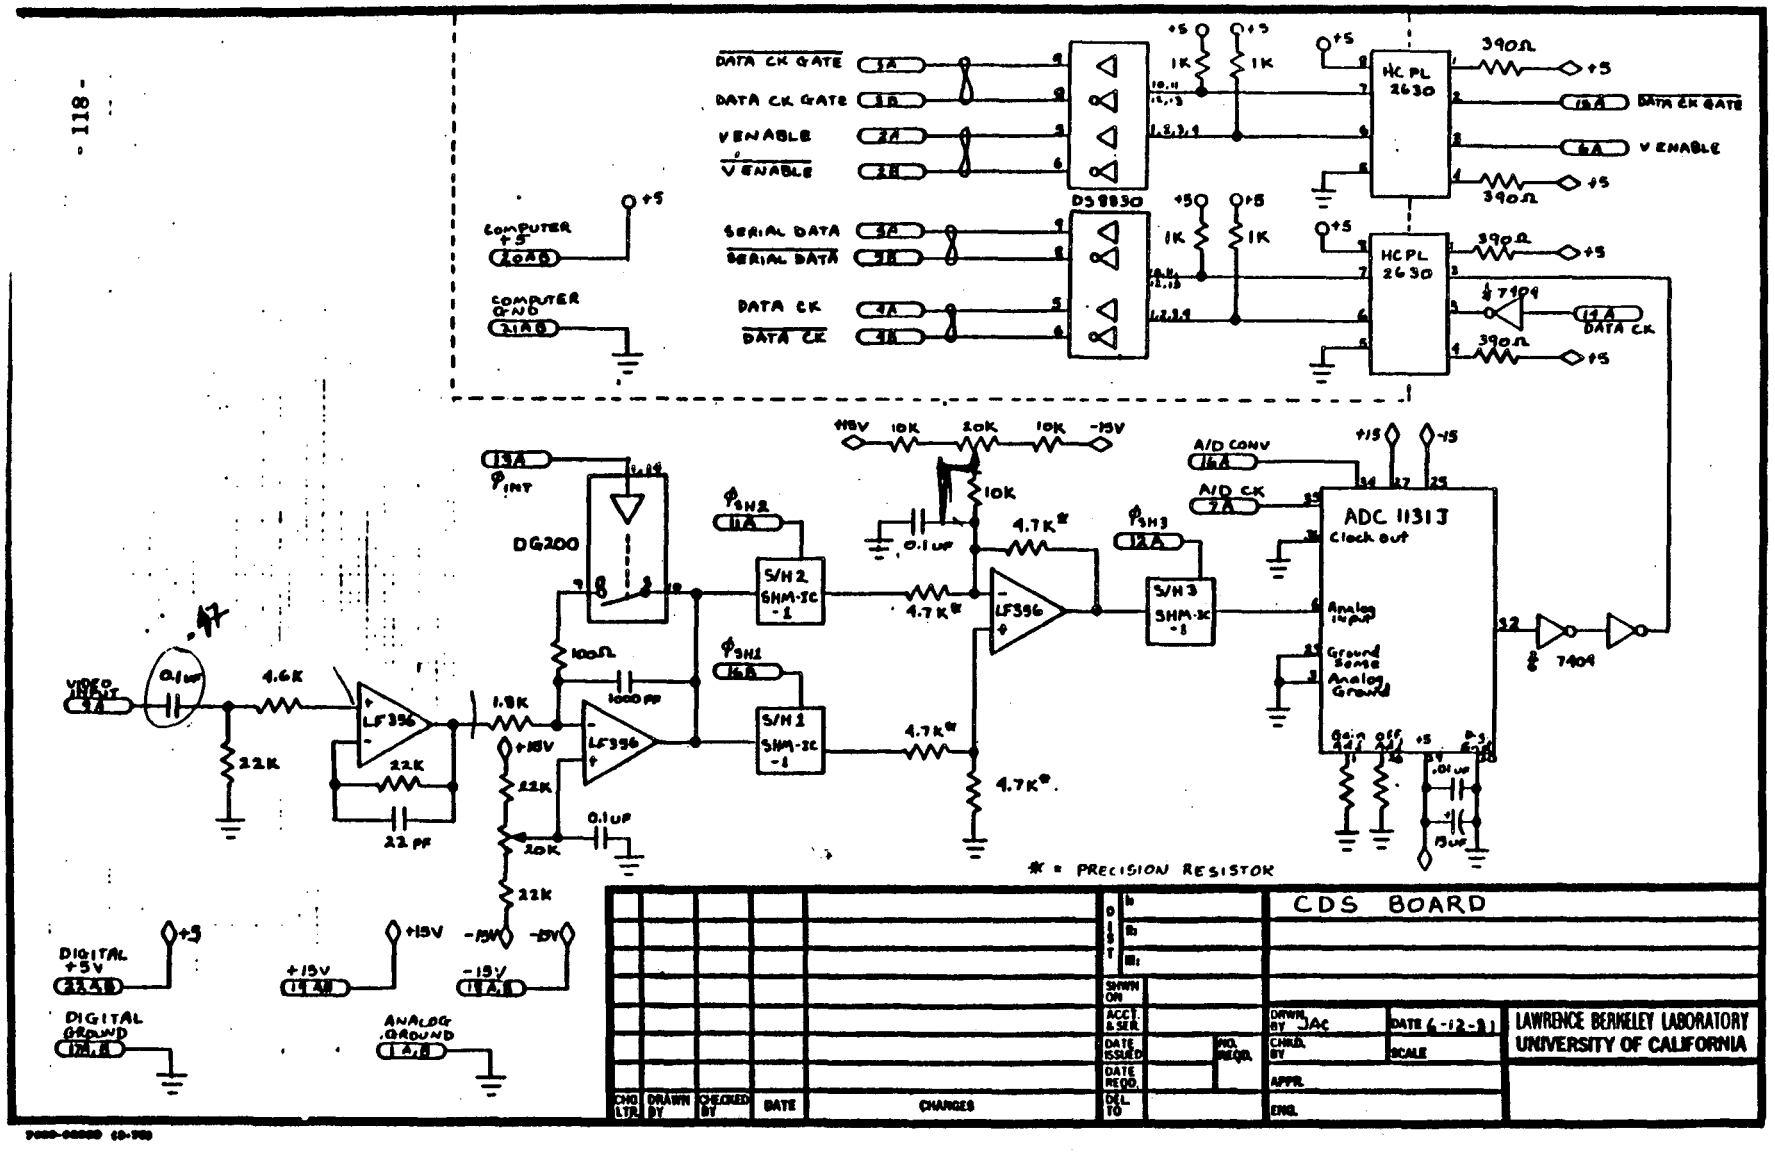
DATA CK GATE
DATA CK GATE
V ENABLE
V ENABLE
SERIAL DATA
SERIAL DATA
DATA CK
DATA CK
DS8830
HC PL 2630
HC PL 2630
390Ω
DATA CK GATE
V ENABLE
DATA CK
7404
COMPUTER +5
COMPUTER GND
DG200
S/H2 SHM-2C -1
S/H1 SHM-2C -1
S/H3 SHM-2C -1
LF356
LF356
LF356
ADC 11313
A/D CONV
A/D CK
Clock out
Analog Input
Ground Sense
Analog Ground
Gain adj
Off adj
VIDEO INPUT
DIGITAL +5V
DIGITAL GROUND
ANALOG GROUND
+15V
-15V
* = PRECISION RESISTOR
CDS BOARD
LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
DRAWN BY JAC
DATE 6-12-81
SCALE
CHG LTR
DRAWN BY
CHECKED BY
DATE
CHARGES

Figure 15

APPENDIX II

VIDEOTAPE DATA RECORDER SYSTEM MANUAL

This Page Included
For
Page Number Sequence
Only

APPENDIX II

Videotape Data Recorder System Manual

1. Specifications

The data recorder system for the Berkeley automated supernova search consists of three parts: a transmitter, a receiver, and one or two videotape recorders. The transmitter accepts digital data in 14 bit serial form (or 16 bit serial or parallel form, with minor modifications), adds error checking and recovery information, and produces a standard black and white video signal. This signal can be recorded on any video recorder or transmitted via 75-ohm cable (or broadcast, should anyone so desire).

Each CCD pixel (14 bits) is converted to a video word (32 bits). Each word is transmitted twice (pass 0 and pass 1) giving redundancy to reduce the error rate. Two such double words occupy one video line. 224 lines of data (up to 448 pixels) are transmitted in each video field (1/60 second). The transmitter operates asynchronously: normally CCD pixels arrive more slowly than video words are transmitted. The unused video words are marked as invalid, and are discarded at the receiver. If CCD pixels arrive more rapidly than they can be transmitted (during vertical blanking intervals) they are stored in a first-in first-out buffer.

The receiver accepts a recorded video signal from the transmitter, recovers the stored information, detects and corrects errors, and sends 14 bit serial data to the supernova search computer. A jog corrector circuit sends dummy pixels to the computer when video words are lost due to dropouts; this maintains the alignment of columns in the CCD image.

We record data on a Sony SL-2000 Betamax videotape recorder, slightly modified for improved performance, operating in the Beta-II mode. A Zenith VR9800 recorder, internally identical to the Sony recorder, is used for playback.

A fiberoptic input adapter allows the transmitter to accept data over HP plastic-core fibr optic links. A modified infrared remote control unit allows computer control of the videotape recorders via an RS-232 serial line.

Figure 1 shows the current configuration of the videotape storage system.

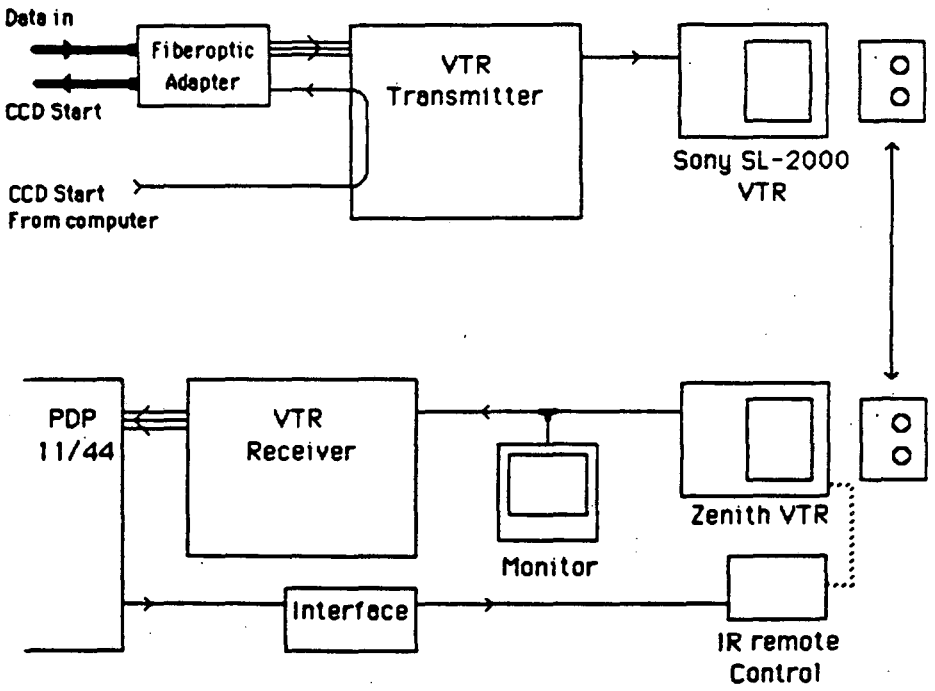

Appendix 2 -- Figure 1
Videotape Data Storage System

2. Circuit description

2.1. Transmitter

Figure 2 is a block diagram of the videotape system transmitter (encoder). The full schematic for the transmitter is given in figure 3.

The input circuit (Integrated circuits U301 - 303, U312, and U313 in figure 3) converts serial data from the CCD to parallel 14 bit words (one word per CCD pixel), and generates a load pulse when each word is complete. The system will also work with 16 bit words; currently the two highest order bits are used to implement the jog corrector (see below). DS8820 line receivers convert balanced input signals to TTL levels. Serial data bits are clocked into the shift register on the falling edge of the data clock. A load pulse derived from the data clock gate (DC gate) signal transfers the complete word into the first-in first-out (FIFO) buffer (U3-U10).

The jog corrector circuit in the transmitter consists of a 2 bit counter driving data bits 14 and 15. The counter is reset to 0 between CCD frames by the timing out of a retriggerable one-shot, so the first pixel transmitted always has the same counter value.

The FIFO buffer stores the output of the CCD for the duration of a vertical blanking interval (approximately 1.5 milliseconds) when no data can be written to the videotape. The maximum CCD readout rate that the videotape store can accept is 448 pixels per video field (1/60 second), or 26,880 pixels per second. At this rate, the buffer must be able to hold 48 pixels worth of data. The recorder was built with a 128-word buffer, but the first 64 words (IC's U3-U6) have been removed, leaving a 64 word buffer. The extra buffers can be replaced if the recorder is ever modified for a higher data rate (4 pixels per line).

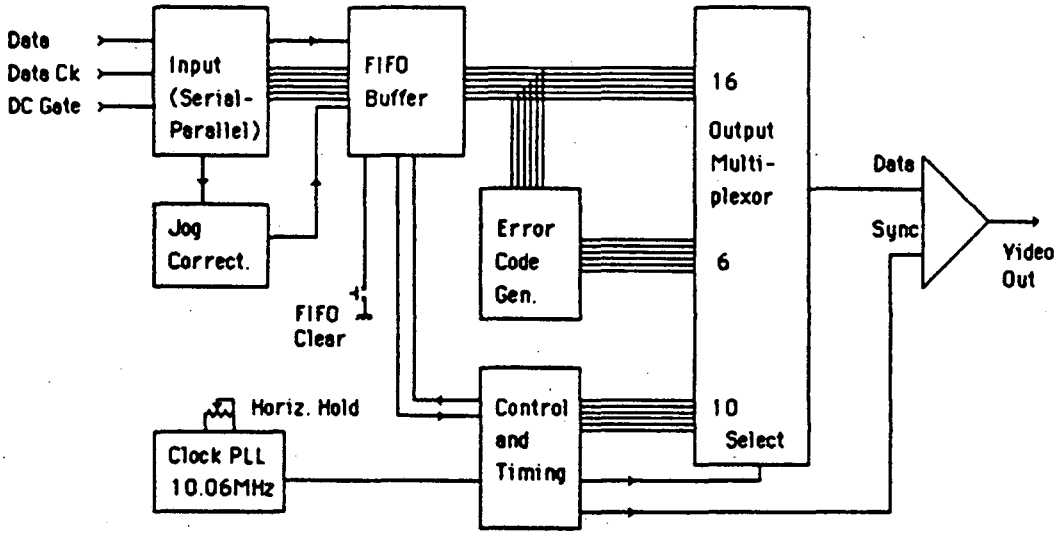

Appendix 2 -- Figure 2
Videotape Transmitter Block Diagram

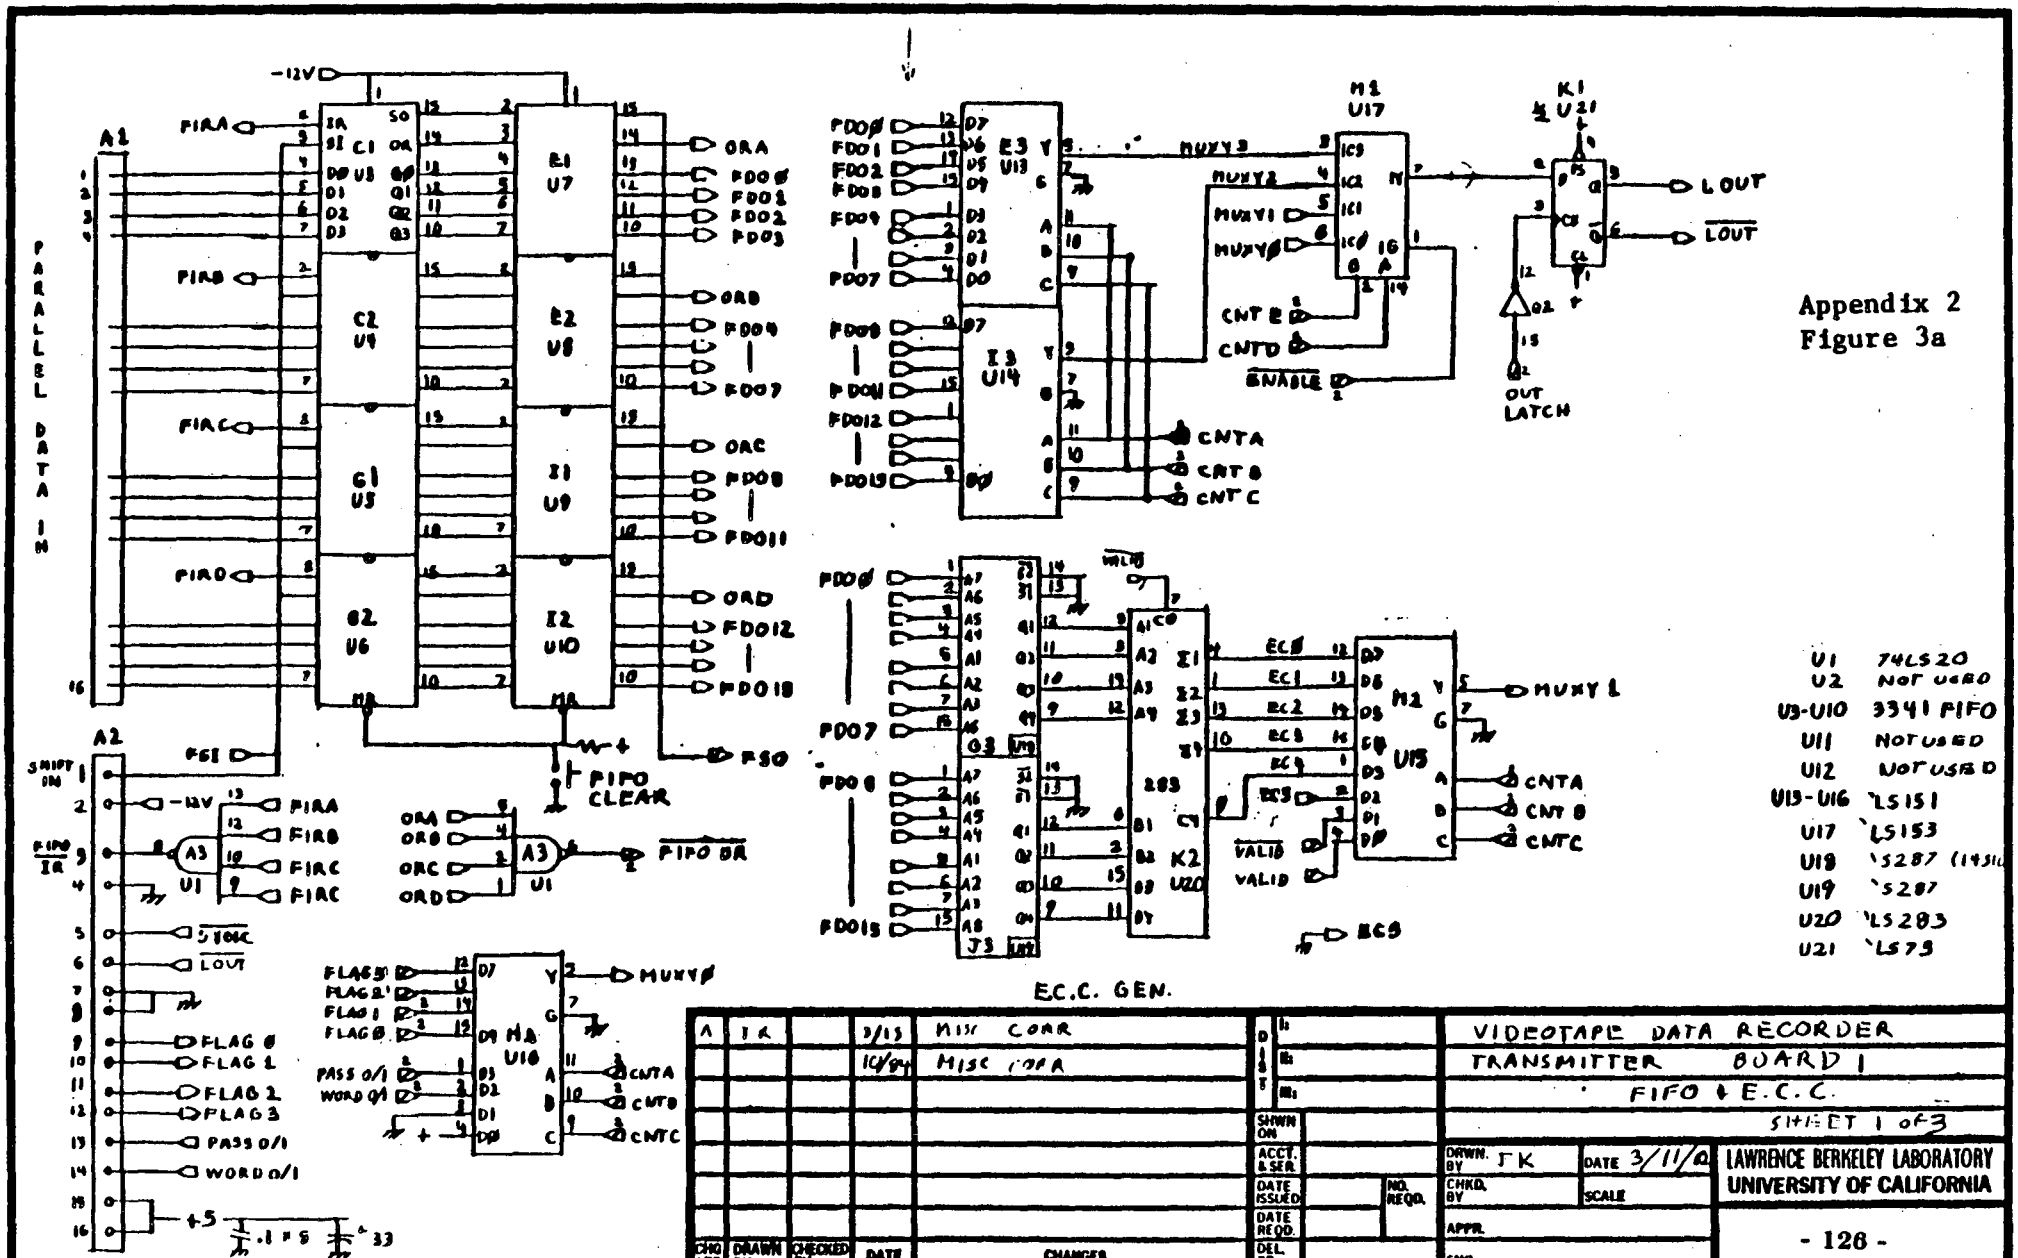

E.C.C. GEN.

Appendix 2
Figure 3a

U1	74LS20
U2	NOT USED
U3-U10	3341 FIFO
U11	NOT USED
U12	NOT USED
U13-U16	LS151
U17	LS153
U18	LS287 (1451)
U19	LS287
U20	LS283
U21	LS75

LTR	DRAWN BY	CHECKED BY	DATE	CHANGES
A	TK		2/15	MISC CORR
			10/84	MISC CORR

VIDEOTAPE DATA RECORDER
TRANSMITTER BOARD 1
FIFO & E.C.C.
SHEET 1 OF 3

DRWN. BY TK DATE 3/11/85
CHKD. BY SCALE
APPR.
ENG.

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA

7000-88860 (8-78)

Appendix 2
Figure 3b

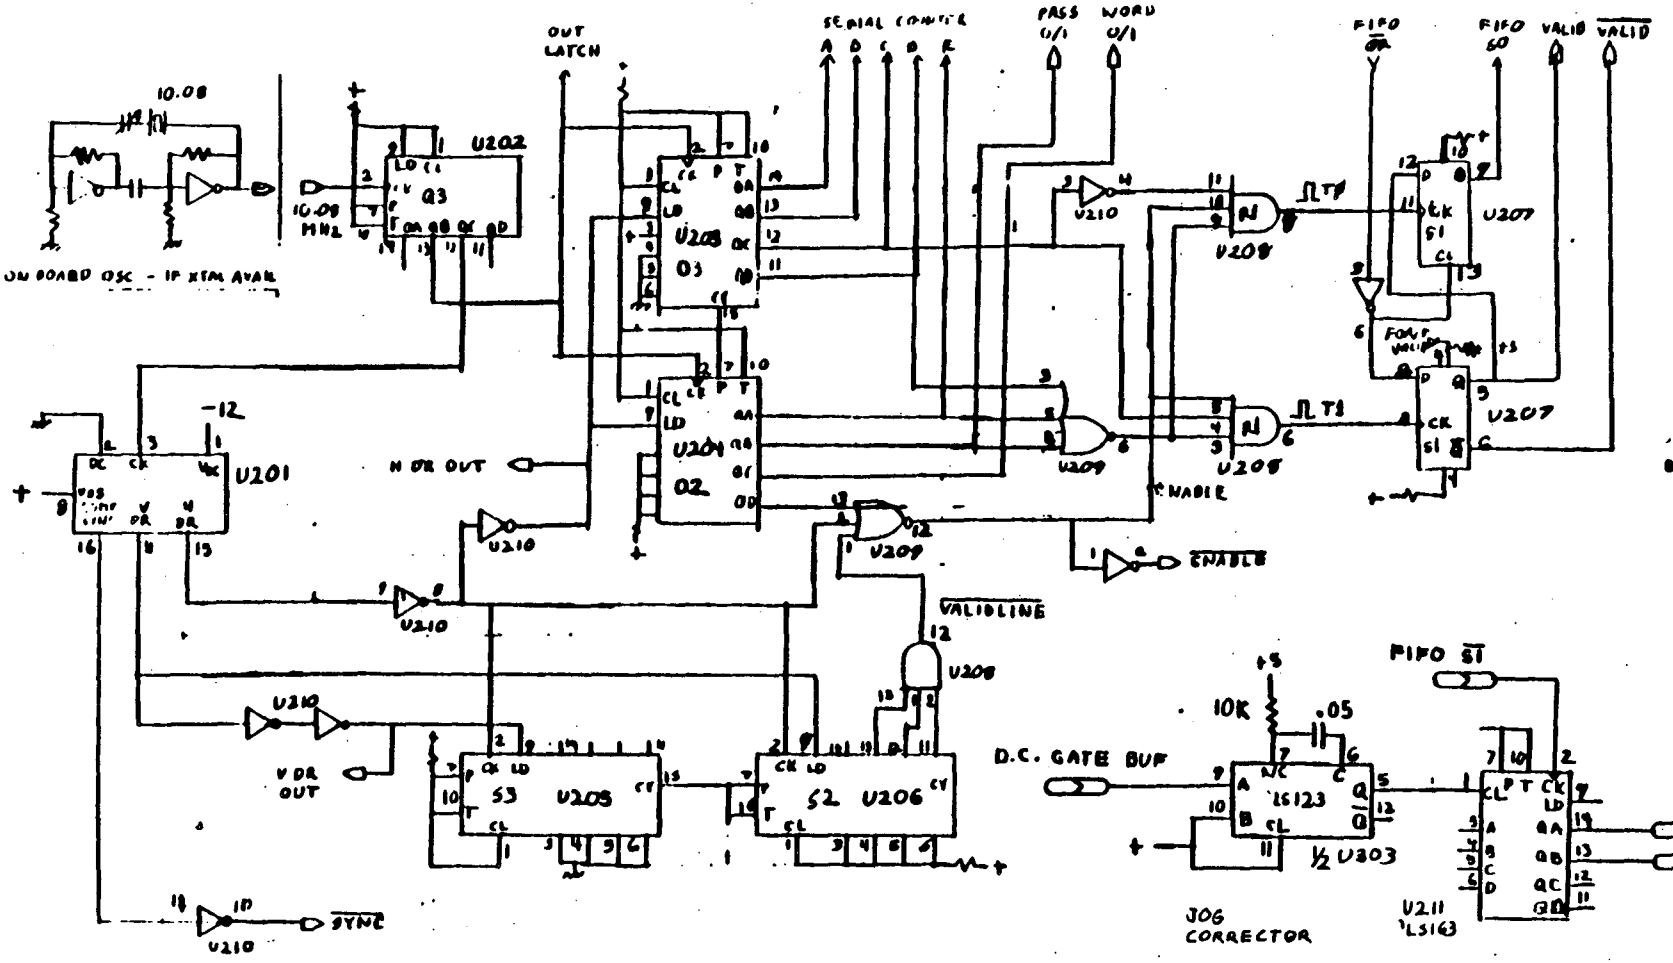

U201	MP15321
U202	74LS163
U203	"
U204	"
U205	"
U206	"
U207	74LS73
U208	74LS10
U209	74LS27
U210	74LS04
U211	74LS163

CHG LTR	DRAWN BY	CHECKED BY	DATE	CHANGES
			10/64	MISC CORR

DIST: H1, H2, H3

SHWN ON | ACCT. & SER. | DATE ISSUED | NO. REQD. | DATE REQD. | DEL. TO

VIDEOTAPE DATA RECORDER
TRANSMITTER BOARD1
CONTROL CIRCUITS
SHEET 2 of 3

DRWN. BY JK | DATE 3/13/82
CHKD. BY | SCALE
APPR.
ENG.

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA

7600-88800 (8-78)

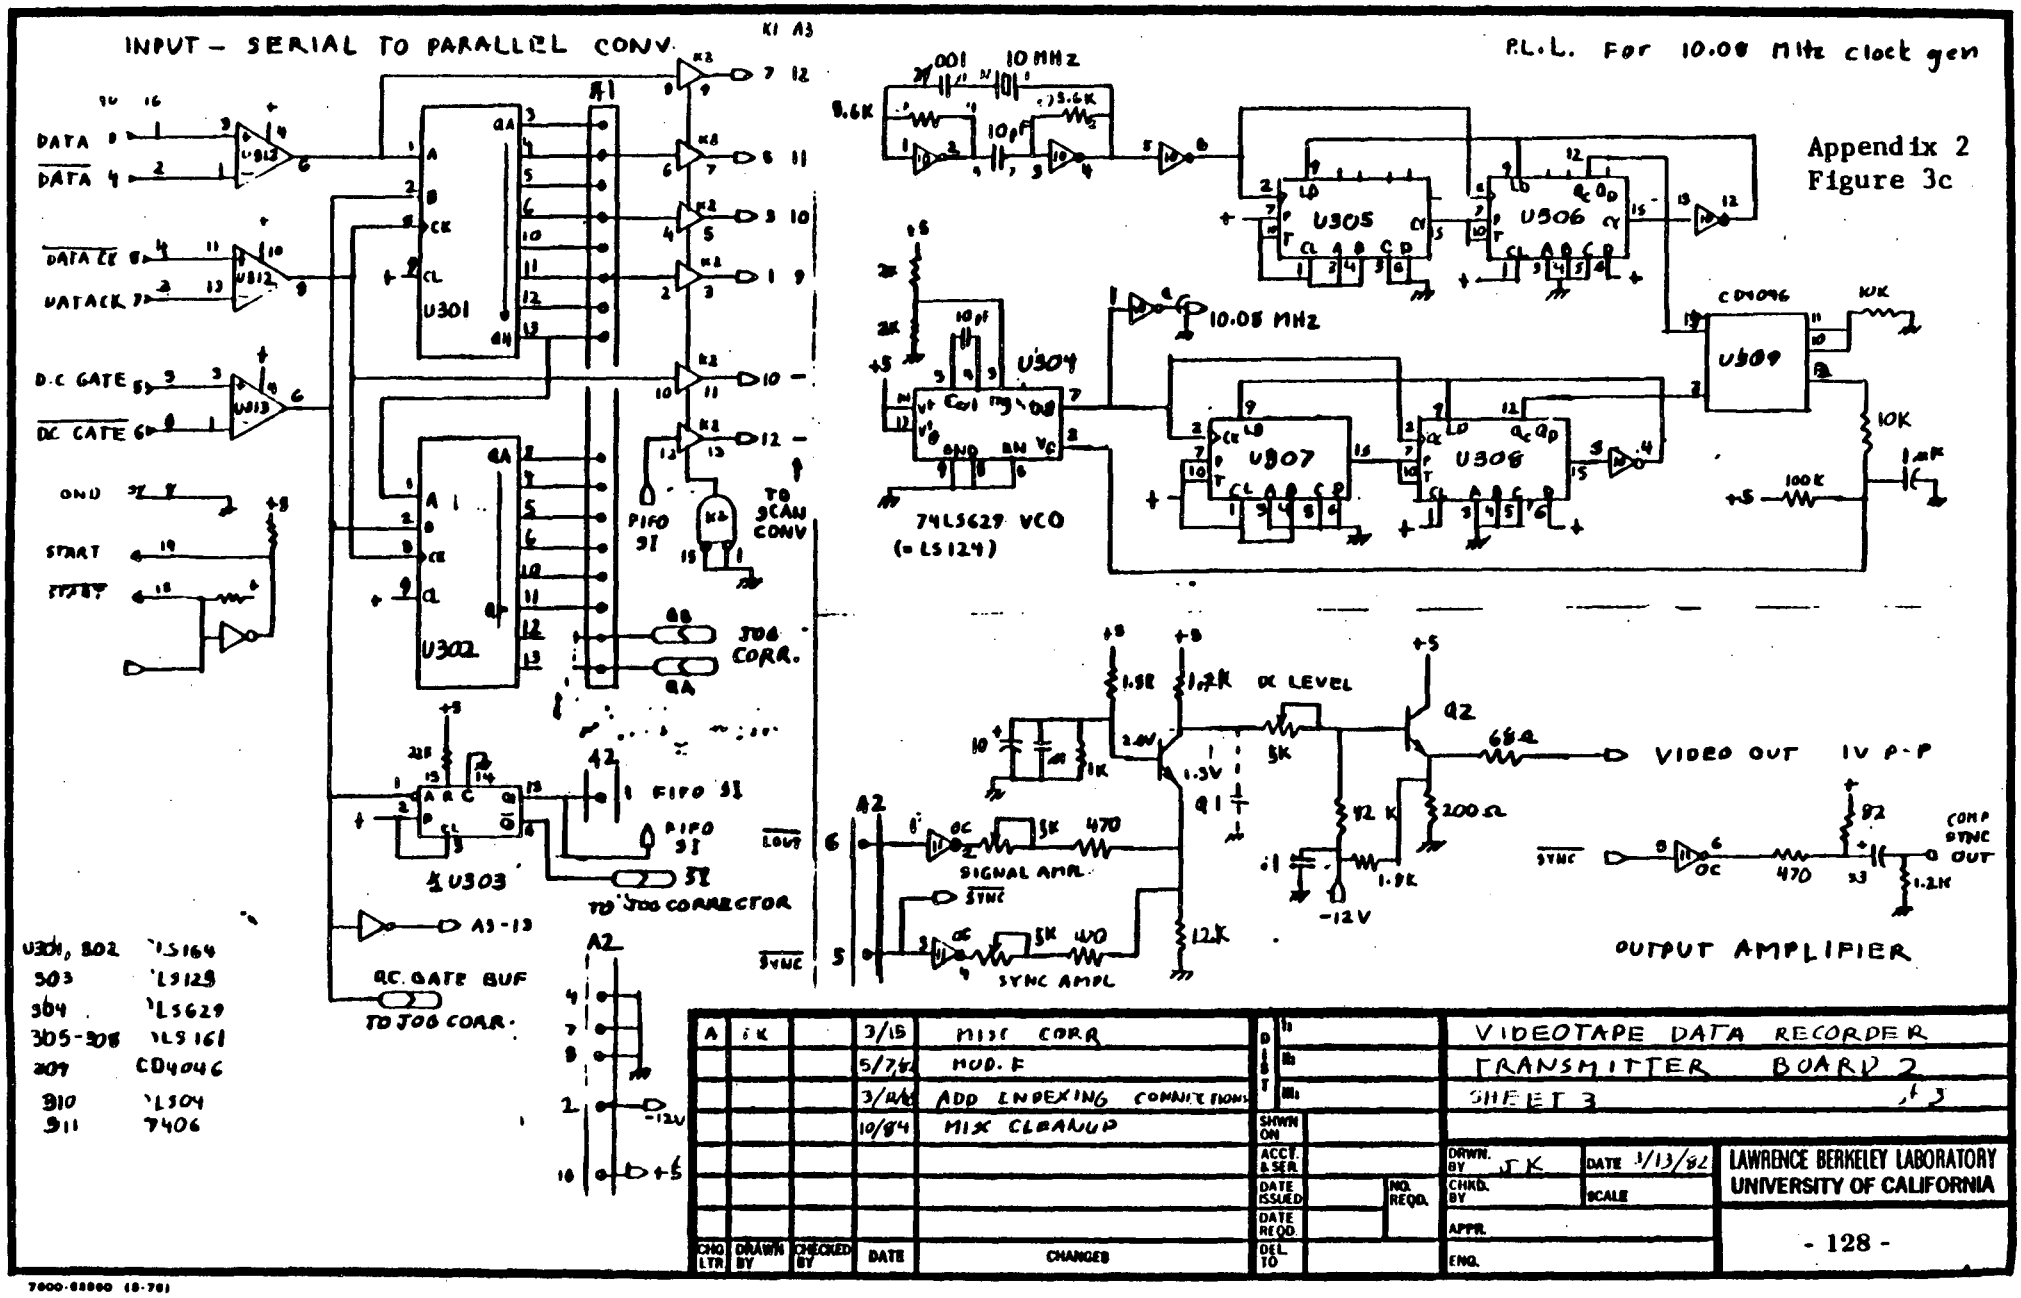
INPUT – SERIAL TO PARALLEL CONV.
P.L.L. For 10.08 MHz clock gen
Appendix 2
Figure 3c
DATA
DATA
DATA CLK
DATACK
D.C GATE
D.C GATE
GND
START
START
U301
U302
U303
U304
U305
U306
U307
U308
U309
74LS629 VCO
(= LS124)
10 MHz
10.08 MHz
CD4046
TO SCAN CONV
PIFO S1
PIFO S1
TO J06 CONNECTOR
J06 CORR.
DC GATE BUF
TO J06 CORR.
DC LEVEL
VIDEO OUT 1V P-P
SIGNAL AMP.
SYNC AMPL
OUTPUT AMPLIFIER
-12V
U304, 302 LS164
303 LS129
304 LS629
305-308 LS 161
309 CD4046
310 LS04
311 7406
VIDEOTAPE DATA RECORDER
TRANSMITTER BOARD 2
SHEET 3 of 3
3/13/84
MISC CORR
5/7/84
MOD. F
3/84
ADD INDEXING CONNECTIONS
10/84
MISC CLEANUP
DRAWN BY JK
DATE 3/13/84
CHKD. BY
SCALE
APPR.
ENG.
LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA

An interesting problem with the FIFO buffer has been observed a few times. The buffer consists of 4 (originally 8) chips, each storing 64 4 bit words. Normally, all 4 chips operate together: words are entered into all chips at once, propagate through the chips in parallel, and are read out of all chips at once. Occasionally, presumably due to a noise spike, one chip acquires one or more "extra" words -- its output buffer is loaded and its output ready line is turned on. When the next real data word (word N) arrives, 4 bits "stack up" behind the false word. As stored on tape, word N contains 4 bits of garbage, word N + 1 contains 4 bits from word N, and so on. This offset is never corrected, because an extra fractional word is always left behind when the last full word has been transmitted.

The problem is rare enough that a manual FIFO CLEAR switch, pushed a few times per night, is adequate insurance against serious loss of data, although an automatic clearing circuit should be added. The problem was, however, spectacularly hard to diagnose on its first appearance, since it produces a "contour map" pattern of pixels randomly good or bad depending on the binary values of neighboring pixels.

The output of the FIFO goes to a multiplexer (U13, U14) for conversion back to serial form, and to the error check code generator. Five error check bits (EC0-EC4 in figure 3) contain the total number of 0's in the data word. This value is generated by two 256 x 4 PROM's, each covering 8 data bits; the outputs of these PROM's are added in a 4 bit adder (U20). 0's counting is used because the expected source of errors is dropouts, which look to the receiver like a string of 1's or of 0's but are unlikely to mix the two values. Simple parity checking cannot reliably detect multiple bit errors. Counting 0's instead of 1's insures that the check code for an all 0 data word is not all 0 also.

I considered an LSI error-correction chip (TI 74LS630), but decided that it was

difficult to obtain and (since single bit error correction was not a requirement) offered no advantages over 0's counting. Many more sophisticated error correction schemes exist, but most require substantial additional hardware, including relatively large buffer memories, and are not well suited to such a simple system. However, low cost Winchester disks have stimulated the development of cheap LSI devices for burst error correction, and a future version of this system might well employ them. A sixth error check bit (EC5) is included in case a different error correction scheme is someday needed.

In addition to the 16 data bits and 6 error correction bits, the output word includes 2 synchronizing bits, 2 valid data bits to indicate the presence of CCD data in a given video word, 2 bits to indicate the position on the video line (PASS and WORD) and 4 flag bits, which are currently not used. The total word length is 32 bits. Each output word is transmitted twice, giving a total of 64 bits transmitted per CCD pixel. Two CCD pixels are transmitted per horizontal video line, for a total of 128 bits per line. The output format is shown in figure 4.

The control logic is driven by a master clock at 10.08 MHz. The master clock is divided by 8 to drive a National MM5321 TV sync generator, which produces standard sync signals for compatibility with all video recorders. Unfortunately, 10.08 MHz is not a standard crystal frequency. The MM5321 could probably be run at 1.25 MHz (a 0.8% frequency error) without affecting the performance of most recorders, but in the interest of guaranteed compatibility, the clock generator circuit includes a phase locked loop to multiply the frequency of a 10 MHz crystal controlled oscillator by 126/125. If this loop unlocks, the output video signal acquires a characteristic "wiggle" pattern when viewed on a monitor; this can be fixed by adjusting the "horizontal hold" potentiometer to vary the free running frequency of the loop oscillator until the loop locks.

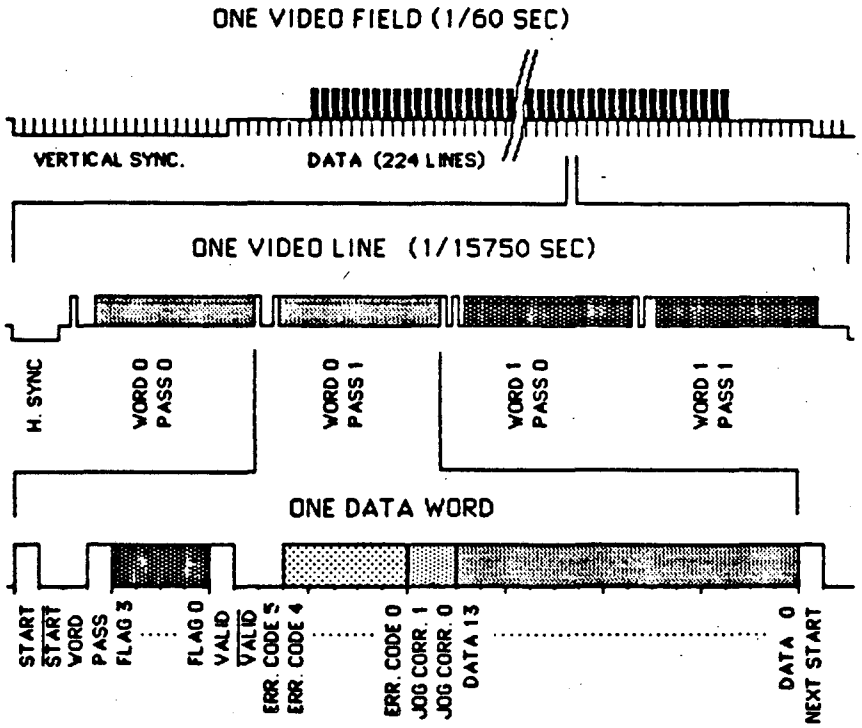

Appendix 2 -- Figure 4
Video Data Format

The horizontal and vertical drive outputs of the MM5241 clear the bit and line counters (U203-204 and U205-206) respectively. The line counter generates a "valid line" signal for lines 16 - 239 (counting from the end of the V drive pulse); only these lines are used to store data. The bit counter, clocked at 2.52 MHz, drives the output multiplexers to select the desired output bit at each point on the horizontal line. The valid data bits and the FIFO output are set separately; a state diagram for the control logic is given in figure 5. The video output bits are latched by U21 and summed with the composite sync from the MM5241 in the output circuit. The gain and DC LEVEL adjustments are set to give a standard video signal (negative sync level at 0 volts, sync amplitude 250 mV, video amplitude 750 mV, total amplitude 1 volt) with approximately 75 ohm output impedance.

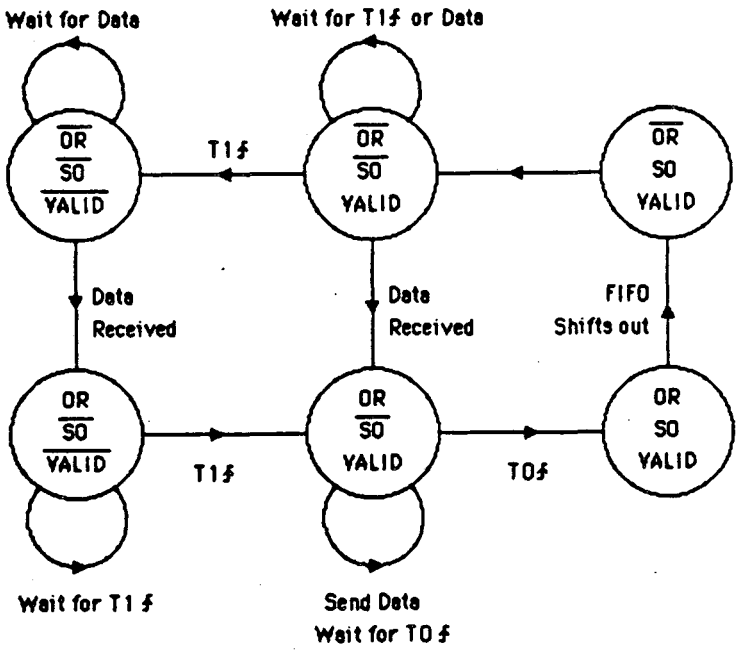

APPENDIX 2 -- FIGURE 5

Note that the actual data format is NRZ (non-return to zero) with no clock information except the two synchronizing bits starting each 32 bit output word. This system thus requires good low frequency performance in the video signal chain (since the digital signal can have a strong DC component) and good speed control (low wow and flutter) in the video recorder. We have not had any trouble with the recorders we use, but cheap video recorders might have trouble, especially if the transmitter were modified for a higher bit rate.

No precompensation (predistortion of the output to cancel out the known distortion of the recording process) is included in the output circuit. One could probably improve the performance of the transmitter with standard video recorders by tailoring the freqency response of the output circuit to match that of the recorder used, or

(with somewhat more effort) by adding "intelligent" precompensation circuits of the sort used in Winchester disk drives, which shift bit transition times slightly to stretch isolated narrow pulses and improve the recovered signal at the receiver.

2.2. Receiver

A block diagram of the videotape receiver is shown in figure 6. The complete schematic is given in figure 7.

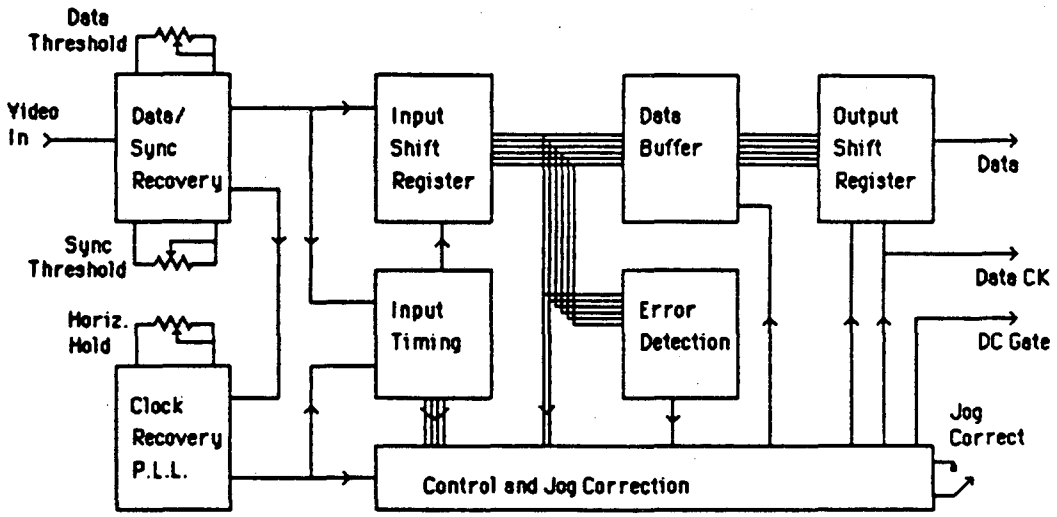

Appendix 2 -- Figure 6
Videotape Receiver Block Diagram

The data and sync recovery circuits buffer the incoming video signal and restore the DC level by clamping the most negative part of the sync signal to a reference value (V/2 ref). This DC-restored signal is fed to a pair of comparators which detect data and sync pulses. The nominal sync threshold is 0.125 volts above the reference level, and the nominal data threshold is 0.625 volts, but both of these are adjustable. In practice, the sync threshold is adjusted until clean sync pulses are produced at the end of the vertical blanking interval (when the DC restorer is most in error), and the

Appendix 2
Figure 7a

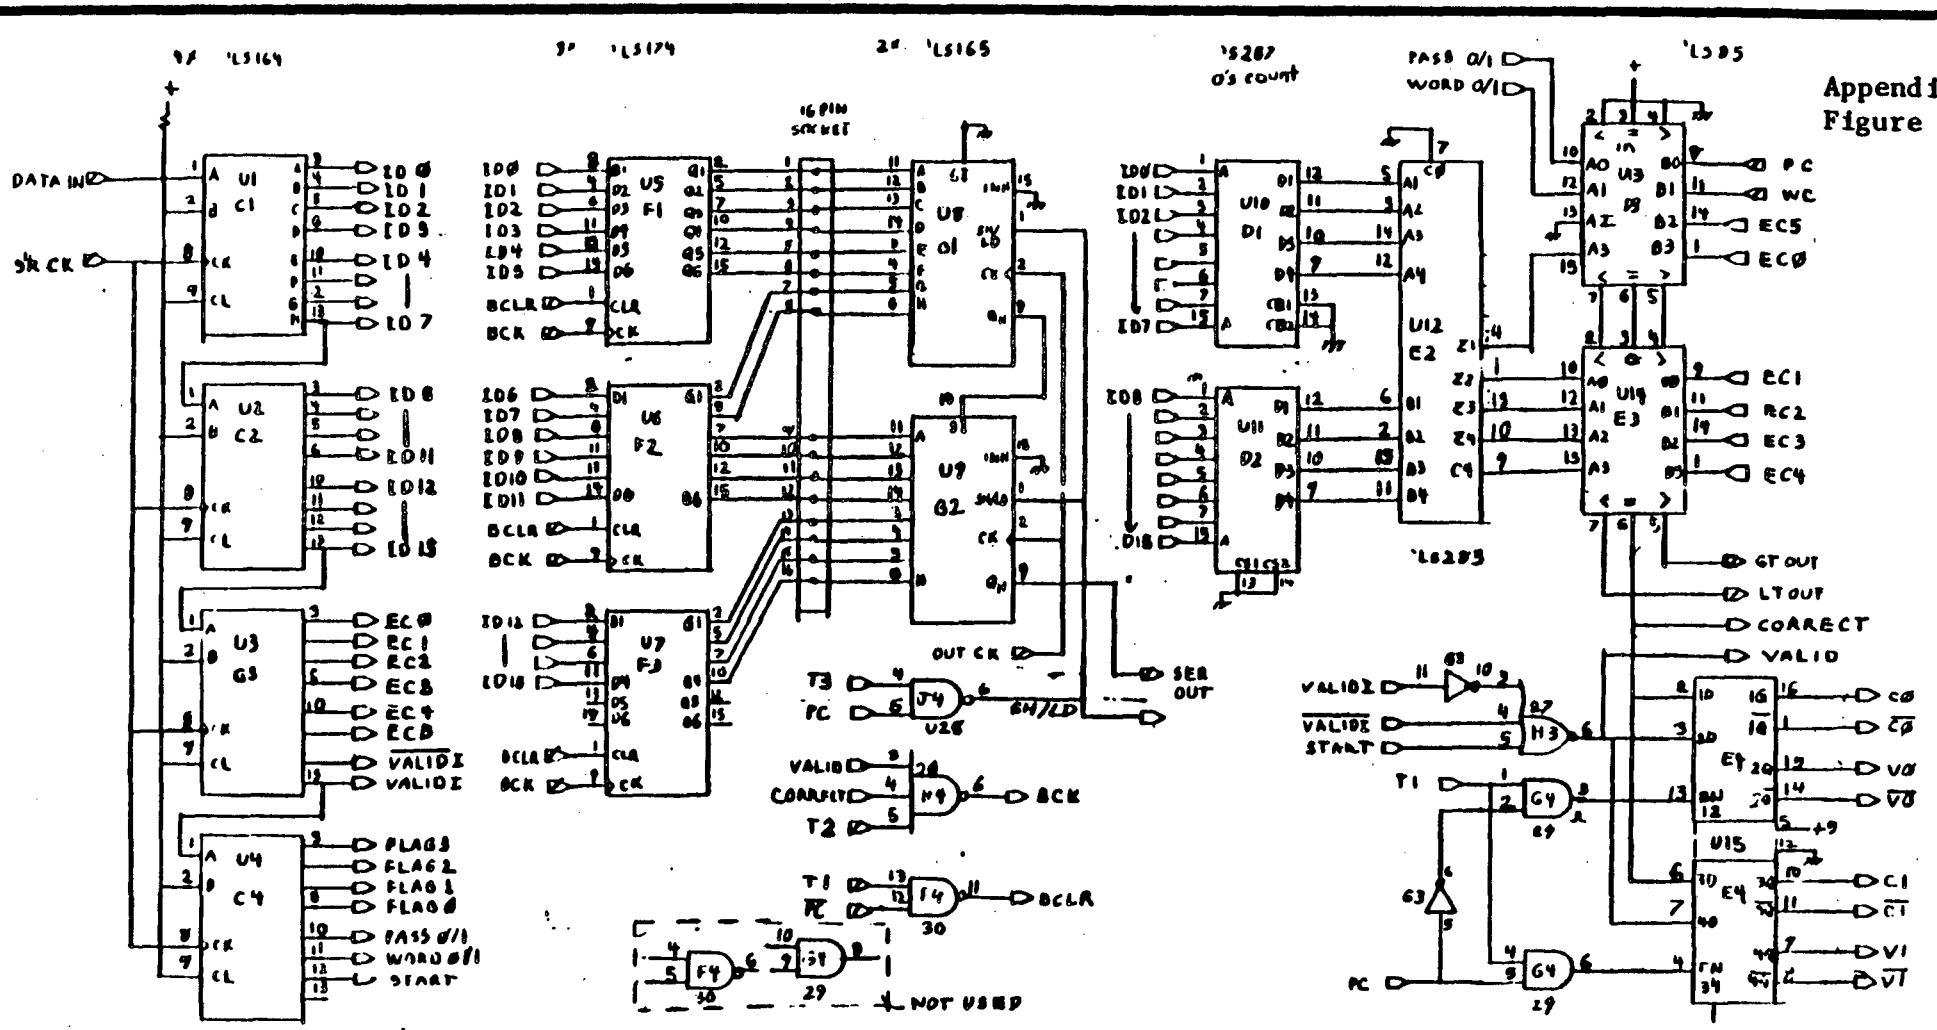

CHG LTR	DRAWN BY	CHECKED BY	DATE	CHANGES
			4/20	Revised as wired

DIST

SHWN ON

ACCT. & SER.

DATE ISSUED | NO. REQD.

DATE REQD.

DEL TO

VIDEOTAPE DATA STORE
RECEIVER
DATA PATH + ECC
SHEET 1 2 5

DRWN. BY J KARE | DATE 3/82

CHKD. BY | SCALE

APPR.

ENG.

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA

7000-88860 (8-78)

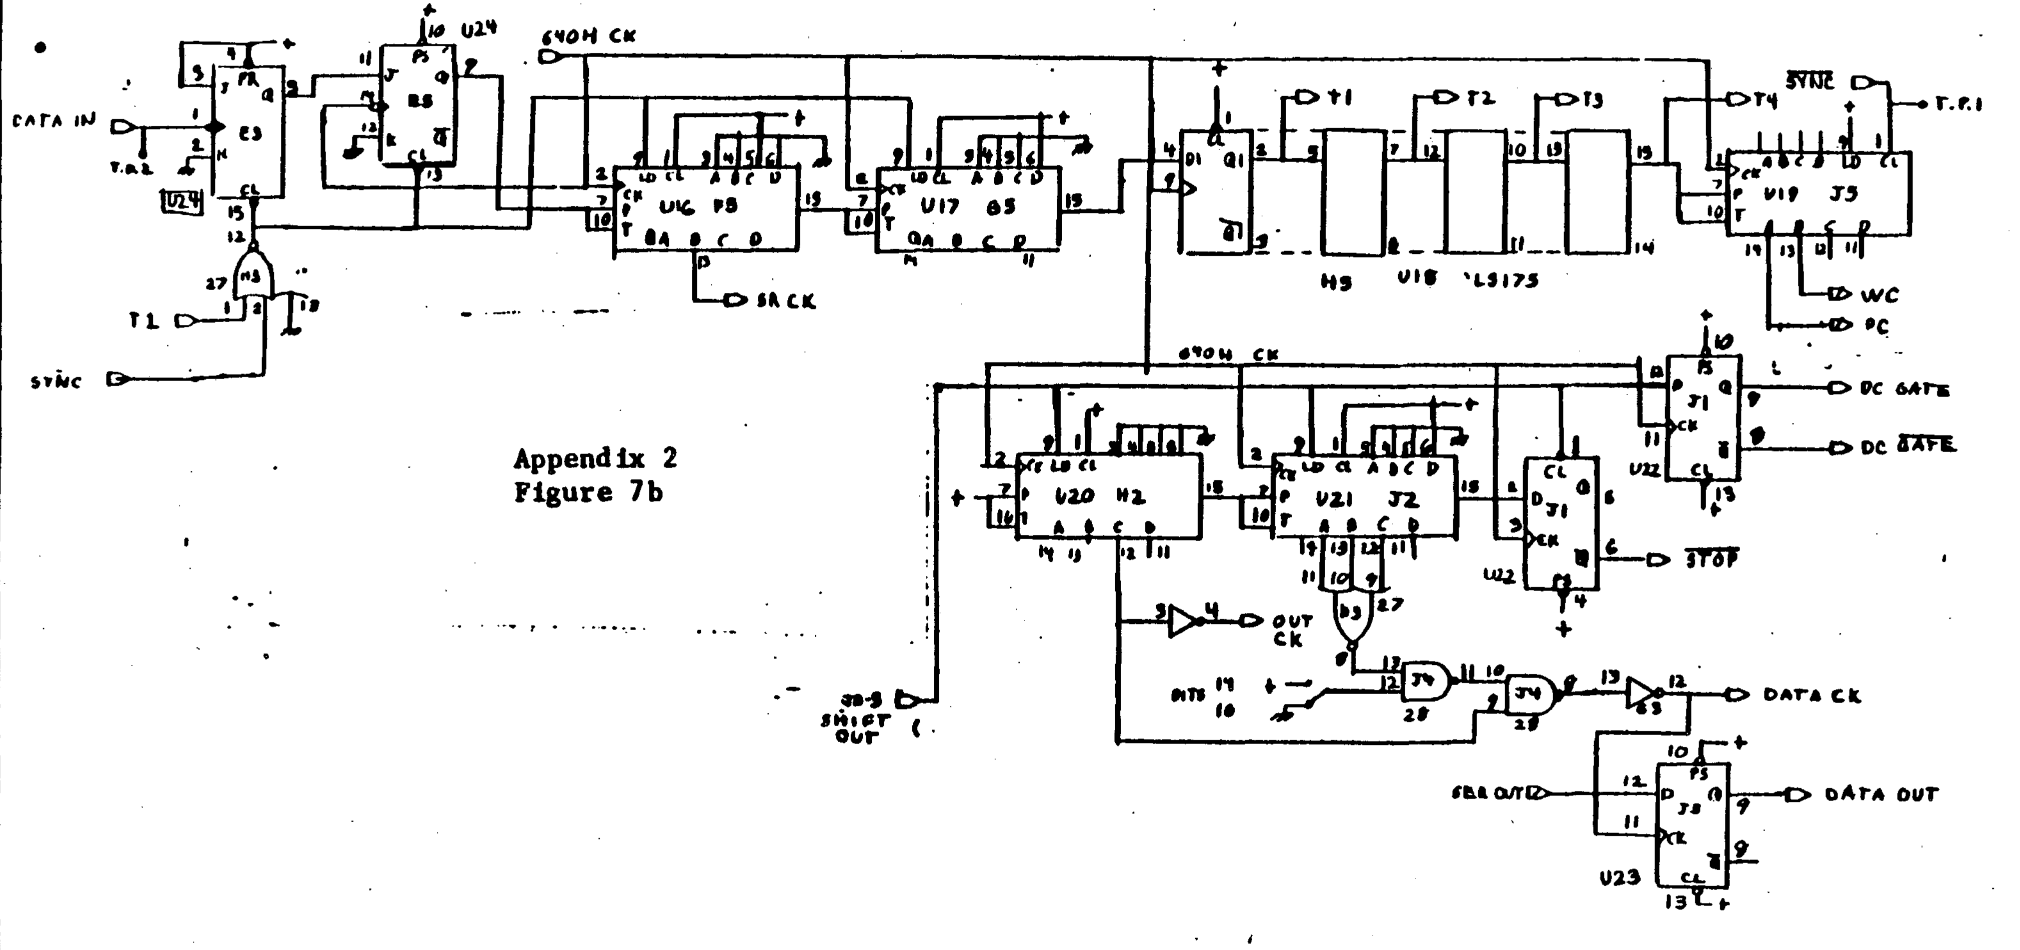

Appendix 2
Figure 7b

U1 - U4	74LS164	U19 - U21	'LS161
U5 - U7	'LS 174	U22, U23	'LS74
U8, U9	'LS 165	U24	'LS112
U10, U11	'S287 O's CT	U25	'LS221
U12	'LS283	U26	'LS10
U13, U14	'LS 85	U27	'LS27
U15	'LS 75	U28, U30	'LS00
U16, U17	'LS161	U29	'LS08
U18	'LS175	U31	'LS04

CHG LTR	DRAWN BY	CHECKED BY	DATE	CHANGES
			4/16	REVISE LOGIC
			5/7	Fix synchronizer
			6/84	Mark additions

VIDEOTAPE DATA STORE
RECEIVER
CONTROL

SHEET 2 of 5

DRWN. BY T KARE DATE 3/82
CHKD. BY SCALE
APPR.
ENG.

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA

7900-03800 (8-78)

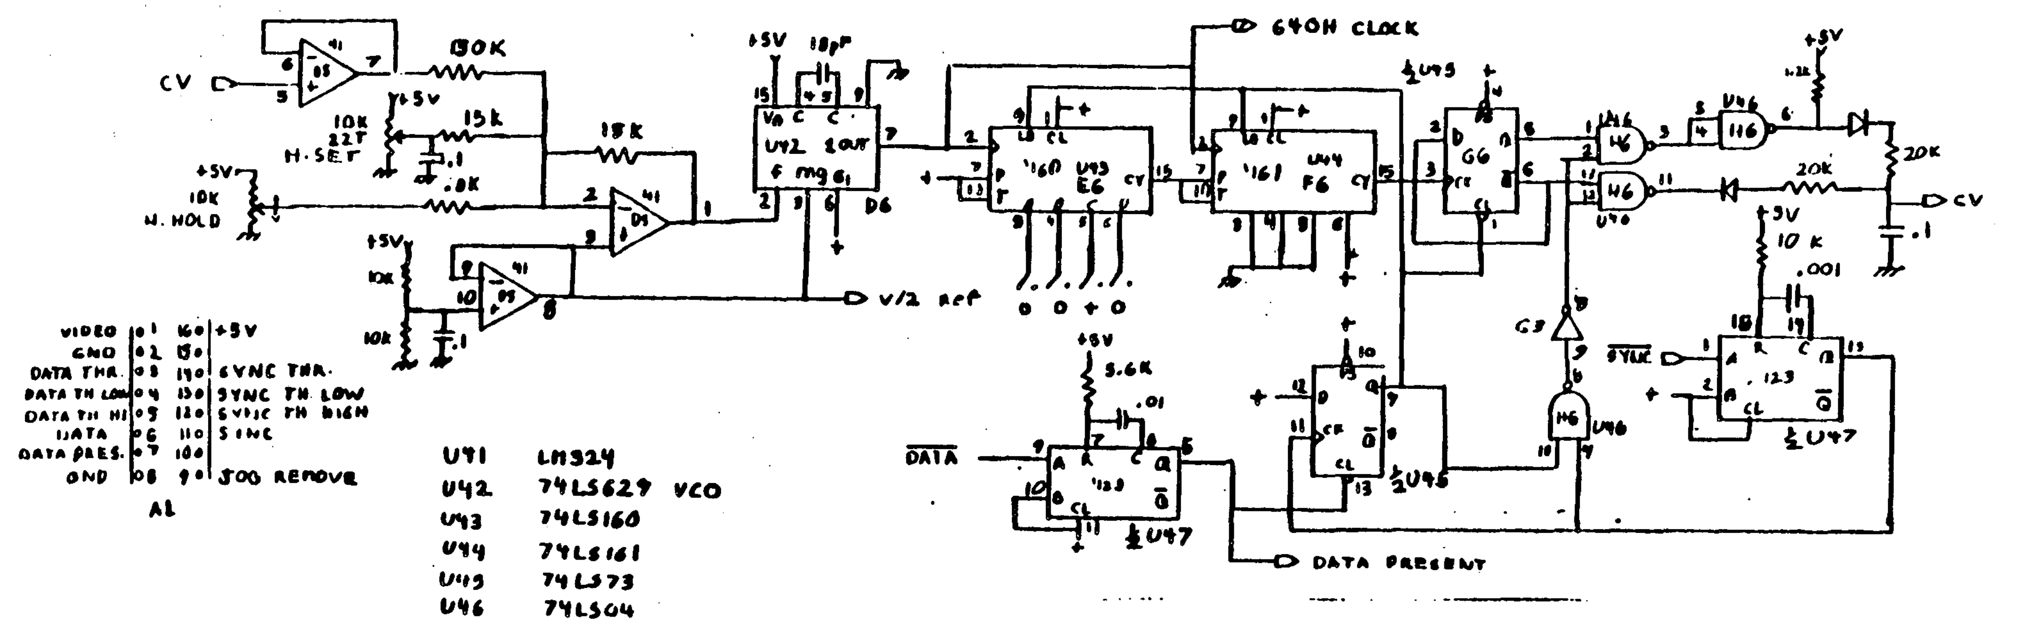

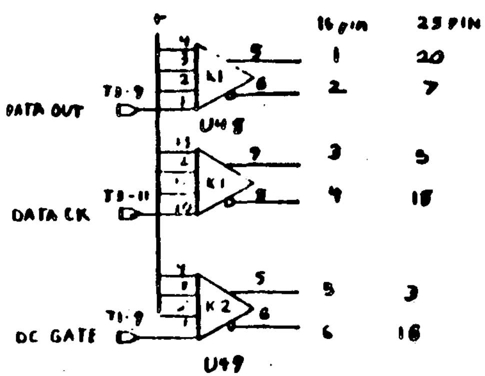

Appendix 2
Figure 7c

CHG LTR	DRAWN BY	CHECKED BY	DATE	CHANGES
			4/20	Revised to as wired
			3/83	lug removed sw

VIDEOTAPE STORAGE UNIT
RECEIVER
INPUT, OUTPUT, CLOCK
sheet 3 of 5

DRWN. BY T KARE DATE 3/82
CHKD. BY SCALE
APPR.
ENG.
SHWN ON
ACCT. & SER.
DATE ISSUED NO. REQD.
DATE REQD.
DEL. TO

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA

7000-02800 (8-76)

PART OF TRANSMITTER

RECEIVER

Appendix 2
Figure 7d

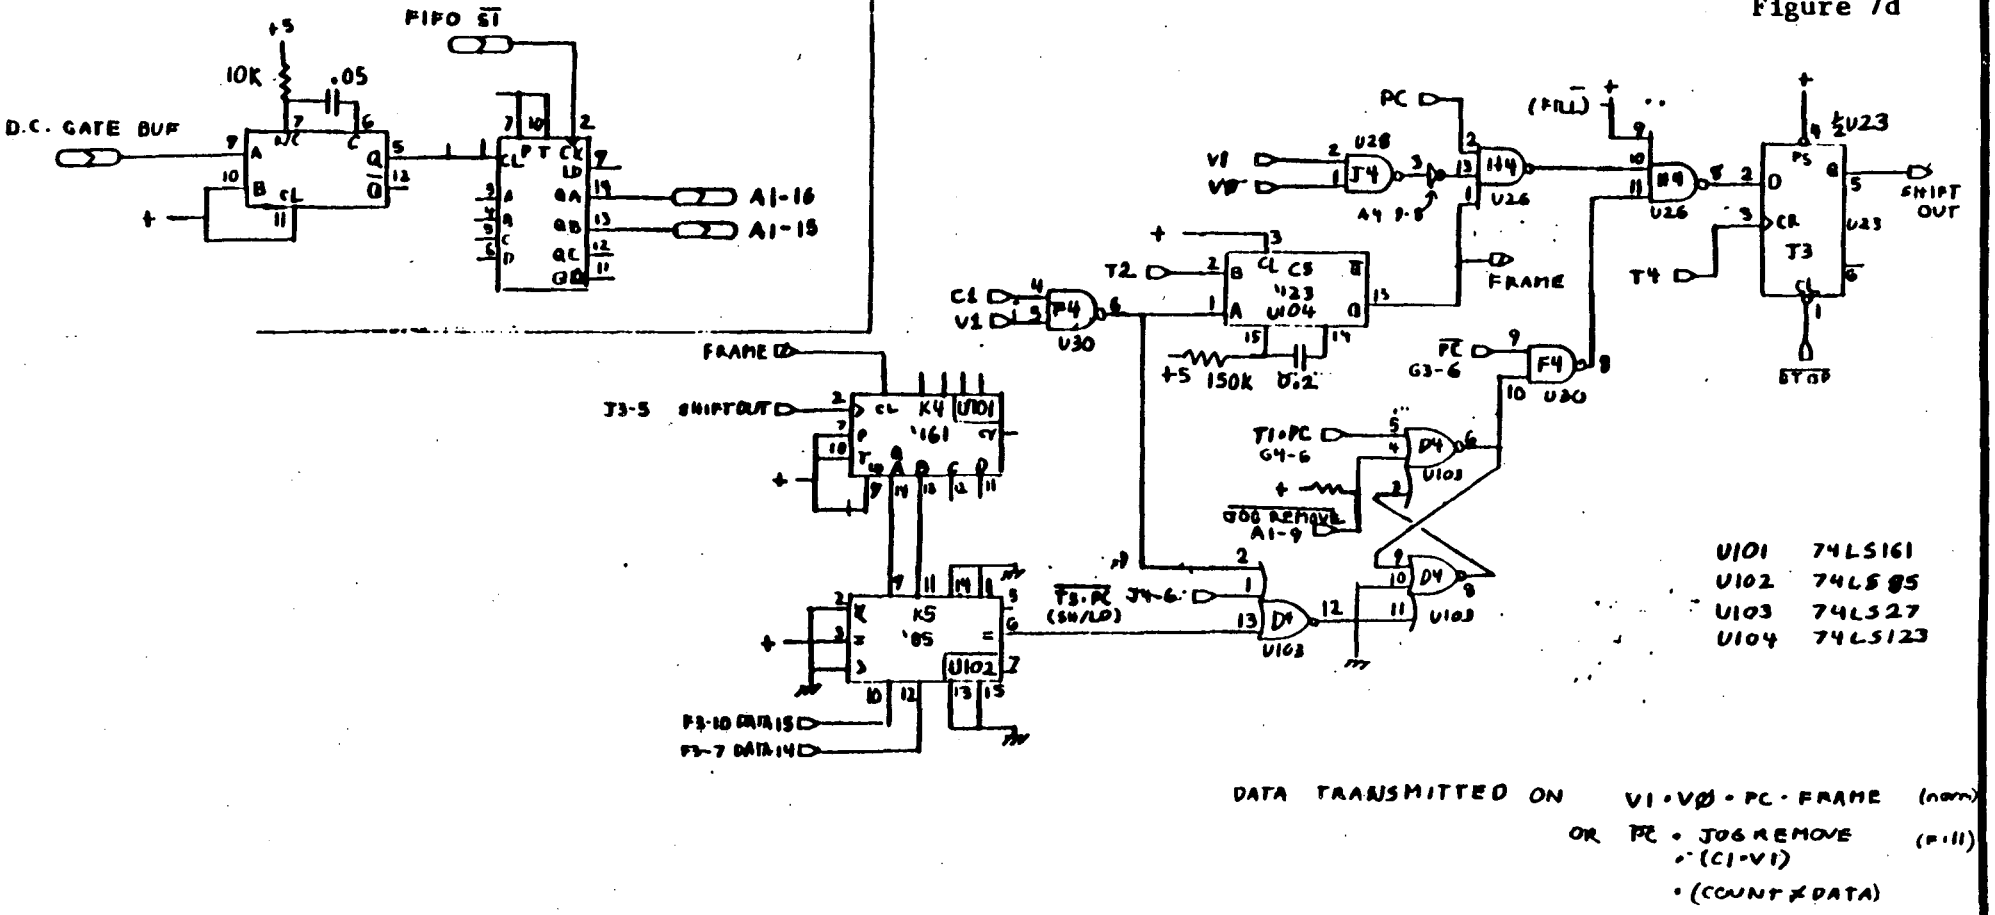

U101	74LS161
U102	74LS85
U103	74LS27
U104	74LS123

DATA TRANSMITTED ON V1·VØ·PC·FRAME (norm)

OR PR · JOG REMOVE (FRU)

· (C1-V1)

· (COUNT ≠ DATA)

CHG LTR	DRAWN BY	CHECKED BY	DATE	CHANGES

DIST I: II: III:

SHWN ON

ACCT. & SER.

DATE ISSUED

NO. REQD.

DATE REQD.

DEL TO

Additions to VIR Trans/Rec.

SHEET 4 of 5 Receiver

DRWN. BY | DATE

CHKD. BY | SCALE

APPR.

ENG.

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA

7600-02800 (8-76)

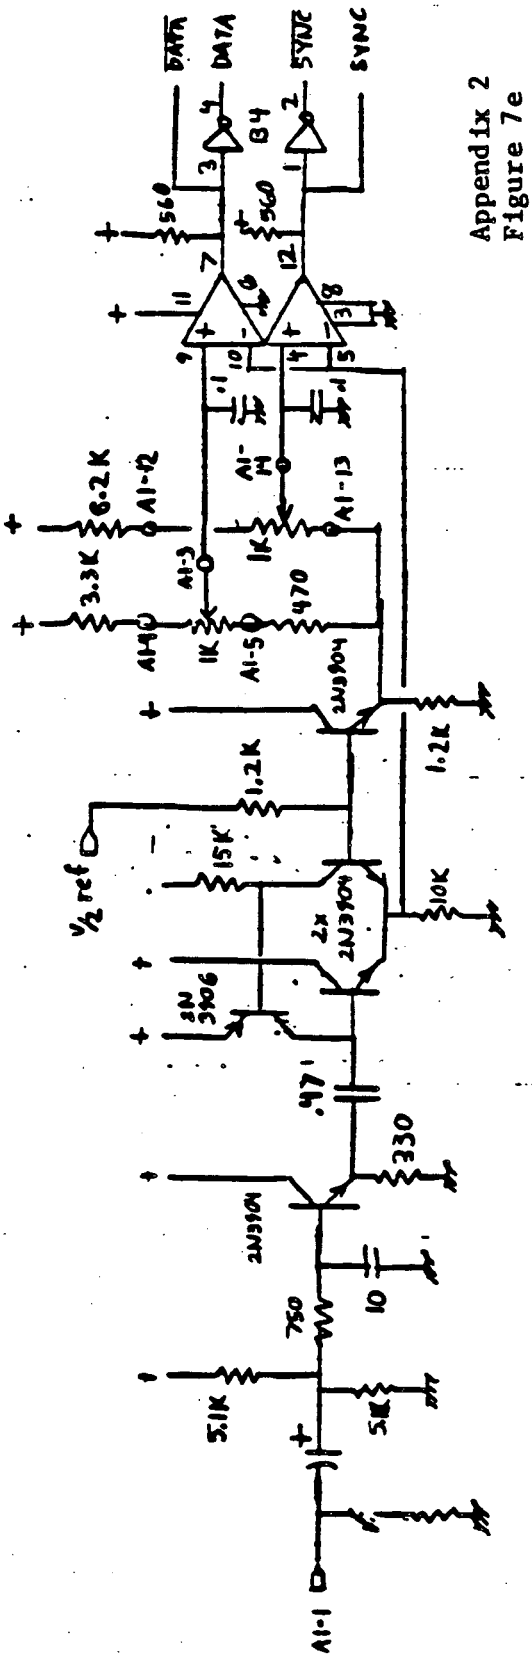

Appendix 2
Figure 7e

VIDEOTAPE RECEIVER
NEW INPUT CKT
SHEET 5 of 5

data threshold is adjusted to give equal width positive and negative single-bit pulses.

Sync pulses are fed to the clock recovery phase locked loop, which is the most complex part of the receiver. The loop oscillator runs at 10.08 MHz, 640 times the horizontal sync frequency. The oscillator output is divided by 320 and the resulting 31,500 Hz clock is compared to a pulse triggered by the recovered sync signal; the loop attempts to hold one transition of this clock at the center of the pulse. As the transition moves away from the center, a net charge is supplied to the loop filter capacitor, adjusting the oscillator control voltage to change the oscillator frequency and drive the transition back to the center of the pulse. 31,500 Hz is used instead of 15,750 Hz (video line rate) so the loop can stay locked despite the half-line shift between even and odd fields in NTSC video (each field is 262.5 horizontal lines). Figure 8 shows the operation of the loop.

This loop should stay locked through the vertical blanking interval, locking to the inverted horizontal sync pulses occurring during blanking. In practice, the loop tended to unlock during the blanking interval, so a shutoff circuit (U47 and related logic) disables the loop and allow the oscillator to free run during this time. In order to rapidly re-lock the loop at the end of the blanking interval, the divider chain producing the 2H signal is pre-set to restart in the correct phase; the frequency drift of the oscillator is small enough that there is no problem recovering data even on the first line after the vertical blanking interval. A horizontal hold adjustment sets the free-running frequency of the VCO in the absence of any signal; this must be adjusted occasionally when the loop begins to unlock spontaneously and the receiver partially or completely ceases to recover data.

The control circuitry detects the first 1-0 transition of the incoming data after each horizontal sync pulse, and starts the timing chain U16-18; the shift register clock

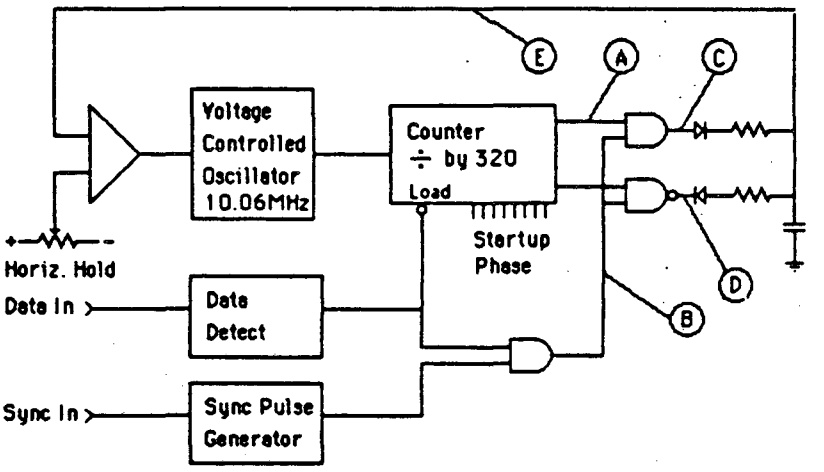

Figure 8 A -- Simplified Schematic of Phase Locked Loop

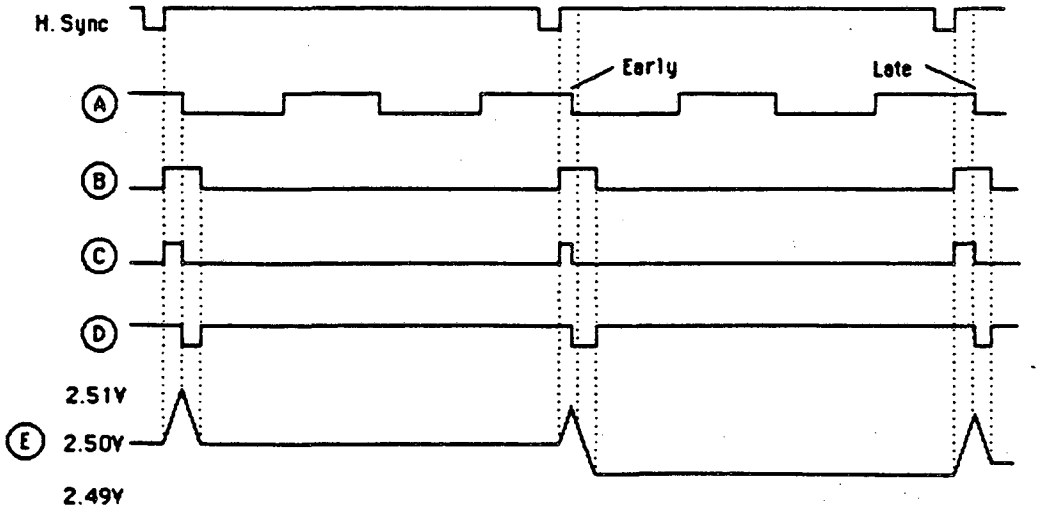

Figure 8B -- Phase Locked Loop Signals

Appendix 1 -- Figure 8
Operation of Receiver Phase Locked Loop

signal (SR CK) is at 1/4 the 640H rate, or 2.52 MHz. Because the timing chain is started by the incoming data, the SR CK signal is automatically set to the correct phase to sample the middle of each data bit -- a common technique for recovering asynchronous data. Incoming bits are loaded into the shift register U1-U4. After all 32 bits of one data word are in the shift register, the timing chain generates 4 timing pulses T1-T4 to control the remainder of the operation.

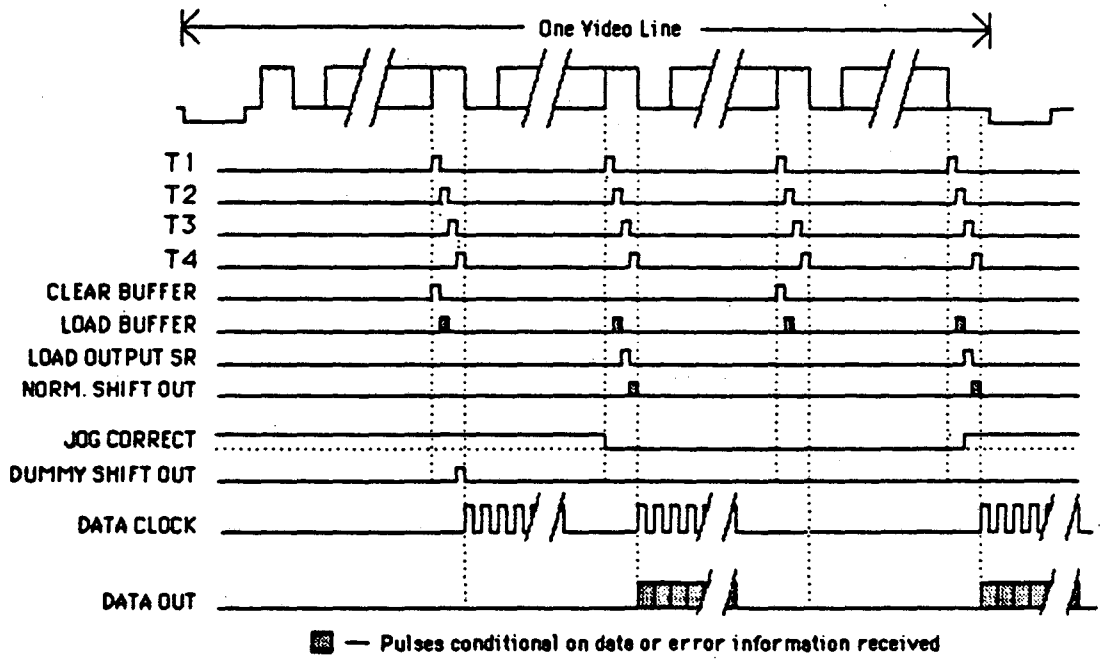

Appendix 2 -- Figure 9
Videotape Receiver Data Output Sequencing

The sequence of events is diagrammed in figure 9. On pass 0, time T1, the buffer (U5-7) is cleared. At T2, the buffer is loaded with data from the shift register if the word is both valid and correct (no errors). Nothing happens at T3 and T4. On pass 1, T1, nothing happens. At T2, the buffer is again loaded with data from the shift register if the word is valid and correct. The buffer now contains either 0 or a

correct data word. At T3, the contents of the buffer are loaded into the output shift register (U8-9). At T4, the shift out process is started if the receiver has detected a valid (not necessarily correct) word. If a valid but incorrect data word has been received, a dummy word of all 0's is shifted out to maintain the alignment of the rest of the image. Currently, the receiver is wired to require both pass 0 and pass 1 to be valid before sending out a word; this relatively strict condition is best because the "jog corrector" circuitry can restore a dropped word, but cannot delete a word added in due to noise.

To reduce the tendency of the receiver to pick up "garbage" while waiting for the start of a CCD frame, the shift out process is blocked until one word both valid and correct is received. One-shot U104 times out if at least one correct word is not received for approximately 10 milliseconds, and must be retriggered by the first pixel of each new CCD frame.

The jog corrector circuit consists of a 2 bit counter (U101) clocked by the shift out signal and a comparator. As long as each incoming pixel is correctly received and sent out, the counter value will match the count stored in data bits 14 and 15. If a valid word is missed due to a dropout, nothing is sent out, the counter is not incremented, and the local count will disagree with bits 14 and 15 when the next word is correctly received. At T3 in pass 1 the state of the comparator is stored; if the local and received count disagree an extra shift-out signal is generated at T4 of the next pass 0, sending an extra dummy pixel (all 0's) to the computer and incrementing the local counter. Up to 3 dropped pixels will be replaced; however, if 4 pixels are dropped the jog corrector will not detect a problem. Also, if the receiver recognizes a noise pulse as a valid word, the jog corrector will insert 3 extra dummy pixels; hence the main circuit is designed to reject false pixels by requiring VALID signals on both pass 0 and pass 1.

The actual shift out circuit is simply a divide by 14 counter clocking the output shift register at 80H (0.63 MHz). The data, data clock, and data clock gate signals are buffered by DM8830 line drivers for transmission to the PDP-11/44.

3. Videotape recorder modifications

The videotape recorder system has been used with an unmodified 1/2" open reel videotape recorder, but consumer videocassette recorders are now much cheaper and easier to use. We selected the Sony SL-2000 recorder as being the most compact and rugged Betamax recorder available on the consumer market when the system was assembled. The Beta recording system was selected primarily because once a cassette is inserted in a recorder and the tape is loaded (wrapped around the recording head) it stays loaded until the cassette is ejected. VHS recorders, the only other type widely available, unload the tape back into the cassette before doing a fast forward or rewind operation, which takes several seconds. Since our use involves a great deal of "shuttling" between different points on the tape, the extra delay and mechanical wear of the VHS system was not desirable. The technical specifications of the Sony recorder were slightly but not significantly better than those of similar VHS recorders.

Two additional features of the SL-2000 are the true real-time counter, which displays tape use in 10 second increments with approximately 1 second accuracy, and the full function wireless remote control unit (part of the TT-2000 power supply/tuner unit). The wireless control unit can be controlled by the PDP-11/44 through an interface which uses CMOS analog switches (CD4066) in parallel with the keyboard switches to activate the various functions.

The major modification of this recorder which must be done is the disabling of the dropout corrector. This circuit uses a 60 microsecond delay line to store one video line. If the videotape recorder detects a large dropout on the tape, it replaces

the current line with the stored line, improving the picture quality. For our purposes, however, it is better to drop pixels than to receive the wrong ones. In the SL-2000, the dropout compensator is disabled by replacing the DOC threshold adjustment potentiometer (RV9) with a fixed 16K resistor.

Other changes do not involve modification of the recorder, and serve only to improve the record and playback response for the digital data signal. Given the limited information in the manufacturer's service manuals, these changes are arrived at by trial and error, varying each trim adjustment in the signal path while monitoring the final reproduced video signal on an oscilloscope. The goal is to obtain clean signals with the fastest possible symmetric rise and fall times (highest frequency response). On the Sony and Zenith recorders, the following adjustments have been made: For recording, the white clip level (RV12) is increased (1/4 turn CCW), and the deviation (RV19) is decreased (1/4 turn CW). Playback is improved by increasing the PB-Y level (RV15, 1/4 turn CCW) and decreasing the de-emphasis (RV10, 1/4 turn CW). As mentioned above, equivalent gains in performance could probably be obtained by modifying the output circuit of the transmitter and the input circuit of the receiver, but adjusting existing trim controls is easier.

APPENDIX III

IMAGE CYCLE ROUTINES

An excerpt from the
Supernova Search Software Manuals

I. Summary of the image cycle

The image cycle is the core of the supernova search software. It processes incoming CCD images, using stored reference data, and yields information on possible supernovae.

In its current form, the image cycle consists of a series of independent tasks (which reside on disk at the moment), each of which performs a single operation. These tasks are:

A. SKYBAK — finds the current sky background.

B. PFFSET — Locates fiducial stars and finds offsets.

C. GGCLSI — 'Cleans' (flat fields) galaxy subimages.

D. LDPRSA — Loads reference subimages from disk

E. GGBLOG — Blurs and logs reference subimages

F. GGSBAC — Subtracts reference from current SI.

G. ??? — Generates list of SN candidates

H. UDLOGS — Updates observation log

In general, there are four main Fortran files for each task. Taking SKYBAK as the example, they are:

A. SKYBAKSPT — Spawn test program -- a control program that can be run directly from a terminal, and that calls the spawn subroutine SKYBAKSPN

B. SKYBAKSPN — Spawn subroutine -- a subroutine which spawns the task itself.

C. SKYBAKTSK — Spawned task header, which handles inter-task messages and calls the main subroutine(s). This is also the file name for the complete task.

D. SKYBAKSBR — Main subroutine for performing the function of the task.

The task itself is task built with the name SKYBAKTSK.tsk, and then installed (made known to RSX-11M) with a 6 character (max) abbreviated name, in this case SKYBTK. Note that if the task name were SKYBAK, it would be impossible to run SKYBAKSPT by typing 'run skybakspt' because (for reasons unknown) the MCR command interpreter will match the first 6 characters of 'skybakspt' with the installed task SKYBAK before looking for a user program called 'skybakspt'. MCR tries to run SKYBAK directly and SKYBAK fails, not having been sent appropriate messages. This problem is avoided by giving the installed task a unique 6 character name. One can also give the installed task the obvious name and write out the full file ID '[311,1]skybakspt.tsk' (or use a command file to do it) -- I did this for LDPRSA, and LDPRSASPT must be run by typing @LDPRSASPT, not 'run LDPRSASPT'. The current installed task names are

SKYBTK
OFSTTK
CLSITK
LDPRSA
BLOGTK

GGSBAC is small enough that it is currently not set up as a spawned task + spawning routine; the test routine GGSBACTST calls subroutine GGSBAC directly.

To install the 5 installed tasks (and some others) after a system reboot or other calamity, log in as a privileged user and type @[311,1]JKINSTALL.

The rest of this manual describes the use of each image cycle task in detail.

[311,2] SKYBAKSPT
Skybakspt
Skybakspn

SKYBTK
Skybaktsk
Skybaksbr

SKY BAcKground calculation

The CCD signal in a given pixel consists, ideally, of 3 parts: the "dark field" value, produced by the CCD in the absence of light, the actual light from stars or galaxies, and the "sky background" of scattered light from the atmosphere, unresolved stars, and so on. The normal flat fielding process consists of subtracting off a standard dark field (presumed not to change with time, though in fact it is apparently quite temperature dependent), multiplying each pixel by an efficiency factor (which includes both actual quantum efficiency and vignetting), and subtracting off the sky background, leaving only the signal from stars and galaxies.

In order to get the correct brightness and position for fiducial stars, we need to know the sky background quite accurately. Using a very simple form of the subroutine GTSTAR, which takes a centroid over a 7 x 7 region, I found that changing the sky background by 5 counts had up to a 1/2 pixel effect on the positions of stars, and up to 20% effect on the calculated sky transmission. This is because that 5 count error is summed over 49 pixels, changing the star brightness by 245 counts. Other versions of GTSTAR are less affected, but it is still desirable to know the sky background accurately.

The simplest algorithm for finding the background is to flat field the entire image and average the flattened image, throwing away all pixels above, say, twice the average. This takes excessive time, and still gets a slightly high value, since some foreground star and galaxy light will be included.

A faster algorithm flat fields only small patches of the image, not the entire image, and includes a test to throw out patches around fiducial stars.

The current version of SKYBAK was a very early effort at writing spawned tasks, and is not by any means a final version. When it was written, the dark and QE fields were still stored on disk, so it was not desirable to access them while finding the background. It processes a set of small subimages, specified in a list file, and checks the appropriate .OBC file to throw out subimages which are close to fiducial stars. In order to avoid using the dark and QE fields, it uses average values of the dark background and the QE for each subimage; these values are pre-calculated by the routine BKLGEN and stored in the list file.

TO USE SKYBAK

RUN [311,2] SKYBAKSPT

SKYBAKSPT requests:

Input region name or abbreviation: region contains image to be processed.

Sky region ID: Currently used only to load .OBC file to avoid area around fiducial stars.

Background list ID: File ID for special background list, generated by the program BKLGEN. List contains one or more items, each specifying a subimage plus a mean dark background and a mean quantum efficiency value.

Tlow, Thi: Threshold values for valid pixels; SKYBAK ignores pixels outside tlow<pix<=thi.

SKYBAKSPT returns:

cskybak	Mean sky background; average of pixels summed
csisma	Standard deviation of pixels; ideally, the image noise level.
npix	Number of pixels summed
nbpix	Number of bad pixels encountered
time	Running time of SKYBAK

[311,2] PFFSETSPT

```
        PFFSETSPT
        PFFSETSPN

PFFSETTSK       (task name OFSTTK)
        PFFSETTSK
        PFFSETSBR
                GGCLBF
                GTSTBF
```

PFFSET is the task which determines most of the properties of a current CCD image: The sky transmission, the seeing, and the offset from a reference image. Its main inputs are the current image itself, a dark field and an inverse QE field, and a reference list of fiducial stars stored in a .OBC (OBject Catalog) file.

PFFSETSPT calls PFFSETSPN, which spawns the task OFSTTK. OFSTTK consists of the header routine PFFSETTSK, which calls PFFSETSBR. PFFSETSBR opens and frames the input region and regions for the dark and QE fields. It loads the fiducial star list from a specified .OBC file and loops thru the list of fiducial stars (up to 10 of them). An area centered on the nominal position of each fiducial star and of size (2 * errw + 2 * intrw + 1) (errw is the expected maximum pointing error and intrw is the size of the integration region for GTSTBF) is loaded into a buffer and cleaned (flat fielded) by subroutine GGCLBF. Once the cleaned buffer is ready, PFFSETSBR calls GTSTBF (GeT STar from Buffer) to determine the actual properties of the star.

GTSTBF is the key subroutine in PFFSET, and if we need to play games with the star positions (due to poor background subtraction, bad pixels, etc.) this will be where we do it. The current standard version takes a simple sum over an integration box (intrw*2 + 1 by intcl * 2 + 1) to get the brightness, and takes first and second moments to get the centroid and the width of the star. An alternate version sums only the center 3 x 3 pixels, and uses a 7 x 7 "hollow box" to find a local background which is subtracted before summing. GTSTBF also returns the width of the star at half max and at a specified threshold; currently the width is just to the nearest pixel, but eventually we should add an interpolation routine.

The values returned by GTSTBF are stored in arrays and processed by PFFSETSBR to produce mean and standard deviation values for offset, sky transmission, and seeing (returns Gaussian sigma = fwhm/2.35). The offset and seeing are returned in pixel * 16 coordinates.

PFFSETTSK files the individual current star values in a file (grid).TL1 (temporary Log 1), and passes the summary values back to PFFSETSPN. PFFSETSPT offers the option of printing out the individual values from the .TL1 file.

To use PFFSET:

PFFSET REQUIRES a .OBC file containing a list of fiducial stars for a given field. This file can be created using EDFSTST, the Fiducial Star Editor.

PFFSET also requires a dark field and an inverse QE field in memory. The default regions are DRKFRM and QEFFRM, which can be loaded with data from disk using GDLOADTST.

PFFSETSPT requests the following inputs:

Sky region	File ID for sky region. Extension will be ignored.
INPUT	Region name or abbreviation for input image.
DARK FIELD	Region name for dark field; "d" gives DRKFRM.
QE FIELD	Region name for QE field; "q" gives QEFFRM.
Dark background	Nominal dark background. Not used except in passing data between GGCLBF and GTSTBF internally; for most purposes, just use 400.
Sky background	Current sky background from SKYBAK
Threshold	Threshold value for recognizing star. Ideally set to 5 * csigma (5 std. dev. above background). For casual use, use 10. Also used to define size of star (width above threshold)

Tlow and Thi are set internally to 1, 16383 to flag 0 pixels as bad.

PFFSETSPT returns the following values:

List of stars:
Ref. #, centroid row, column, total brightness, flag (normally 1)

zskytr	Sky transmission: mean of ratios of current star brightness to reference star brightness. Note that all stars are currently weighted equally (bright stars do not dominate) and that current telescope transmission and exposure time are set to default values inside PFFSETSPT.
zskytv	Sky transmission variation: standard dev. of ratios of different stars (again, unweighted).
zcsens	Equal to zskytr, since (zcexp * zteltr) is set to 1 (zcsens = zcexp * zskytr * zteltr)
cblurx	'Seeing' in 1/16 pixels, defined as mean of FWHM of stars, divided by 2.35. This is approximately sigma for gaussian stars.
offset	Offset between current and reference position, rows and columns, in pixels (division by 16 is done already). Offset is calculated for each star as (current centroid - reference position), and the PFFSET result is the average of the individual offsets, NOT weighted by star brightness.
var.	Variations (standard deviation) in offsets, rows and columns. Note that this is not a very precise statistic for small numbers of fiducial stars, and is intended only to flag very bad fits. Should be less than 0.5.
n star	# of stars found. This may be less than the number of reference stars if bad pixels cause one or more stars to be thrown out. If more than 2 stars are thrown out, something is probably seriously wrong.
errf	Error flag
time	PFFSET running time

[311,1] GGCLSISPT

GGCLSISPT
GGCLSISPN

GGCLSITSK (task name CLSITK)
GGCLSITSK
GGCLSISBR

GGCLSI prepares the current galaxy subimage for subtraction. It subtracts off a dark field value from each pixel, multiplies by an inverse QE value (which corrects for both quantum efficiency and optical vignetting), subtracts off a sky background, scales the image to a specified sensitivity, and adds back a standard dark background (normally 400). It also copies the galaxy subimage into an output region. GGCLSISPT controls the framing of the regions, and frames the output region such that the remaining offset between the current subimage and the reference subimage will be between 0 and 1 pixel on each axis.

GGCLSI obtains dark field and QE information from regions, normally DRKFRM and QEFFRM. GGCLSISPT obtains subimage framing information from a specified .RSP (Reference Subimage Parameter) file; the actual task CLSITK does not access disk at all.

The algorithm for framing the subimage and output region is as follows:

ioffrw = integer offset = (offrx + 16000)/16 - 1000

(the 16000 insures that the rounding process rounds DOWN)

sbrwbg = rfrwbg + ioffrw
sbrwed = rfrwed + ioffrw
sbclbg = rfclbg + ioffcl
sbcled = rfcled + ioffcl
(subimage rwbg = reference rwnbg + integer offset, etc.)

rorwbg = sbrwbg, etc.
(Output rwbg = subimage rwbg)

Example: the reference framing is 101,200,101,200, and the reference galaxy center is at 150,150. The offset is 1.8, -2.4, so the current galaxy center is at 151.8, 147.6. The integer offset will be 1, -3 , (leaving a fractional remainder of 0.8, 0.6 for GGBLOG to fix). The current subimage (and output) framing will be 102,201,98,197. When the output region is RE-FRAMED (by GGSBAC) to the reference framing, 101,200,101,200, the galaxy center will appear at 150.8, 150.6. A fraction of a row of the galaxy is lost; one could save that information by making the subimage 1 row (and 1 column) larger than the reference subimage, but that would make the offset calculations even more confusing than they are now.

To Use GGCLSI

RUN GGCLSISPT

GGCLSISPT requires 4 regions and a disk file. The regions are the Input, Dark Field, QE, and Output regions; the disk file contains the Reference Subimage parameters. In order, GGCLSISPT requests:

Sky region	32 char. filename; extension is set to .RSP
Input region	Contains RAW CCD image
Dark frame	Default is 'd'
QE frame	Default is 'q'
Output region	Must be large enough to contain galaxy subimage; does not need to be larger. Framing info. from REGIONFRM.DAT is disregarded; framing is set to reference subimage framing + offset. Recommended output is CURSIP ('csp' -- Current SubImage, Processed)
Sky background	Value from SKYBAK. This is subtracted off after flat fielding and appears only as an additive constant.
ZCSENS	Current sensitivity. Output will be scaled by 1/ZCSENS.
Offset	Value from PFFSET. This is used with data from the .RSP file to determine the subimage and output framing. Offset should be entered as row*16, col*16; i.e. if offset is 1.75 rows and -.375 col., type 28, -6.

CBKGND, the background added to the output subimage, is set internally to 400. Tlow and Thi are set to 1, 16383.

In addition to the output region, GGCLSISPT returns:

Reference framing:	The framing retrieved from the .RSP file
Output framing:	The framing used for the output region
Scaling:	1/ZCSENS
nbpix:	Number of bad pixels (outside (tlow, thi] encountered
errf:	Error flag
time:	Execution time of CLSITK

** DR0:[311,1]IMAGECYC2.MAN;5

[311,1] CURIMGTST

CURIMGCTL

SKYBAKSPN
PFFSETSPN
GGCLSISPN

CURIMGCTL is a control routine for the first half (roughly) of the image cycle, consisting of the three tasks SKYBAK, PFFSET, and GGCLSI. It spawns the three tasks in order, and takes care of passing data between the tasks. The major inputs to CURIMGCTL are a new raw CCD image, dark and QE fields in memory, and a fiducial star list (.OBC file) The output is a "cleaned" current galaxy subimage, aligned to the nearest pixel with the reference subimage, plus values for offset, seeing, etc.

CURIMGTST is a test routine which allows the user to run CURIMGCTL directly. CURIMGTST loads a .RSP file to determine subimage framing, and accepts other information from the keyboard.

TO USE CURIMGCTL

RUN CURIMGTST

CURIMGTST requires:

Sky region ID	32 character file ID for retrieving .RSP and .OBC files
Bkgnd list ID	32 character file ID for background pixel list for SKYBAK. File must be generated using BKLGEN
INPUT region	Region containing raw CCD image
DARK FRAME QE FRAME	Defaults are "d" and "q"
OUTPUT region	Region which will contain galaxy subimage. Output region framing will be Ref. Subimage framing from .RSP file.
Threshold	Threshold for recognizing fiducial stars. For casual use, use 10. When SKYBAK yields correct values for csigma (background noise), the normal value will be -5 (threshold = 5 * csigma).
Exposure	Exposure in seconds. Telescope transmission is set internally to 0.2. Note that changing the exposure will not change anything except the calculated sky transmission; all the scaling is done based on fiducial star brightnesses.

The output background is set internally to 400; Tlow and Thi are set to 1,16383.

CURIMGCTL (not Curimgtst) provides several blocks of output when compiled in Debug (/de) mode:

Startup information:

arifrm	input framing (normally 1 512 1 320)
arofrm	output " (from .RSP file)
adkfrm	dark field " (normally 1 512 1 320)
agefrm	QE field " "
cdrkbk, etc.	Printout of input and default variable values

After running SKYBAK

CSKYBK	Sky background
CSIGMA	Std. dev. of sky background
NPIX	# of valid pixels summed
NBPIX	# of bad pixels encountered

After running PFFSET

ZSKYTR	Sky transmission (assuming ZTELTR = 0.2, input ZCEXP)
ZSKYTV	Std. dev. in sky transmission among fiducial stars
ZCSENS	ZSKYTR * ZTELTR * ZCEXP
CBLURX	'seeing', gaussian sigma, in 1/16 pixels
OFFSET	Mean Current - reference star positions: row, col
Off. VAR.	Std. Dev. of OFFSET among fiducial stars: row, col
NSTAR	# of fiducial stars matched

After running GGCLSI

Subimage framing	Ref. subimage framing + integer offset
Int. Offset	Nearest smaller integer offset
Frac. Offset	Remaining (positive) fractional offset
Scaling	Scaling factor, ZRSENS/ZCSENS, multiplies each pixel to make current image brightness equal to ref. brightness
NBPIX	# bad pixels in subimage

If fiducial star list printed, for each star

Ref. # centroid row, col, total brightness (not scaled), flag (norm. 1)

```
[311,1] LDPRSASPT       (Run via @ldprsaspt)
                LDPRSASPN

        LDPRSATSK       (task name LDPRSA)
                GCSET
                GCSUBX
                MIDKGN
```

LDPRSA is the task which loads a subimage, stored in RW's format (NOT in .RAI format), from disk into a region, in preparation for blurring and fogging by GGBLOG. LDPRSA performs 3 functions: Clear region, load subimage, and subtract reference background. The clear function is performed by GCSET, and prevents extraneous data from appearing in the border around the subimage (recall that GGBLOG/CONVOLVE acts on up to 3 pixels outside the input subimage when doing a 6 x 6 convolution). The subimage is loaded by RW's routine MIDKGN. Finally, GCSUBX subtracts a constant (rbkgnd) from the subimage. The subtraction ideally sets the reference image background level to 0; this both speeds up and improves the accuracy of CONVOLVE.

LDPRSASPT loads a .RSP file to determine the reference subimage framing and the reference background. IF THE REFERENCE BACKGROUND is NOT SET in the .RSP file, the processed reference subimage will come out with double backgrounds (e.g. a background level of 800).

The output region framing will be set to the .RSP framing PLUS A BORDER for the benefit of CONVOLVE. The border is currently 3 pixels; e.g. if the ref. subimage framing is 101,200,101,200, the output region framing will be 98,203,98,203.

TO USE LDPRSASPT type

@LDPRSASPT

RUN LDPRSASPT WILL NOT WORK! because the installed task name is LDPRSA, and MCR will find that name first and try to run the installed task directly. You can enter RUN [311,1]LDPRSASPT.TSK, which is what @LDPRSASPT does.

LDPRSASPT requires the following inputs:

OUTPUT region	Region to put subimage in. MUST BE LARGE ENOUGH to hold ref. subimage plus 3 pixel border.
Sky Region ID	32 char. file ID for .RSP and .RSA (ref. subimage array) files.

TLOW and THI are set to 1, 16383 internally. All bad pixels are set to the background level (i.e. 0), to minimize effect on CONVOLVE, which DOES NOT test for bad pixels at this time.
A later version may set bad pixels to an average of neighboring pixels.

LDPRSASPT returns the reference subimage in the output region, plus the following information:

ri frame	Actual framing of reference subimage
rbkgnd	Background from .RSP file
nbpix	# of bad pixels (outsid (tlow, thi]) encountered
time	running time of LDPRSATSK
errf	error flag

[311,1] GGBLOGSPT
GGBLOGSPN

GGBLOGTSK (task name BLOGTK)
GENCOA
CONVOLVE

GGBLOG performs the functions of blurring the reference subimage to match poorer current seeing, and jogging the subimage by less than 1 pixel.

GGBLOG uses CONVOLVE to convolve the input array with a 6 x 6 array generated by GENCOA (GENerate COnvolution Array. Element i+3, J+3 of the convolution array is given by

```
c(i+3,J+3) = sum over dx = 0, 1;  dy = 0, 1
             [   ERF ((i + dx + 0.5) * (sigmax/sqrt(2)))
               -  ERF ((i + dx - 0.5) * (sigmax/sqrt(2))) ]
         *   [   ERF ((J + dy + 0.5) * (sigmay/sqrt(2)))
               -  ERF ((J + dy - 0.5) * (sigmay/sqrt(2))) ]
         *  jog(dx, dy)
```

where

sigmax = std. deviation "width" of gaussian distribution by which image is to be blurred

```
jog(0,0) = (1 - offsetx)(1 - offsety)
jog(0,1) = (1 - offsetx)(offsety)
etc.
```

and ERF(x) is a polynomial approximation to the error function

The array returned is a fixed point array, scaled so that the sum of all pixels is 1000 (+/- a roundoff error of up to 36). If the blurring width is too large (bigger than about 1 pixel), significant light would be "lost" from a 6 x 6 convolution; this overflow is ignored in the rescaling process, so one actually convolves by a truncated gaussian scaled to unit area, instead of by the full gaussian.

Once the convolution array exists, GGBLOG calls CONVOL to perform the convolution. CONVOL simply sums the products of 36 pixels from the input region with the 36 elements of the convolution array, divides the sum by 1000, and stores the result in the output region. To minimize remapping memory, CONVOL processes blocks of pixels together, and does some tricks such as testing input pixels and discarding those which contribute negligably (less than 1/2 count) to the output.

LIMITATIONS OF GGBLOG

GGBLOG is a slow routine, simply because straightforward convolution involves lots of arithmetic. With no arithmetic, it takes 2 seconds to load and store a 100 x 100 array. With just jossing (so the convolution array has only 4 nonzero elements), a 100 x 100 array is processed in 4 seconds (approx.). With a full 6 x 6 convolution, a 100 x 100 array takes upwards of 10 seconds.

One approach to speeding this up is to take advantage of the orthogonal nature of gaussians (and of our current jossing process) and do separate X and Y convolutions, reducing the worst case calculations from 36 multiplies to 12 at the expense of some extra memory mapping and copying. Another is to store two or more reference pictures with different blurrings, to reduce the amount of additional blurring needed.

Jossing an image introduces additional blurring; a single bright pixel is spread over 4 pixels in the output. The worst case of this occurs when, in the reference image, a star lands on the junction of 4 pixels. When the image is jossed, the star is spread over 9 pixels, while in reality the star may fall almost entirely in 1 pixel:

GGBLOG						Reality (extreme case)				
100	100		25	50	25	100	100		0	0
100	100	=>	50	100	50	100	100	=>	0	400
			25	50	25					

One cannot completely prevent this; in the other extreme (a perfectly square, flat topped star!) GGBLOG would be right. Given an undersampled image (in the Nyquist sense) one cannot do perfectly well. However, given some knowledge of the shapes of star images, one should be able to do better, and slightly fancier convolution-based algorithms are in the works. Still fancier processes (fourier transforms, fitting routines, etc.) could do better still, with additional complexity (though not necessarily much additional time; ideally (meaning if we had a 32 bit machine) 2D 128 x 128 fourier transform convolution would involve only (7 + 7 + 1 + 7 + 7) = 28 multiplications per point. Unfortunately, a) we don't have a 32 bit machine and b) they're COMPLEX multiplications). However, the easiest thing to do is do as well as is convenient, and then collect and ignore the remaining bad pixels....

One thing which is very easy to do is alter the blurring algorithm so that, if one IS blurring the reference image, the amount of blurring is reduced as the amount of jossing increases. At the moment, however, the reference seeing is no better than the current seeing, so that doesn't help any.

TO USE GGBLOG:

RUN GGBLOGSPT

GGBLOGSPT requests:

INPUT region Region containing image to blur and jog. Background level, and all bad pixels, should be set to 0, e.g. by LDFRSA. Region must be at least 3 pixels larger than subimage on all edges.

OUTPUT region Need only be large enough to contain subimage.

Subimage framing

NOTE: GGBLOGSPT DOES NOT access a .RSP file to find subimage or region framing. Therefore, one will in general have to enter peculiar framing values for input, output, and subimage framing. For example, a reference image with framing 100,200,110,190 may be stored in region REFSIM, region framing 97,203,107,193 (remember the 3 pixel border), so one would have to enter for:

```
INPUT
        *               (specifies user input region)
        REFSIM          (region name)
        97,203,107,193  (region framing)
OUTPUT
        *
        REFSIP
        100,200,110,190
SUBIMAGE
        100,200,110,190
```

Desired blur Sigma of desired gaussian blur, for rows (vertical blurring) and columns (horiz. blurring), in 1/16 pixels. GGBLOGSPT does not keep track of reference image seeing, so this is the ADDITIONAL blurring requested, not the desired seeing in the output subimage.

Desired jogging In 1/16 pixels, range 0 to 16, rows and columns. Images are jogged in the direction of increasing row and column; e.g. a star centered on the middle of pixel 100,100 and jogged by 8,8 will be centered at 100.5, 100.5 (junction of 4 pixels) in the output.

Ref. background Constant added to output subimage; normally 400

Tlow and thi are set internally to 1,16383, but CONVOLVE currently does not check for bad pixels at all.

GGBLOGSPT returns:

BLROVF Blur Overflow; the amount of light which falls outside a 6 x 6 convolution array. If this is larger than about 100, the gaussian convolution array is badly truncated at the edges.

Error Error flag

Time Execution time of GGBLOGTSK

This report was done with support from the Department of Energy. Any conclusions or opinions expressed in this report represent solely those of the author(s) and not necessarily those of The Regents of the University of California, the Lawrence Berkeley Laboratory or the Department of Energy.

Reference to a company or product name does not imply approval or recommendation of the product by the University of California or the U.S. Department of Energy to the exclusion of others that may be suitable.

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