THE DESIGN AND CONSTRUCTION OF A CHARGE-COUPLED DEVICE IMAGING SYSTEM

DONALD H. GUDEHUS AND DENNIS J. HEGYI
The Harrison M. Randall Laboratory of Physics, The University of Michigan, Ann Arbor, Michigan 48109
Received 20 June 1984; revised 5 October 1984

ABSTRACT

The design and construction of a charge-coupled device (CCD) imaging system optimized for the detection of low surface brightness objects is discussed in detail. In particular, we describe the design of the Dewar and the CCD electronics, as well as the choice of CCD operating voltages, clock timings, and operating temperature. In addition, we discuss the measurement of readout and spatial noise for an RCA SID53612 thinned CCD (320 by 512 pixels). For this chip, we find a readout noise of 64 electrons and obtain a well depth in excess of 350 000 electrons with essentially no blemishes or bleeding of charge on bright star images. The threshold level for trailing of charge is 300 electrons in the center of the chip and is essentially zero in the corner of the chip which is read out first. On broadband exposures, fringing from night-sky lines is typically seen only in the infrared ($\sim$8700 Å) at about 2% of the night-sky level. Our twilight flat-field exposures taken on different nights are consistent with each other at a level varying from one part in 100 000 to a few parts in 10 000 on the large scale, and the zero level is stable to $\pm$ 10 electrons over a night. The system is linear over its entire range, and we can currently do photometry to better than 0.8% using standard stars.

I. INTRODUCTION

The invention of the charge-coupled device (CCD) by Amelio, Tompsett, and Smith (1970), and its subsequent development over the past decade, has provided astrophysicists with an almost ideal photon detector between the wavelengths of 0.35 and 1.0 μ. A CCD is a solid-state device which, when struck by photons, will produce electron-hole pairs with high efficiency. The electrons are stored in a mosaic of electrostatic potential wells formed by a polysilicon electrode structure and ion implants, while the holes are repelled from the wells. In the thick version of the RCA CCD (SID52501), the photons must pass through the polysilicon electrodes whch are highly absorbing in the UV, whereas in the thinned CCD (SID53612), the silicon substrate is etched to extreme thinness so that the photons can enter through the backside of the chip, avoiding the absorptive electrode structure. In both CCD's, the entire imaging area is light sensitive; there are no "dead regions" along the edges of pixels. Also, because the thick CCD is many times more sensitive to cosmic rays than the thinned chip, the thinned CCD has been preferred for use as a photodetector by astrophysicists when long integration times are required.

To read out an image stored in the CCD, the electrode structure which helps form the potential wells is used to transfer the charge from the pixels to a charge-sensitive output amplifier located on the chip. In operation, charge is first shifted vertically, one row at a time, to a horizontal register. The charge is then shifted out of the horizontal register one column at a time to the on-chip amplifier consisting of a small capacitor ($\sim$0.5 pF) whose voltage is sampled by an on-chip FET source follower. The change in voltage of the capacitor is directly proportional to the amount of charge deposited, and, consequently, is directly proportional to the number of photons incident on a pixel.

The advantages of a CCD are high quantum efficiency (80% peak), spatially digitized output, geometric stability, large dynamic range (5500), and excellent linearity. All existing CCD's, however, suffer from nonzero readout noise ranging from about 10 to 100 electrons and, for very low light levels, a less than perfect charge-transfer efficiency which gives rise to a "threshold effect" which can complicate applications to spectroscopy and astrometry. However, at higher background light levels, the threshold effect is no longer a problem, and at still higher background light levels, the readout noise is dominated by shot noise. Blemishes and fringing (the interference of light between the front and back faces of a CCD) can further complicate the data analysis. From the experiences of various groups who have built CCD imaging systems (Low 1977 and 1981; Dewey and Ricker 1980; Leach, Schild, Gursky, Madejski, Schwartz, and Weekes 1980; Marcus, Nelson, and Lynds 1979; Meyer and Ricker 1980; Robinson 1981; Wright and Mackay 1981), from our own experience, and from inspection of published CCD pcitures, it is clear that the above problems are serious. Accordingly, we shall discuss aspects of the CCD system relevant to these issues in greater detail.

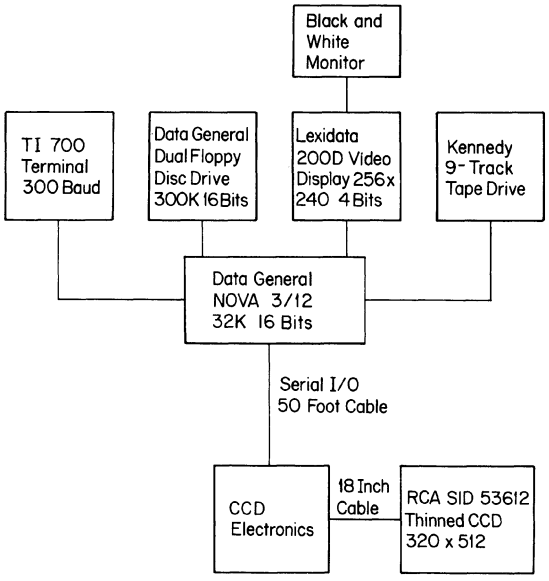

FIG. 1. Block diagram of the CCD operating system.

 0004-6256/85/010130-09$00.90

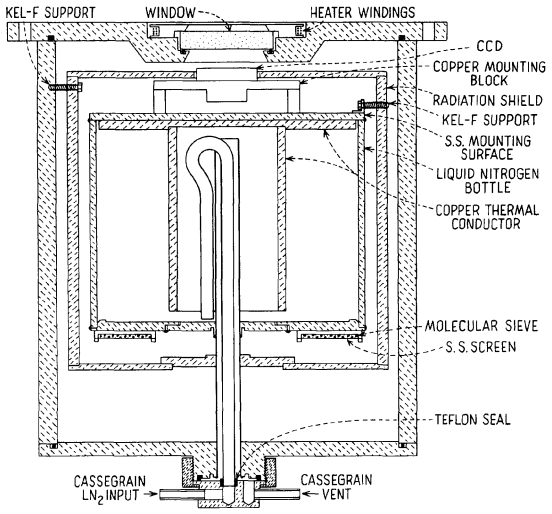

FIG. 2. Schematic diagram of the CCD Dewar.

II. SYSTEM CONTROL

Figure 1 shows an overview of our CCD system. We use a Data General NOVA 3/12 minicomputer to provide the overall control of the CCD system. The selection of a NOVA as a host computer may seem somewhat dated at the present time, but in 1978, when the design and construction was started by one of us (DH), it was a reasonable choice and was available from another project. The NOVA supervises an exposure by loading the memory of the CCD control microprocessor located near the CCD by opening, timing, and closing the shutter, by initiating the readout, and by transferring the data to the magnetic-tape unit for storage. In addition, data can be transferred from tape to floppy disk for processing, to a TV monitor via the Lexidata for display, or to the TTY. All the NOVA programming has been carried out in FORTH which has proven to be very convenient.

We shall now discuss the system in more detail. Communication between the NOVA minicomputer, located in an observing room near the dome, and the CCD electronics, mounted beneath the Dewar on the telescope, is via a serial link over a 50-ft cable. The NOVA transmits (downloads) machine code over the cable in the form of 256 48-bit numbers to the memory of an AM2910 microprocessor which generates the necessary timing signals for control of the CCD. We chose this particular microprocessor because it was the fastest microprocessor available at the time we designed the system. Its cycle time is 100 ns. Additional machine code is used for control of the shutter, for the filter wheel, and to adjust voltages on the CCD.

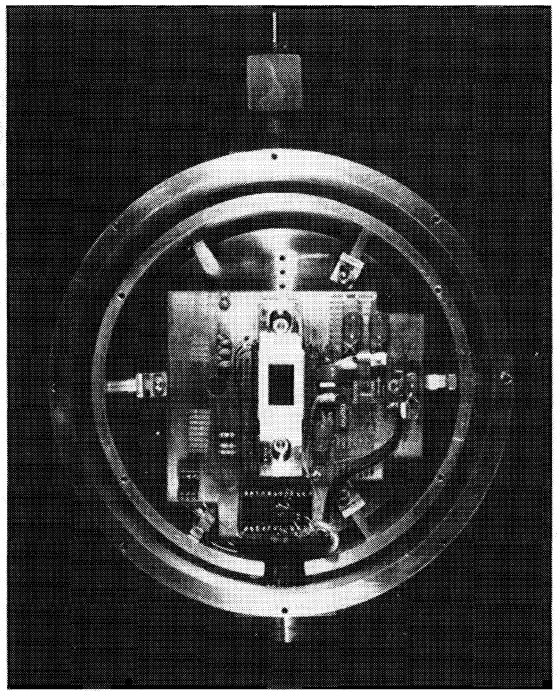

FIG. 3. View of the top of the opened Dewar.

Readout of the CCD, sampling and digitization of the output, and transfer to magnetic tape is accomplished in 50 s by making use of a double buffer programmed into the NOVA's memory. During an exposure, the time remaining in the integration is displayed on the monitor. A maximum of 55 full-frame exposures can be stored on a 1200-ft reel of tape.

Through the TI Silent 700 hard-copy terminal, a user can enter a variety of commands into the system for displaying gray-scale data, graphing data, rapidly reading out a 64×64-pixel region for focusing the telescope, performing speckle interferometry, doing photometry, and performing picture arithmetic. (The hard copy has proven useful for resolving discrepancies in the observing log.)

Under software control, the user may display any four bits of the 14 bits of digitized intensity information from any 256×240-pixel subsection of the CCD frame by loading the data into the Lexidata 200D video generator. The user can also graph horizontal, vertical, or diagonal cuts through the data to further examine a feature. These cuts can be used to calculate the full-width-half-maximum (FWHM) of a feature, which is useful for several modes of operation, but most typically, focusing. In the focusing mode, a repetitive cycle is

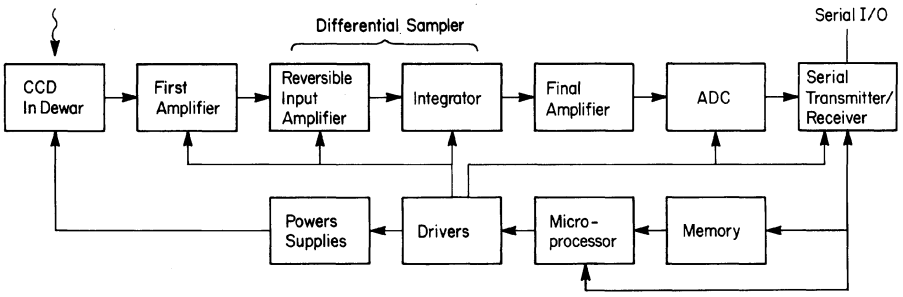

FIG. 4. Block diagram of the CCD electronics.

established in which a 64×64-pixel section of the CCD surrounding a star or a spectral line is displayed and optional cuts through the image in the horizontal and vertical directions with corresponding FWHM's may also be displayed. The 64×64-pixel region is also used for speckle interferometry.

A cursor can be displayed on the video monitor to mark objects to be photometered, pixel values to be read out, and regions to be averaged. Some simple picture processing routines are available so that subtle features in images may be studied while still at the observatory to see if additional exposures are necessary. Routines are available to perform flat-fielding, dark current and bias subtraction, and to co-add images. Fully automatic operation is possible for which the user is not required to be present; this is useful for taking bias frames and dark frames.

III. THE CCD DEWAR

The liquid nitrogen Dewar for the CCD is a cylindrical chamber 7.75-in. in diam. and 8.5-in. in length with a 9-in. diam. faceplate (see Fig. 2). It is designed to work in an upward- or downward-looking position by interchanging only the external fill and vent tube connections. The CCD socket is screwed to a copper block and the CCD chip is held in place with two copper fingers. The copper block is attached by means of stainless steel standoffs to a 1-liter welded stainless steel liquid-nitrogen bottle. The nitrogen bottle is supported by a combination of Kel-F standoffs and by the 0.5-in. diam 20-mil stainless steel fill/vent tube. Three additional Kel-F rods connect the cylindrical radiation shield to the outer Dewar wall. Flexure is undetectable in this arrangement, and the nitrogen holding time is not seriously degraded by our rigid structure; one filling lasts for 10.5 hr. A photograph of the top of the opened Dewar is shown in Fig. 3.

IV. THE CCD ELECTRONICS AND LOGIC

Figure 4 shows a block diagram of the CCD electronics. The digital signal-processing circuitry is enclosed in a heavily shielded and temperature-controlled chassis divided into separate sections and connected to the Dewar by a triply shielded 18-in.-long cable. Power supplies are mounted outside the box and power is brought in through filtering feedthroughs. Although we initially placed a preamplifier inside the Dewar, we found that greater stability could be achieved without it and at no increase in readout noise. Our output bias level is presently stable to ± 10 electrons over a night.

In the following paragraphs, we discuss the preparation and readout of the CCD and trace the signal through the electronic circuitry. In order to prepare the system for an exposure and readout, the NOVA downloads a variety of subroutines in machine code into the microprocessor memory. The first microprocessor memory location is reserved for a jump instruction which moves the program pointer to the desired program. When the Nova initiates a particular command, it only loads the address of the chosen program into the first microprocessor memory location and starts the microprocessor. An exposure and its preparation consist of the execution of FORTH commands on the NOVA and microprocessor programs which (i) rapidly read out the CCD six times, after which all residual charge from a previous exposure is flushed from the system, (ii) apply the correct voltages for an integration of the CCD electrodes, and finally, (iii) open the shutter for a predetermined period of time. Immediately upon completion of an exposure, a different jump instruction is sent to the CCD and it is clocked in its normal slow-readout mode.

We turn now to the details of shifting the charge on a pixel across the CCD substrate to the output amplifier by stepping voltages on the horizontal and vertical electrodes, and to the sampling and digitization of the charge. The charge on a pixel is most accurately determined by performing a differential sample of the voltage of the on-chip CCD capacitor before and after charge is transferred to it. In our system, we use an up-down integrator for this task.

We shall outline the essential steps in the differential sample, referring the reader to Figs. 5–7 for more specific details of the waveforms. (Our differential sampler is based on a design by Marcus *et al.* 1979). First, the CCD capacitor is reset to a reference voltage, and that voltage is sampled by up-integrating, when switch #2 (integrator) is turned on. Then the integrator is turned off and signal charge is transferred to the CCD capacitor by driving the three horizontal electrodes through a cycle. To perform the down-integration, switch #1 is reversed, and the integrator is turned on again. After the end of the down-integration, the voltage on the integrator is digitized to 14 bits and the integrator reset with switch #3 in preparation for the next cycle. After 320 horizontal advances, charge is shifted by one row in the vertical direction loading the horizontal register, and another 320 horizontal shifts are performed.

As may be seen in Fig. 7, we use three distinct voltage levels on the verticals, in contrast to other groups that use only two voltage levels. The intermediate state is used on all vertical electrodes during the exposure and on vertical electrode #2 during the horizontal transfer. We found that the presence of an intermediate state increases the charge-transfer efficiency and the depth of the potential wells.

Including the timing variables shown in Fig. 8 which show the 8 τ's for the fast reads used for reading the area surrounding a 64×64-pixel area, there are a total of 29 timing values. These values are software variables and, consequently, are easily changed. There are nine independent electrode voltages which are set using hardware.

Trimming the differential sampler makes use of both hardware and software adjustments. The product of gain ×time must be the same for both the up- and down-integrations during sampling. Three potentiometers (not shown) in the output arms of the switches in Fig. 5 can be adjusted to balance the gain. Final adjustment is most conveniently made by fine-tuning the ratio of the time spent in the up and down portions of the integration. For example, our time interval for up-integration τ_3 is 124.7 μs and the time interval

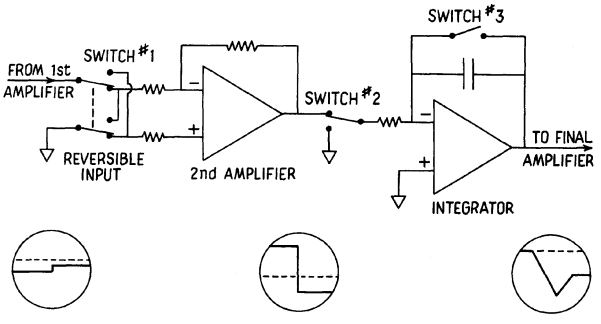

FIG. 5. Schematic diagram of the differential sampler. In the depiction of the waveform at each stage, the dashed line represents zero volts.

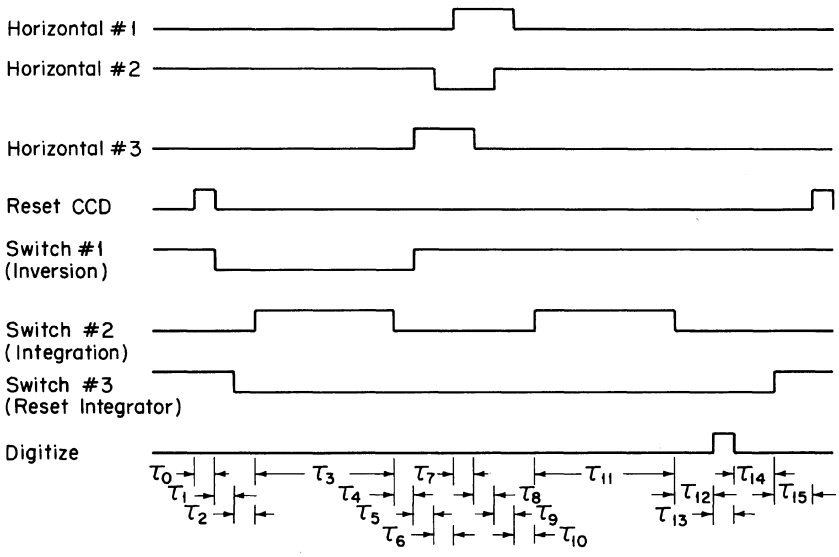

FIG. 6. Timing logic for horizontal transfer and switching.

for down-integration τ_{11} is 125.0 μs. Other timings are discussed in detail in Sec. V, on optimization. With the system balanced, the noise due to the electronics is about five electrons. Compared to the readout noise of 63.9 electrons, the noise due to the electronics makes a negligible contribution to the overall noise since these components add in quadrature.

Because differential sampling is not widely understood, we wish to emphasize the importance of keeping the ratio of integration time to transfer time large when $1/f$ noise is significant, as is the case with CCD on-chip amplifiers. While the contribution to $1/f$ noise is independent of integration time if one neglects the transfer time, a nonzero transfer time results in a signal-to-noise ratio (SNR) proportional to $(1-x)^2$, where x is the ratio of transfer time to total integration plus transfer time (Loh 1977; Hegyi and Burrows 1980). Our horizontal transfer time is 3.4 μs; hence, the SNR is degraded by 2.7% compared to an infinitesimally small transfer time. If the transfer time is 12 μs and the integration time 18 μs for each portion of the differential sample, as in some systems, the degradation is about 44%. It is interesting to compare the SNR obtained with the differential sampling function which we use for our CCD with that of a sampling function optimized for a combination of $1/f$ and white noise (Hegyi and Burrows 1980; Hegyi and Burrows 1984). For small signals, an optimal sampling function could improve the SNR by ~40%. To derive this result, we have estimated that the white-noise power is equal to the $1/f$ noise power at about 100 KHz.

Two problems we did not anticipate during the initial design of the CCD electronics which deserve mention are the generation of light in the on-chip amplifier under certain operating voltages and the drainage of charge during long exposures from several rows of the CCD along the horizontal register. We have since heard that the latter problem has been observed by others. Our solution to both these problems has been to disconnect the on-chip amplifier output and keep the phase reset voltage low during an exposure. A large load resistance for the CCD is only a partial solution to these problems because although it modestly increases the Johnson noise at the input of the first amplifier, it, more importantly, decreases the signal power received from the CCD relative to the background-noise power. A lower CCD temperature also decreases the charge leakage, but the temperature is constrained by factors discussed in Sec. V. By disconnecting the amplifier output we are able to use a moderately low resistance in the source follower, 2000 Ω.

V. OPTIMIZATION OF CCD PERFORMANCE

We spent considerable time and effort optimizing the performance of our CCD. Our CCD exposures are now characterized by excellent resolution, large dynamic range, and an imaging area essentially free of blemishes. The desired characteristics for the performance of a CCD system are summarized in Table I. To achieve this, the parameters in our system listed in Table II were designed to be adjustable.

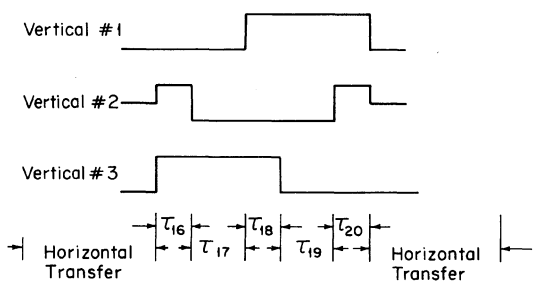

FIG. 7. Timing logic for vertical transfer.

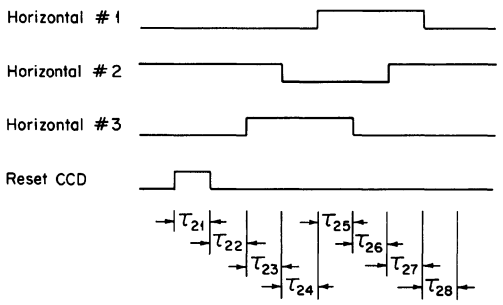

FIG. 8. Timing logic for fast readout.

TABLE I. Desired characteristics of a CCD system.

A.	The CCD
	1. Large dynamic range (low readout noise and deep potential wells)
	2. High resolution
	3. Small number of blemishes
	4. No leakage of charge from wells
	5. Low dark current
	6. Low noise pickup
	7. Time stability
B.	The Electronics
	1. Lowest possible noise
	2. No saturation
	3. Perfect differential sampler
	4. Time stability

TABLE II. Available control parameters.

A.	Voltages
	1. 6 Horizontal (2 states for each of 3 electrodes)
	2. 9 Vertical (3 states for each of 3 electrodes)
	3. Reset drain, output drain, and phase reset
B.	Timing Values—29 Total
C.	Timing and Switching Logic
D.	Temperatures of CCD Chip and Electronics

After the system was constructed, we first adjusted the parameters to obtain a usable picture from the CCD. Next, we began a long series of tests to fine-tune the system. The process of optimization is complicated by the fact that a single control parameter can affect more than one CCD characteristic and, consequently, adjusting one control parameter may necessitate readjusting another. These interactions, however, are not serious, and in most cases the CCD's characteristics are unchanged by a small variation in the value of a parameter.

A large dynamic range is achieved by optimizing the CCD so that the potential wells can be filled to their maximum capacity. With a readout noise of about 64 electrons for the RCA thinned chip, a dynamic range of about 5500 can be attained. Whereas the readout noise is mostly determined by the $1/f$ noise spectrum of the on-chip amplifier and is sensitive to timing ratios in the differential sampling function, the ability of the potential wells to hold charge is strongly determined by the voltages on the horizontal and vertical transfer electrodes.

In order to determine a first approximation to the operating potentials of the electrodes and so obtain a usable picture, we exposed the CCD to light of an intensity and duration such that the wells would be filled to their maximum level. The empirically determined voltage ranges to prevent bleeding of charge are given inTable III. The voltage limits were less stringent at lower exposure levels; hence, adjustment at a high exposure level guarantees a large dynamic range.

Having determined a range of electrode potentials which gave the largest dynamic range, the next step was to determine the potentials whch gave the best resolution. Because the resolution is adversely affected by the number of transfers of a charge packet before it is deposited onto the on-chip capacitor of the CCD and by low light levels (the threshold effect), we achieved the greatest sensitivity for improvement of resolution by imaging a faint point-like source (a pinhole in aluminum foil whose image was larger than one pixel) onto the last area of the CCD to be read out. The resolution was most sensitive to the potential on horizontal electrode #3. By making the voltage swing during transfer larger on this electrode than on those of the other horizontal electrodes, i.e., -7 to $+6$ V compared to -5 to $+5$ V, the resolution was noticeably improved. The potential on horizontal electrode #2 had less of an effect, but when the voltage swing was set to only -4.5 to $+4.5$ V, defects appeared on the screen and the resolution deteriorated.

The resolution is also strongly affected by the time spent transferring charge. In general, a longer transfer time gives higher resolution, but at the expense of increased blemishes. In addition, the vertical timing values and the vertical electrode potentials are weakly coupled. Referring to the timing logic diagram in Fig. 6, we found that the horizontal transfer times τ_5, τ_6, and τ_9 had little or no effect upon either the resolution or the blemishes. Making τ_7 longer (up to 1.2 μs) improved the resolution; however, pulse feedthrough on the chip increased the readout noise. A good compromise for τ_7 was 0.5 μs. Making τ_8 longer also improved the resolution, but at the expense of increased blemishes. Its value was set to 0.3 μs.

For the vertical transfer times (see Fig. 7), increasing τ_{16} and τ_{18} improved the resolution slightly, but also increased the number of vertically oriented blemishes. Increasing τ_{17} and τ_{19} improved the resolution, but also increased the number of blemishes, though τ_{20} had no effect on resolution and blemishes. By making the high state voltage on vertical electrode #3 more negative, the number of blemishes could be reduced, but at the expense of reduced well capacity (observed as bleeding of the image at high exposure levels). The voltage levels at which the number of blemishes were reduced and at which bleeding appeared were found to be affected by the vertical transfer times. The best compromise for electrode potentials is listed in Table IV. The effect of temperature on resolution is discussed at the end of this section.

When the CCD is operated in its focusing mode, a 64×64-pixel area is read out with the same electrode voltages and timing values as for the full frame, but the surrounding pixels are read out at a much faster rate. The time intervals τ_{21} through τ_{28} control this fast readout and the timing logic is

TABLE III. Electrode potentials for which bleeding of charge did not occur.

Electrode number	High state	Low state	Intermediate state
Horizontal 1	$3.5 < V < 6.5$	$-5.9 < V < -1.7$	
Horizontal 2	$-3.0 < V$	$V < -3.0$	
Horizontal 3	$1.7 < V$	$-7.9 < V < -3.4$	
Vertical 1	$-0.8 < V$	$V < -8.0$	$V < -2.4$
Vertical 2	$V < -2.0$	$V < -8.0$	$-4.0 < V < 0.0$
Vertical 3	$-0.6 < V$	$V < -8.0$	$V < -1.0$

Note to TABLE III

Only the vertical electrodes have an intermediate state.

TABLE IV. Best CCD voltages.

Electrode number	High state	Low state	Intermediate state
Horizontal 1	5.0	−5.0	
Horizontal 2	5.0	−5.1	
Horizontal 3	6.0	−7.0	
Vertical 1	0.5	−9.4	−9.0
Vertical 2	−4.0	−9.4	−2.0
Vertical 3	−1.5	−9.4	−9.0
RD = 12.8			
OD = 10.8			
ϕ_R = 8 and −8			

shown in Fig. 8. We found that the readout noise is slightly dependent on τ_{22} and τ_{28} and that values of 4.2 μs give about a 4% reduction in noise compared to 0.3 μs.

We also operate the CCD in a mode called "speckle" which is designed for short exposures of the highest possible resolution. This mode is similar to the focus mode described above, except that the 64×64-pixel area is placed in the upper left of the CCD frame (the first area to be read out), and hence undergoes the smallest number of transfers before being sampled. In this mode, 32 consecutive frames are recorded on tape. We were able to increase τ_5 through τ_9 so that the total horizontal transfer time was 6 μs and to increase τ_{16} through τ_{20} to 1.2 μs without producing blemishes. The improvement in resolution was 27% as determined by measuring the FWHM of a point-like light source imaged onto the CCD. For our full-frame format and the long exposures required by astronomical objects, the resolution is as good as in the speckle mode; this is because the longer exposure times allow a sky background to build up which eliminates the threshold effect.

The threshold effect, which is eliminated with a background of about 300 electrons, was investigated by imaging faint point-like images on an adjustable background. In the upper left corner, however, there is no measurable threshold effect. The fast read τ's, τ_{24}, τ_{25}, and τ_{26} were, in addition, increased to 1.2, 0.7, and 0.6 μs in order to improve the uniformity of the bias exposure (zero incident light, zero exposure time).

Several schemes for the timing and switching logic were tried. Our first arrangement (based on the RCA data sheets) used four fewer time intervals for the horizontal transfer. A shorter total time for transfer was expected to give lower readout noise, as explained in Sec. IV. With this switching logic, the electrodes have different potentials before and after charge transfer and, because of feedthrough effects in the CCD, an offset voltage developed which caused saturation in the amplifier chain. Because the time between integrations is constrained by its effect on resolution and blemishes, the $1/f$ noise and the white noise can be made smaller by lengthening the integration times (the integration times, τ_3 and τ_{11}, also control the gain of the amplifier chain). Other optimizations of the timing included determining the best values to avoid feedthrough and pickup of timing signals by the integrator. In particular, τ_1, τ_2, and τ_{10} should be slightly greater than 1.2 μs. The best values for the timings are listed in Table V.

We found that the chip temperature, in addition to affecting the loss of charge near the horizontal register mentioned in Sec. IV, also controls the quantum efficiency and the resolution. For each degree drop in temperature, the quantum efficiency decreases by about 0.3% (tested at a wavelength of 5500 Å and between 107 and 173 K). The resolution was found to be best at 152 K, with the FWHM of a point-like image being 41% larger at 10 deg higher or lower than the optimal temperature (see Fig. 9). Note that the minimum FWHM in Fig. 9 is constrained by the size of the point-like image we used. Although we can obtain a FWHM of unity, e.g., on spectral emission lines, the actual value depends on the placement of the image relative to a pixel's boundary. Here we avoided that problem by using a larger image. We estimate the accuracy of our temperature measurements to be approximately three degrees.

TABLE V. Timing values.

Tau	Time (μs)
0	0.3
1	1.5
2	1.6
3	124.7
4	0.3
5	0.3
6	0.3
7	0.5
8	0.3
9	0.3
10	1.4
11	125.0
12	0.6
13	0.3
14	13.0
15	28.6
16	0.5
17	1.0
18	0.5
19	1.0
20	0.7
21	0.3
22	4.2
23	0.3
24	0.3
25	0.5
26	0.3
27	0.3
28	4.2

Precise control of the CCD temperature is important not only for assuring that one is always within the optimum range of operation, but also to stabilize the bias level at the input of the amplifier chain. For each degree change in temperature of the CCD, the input bias (equal to the voltage drop across the CCD load resistor) changes by 3.1 mV. As we mentioned in Sec. IV, the output bias level is made relatively insensitive to this change by carefully balancing the differential sampler by adjusting τ_3. We stabilize the tem-

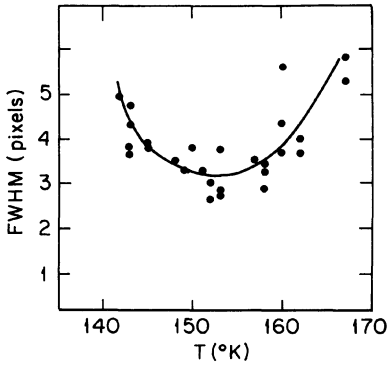

FIG. 9. Dependence of resolution on CCD temperature.

perature of the CCD by a very high-gain feedback loop which drives a heating element, two resistors mounted inside a copper block onto which the CCD is attached. Operation at an average power input of about 1 W gives some additional sensitivity (depending on the square root of the power) and does not seriously decrease the holding time for liquid nitrogen in our Dewar.

VI. DETERMINATION OF CCD NOISE

The readout and spatial noise of a CCD are of importance because the former determines the minimum detectable signal for situations with a small number of collected photons per pixel such as spectroscopy and speckle interferometry, while the latter sets the limit for a large number of collected photons such as the detection of faint features superimposed on a night-sky background. Although spatial noise can be corrected by dividing by a flat-field exposure, the variation in quantum efficiency with wavelength from pixel to pixel is sufficiently large that differences between the spectrum of the signal and that of the flat field become important at roughly the level of a few parts in 1000, for wavelengths near bright night-sky lines.

The measurement of readout noise can be done in two ways. The methods start from the basic equation

$$\sigma_{t,e}^2 = \sigma_{r,e}^2 + \sigma_e^2 + \sigma_{s,e}^2, \tag{1}$$

where $\sigma_{t,e}$ is the total noise, $\sigma_{r,e}$ the readout noise, σ_e the shot noise, and $\sigma_{s,e}$ the spatial noise. All quantities are on a per pixel basis and are measured in units of electrons. If we define the gain g as ADU/n_e, where ADU represents observed analog-to-digital units and n_e is the corresponding number of electrons, and define F_s, the fractional spatial noise, as $\sigma_{s,e}/n_e$, then because $\sigma_e = \sqrt{n_e}$ we find the total noise in ADU to be

$$\sigma_{t,\mathrm{ADU}}^2 = \sigma_{r,\mathrm{ADU}}^2 + g\cdot\mathrm{ADU} + F_s^2\mathrm{ADU}^2. \tag{2}$$

The first method makes use of a least-squares fit of the observed $\sigma_{t,\mathrm{ADU}}^2$ versus ADU data to the above second-order equation. It is not difficult to show that each point must be weighted by

$$\frac{(n_p - 1)n_0}{2\langle\sigma_{t,\mathrm{ADU}}^2\rangle^2}, \tag{3}$$

where n_p is the number of pixels in each exposure and n_0 is the number of repeated observations of each point. If information about the spatial noise is not needed, some computation can be saved by fitting to a first-order equation after removing the spatial noise from the data. There is an advantage to this second method in that any nonlinearities become readily apparent. The spatial noise is removed by averaging many exposures and computing the mean error about the mean value in ADU for each pixel. A variation of the second method and our chosen approach removes the spatial noise by taking only two exposures. It can be shown that the best estimator of the total noise minus the spatial noise, $\sigma_{t\text{-}s,\mathrm{ADU}}^2$, is

$$\sigma_{t\text{-}s,\mathrm{ADU}}^2 = \frac{\sum_i (\mathrm{ADU}_1 - \mathrm{ADU}_2)_i^2}{2n_p}, \tag{4}$$

where the sum is over all pixels n_p in the area. We expose twice on an area of 50×50 pixels. A sample plot from which the gain and readout noise can be obtained is given in Fig. 10.

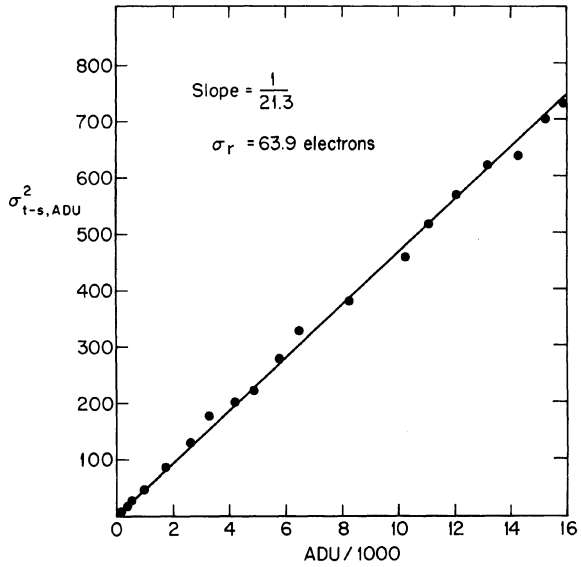

FIG. 10. Determination of CCD gain and readout noise. $\sigma_{t-s,\mathrm{ADU}}$ is the total noise less the spatial noise measured in ADU.

In this case, the reciprocal of the gain $1/g$ is 21.3 electrons/ADU and the readout noise is 63.9 electrons.

As a check to verify that we understand the origin of the readout noise, we can calculate the noise using the measured CCD capacitance and the measured $1/f$ noise of the on-chip amplifier. By measuring the current into the CCD through RD, the voltage across the load resistor as the CCD is being read out, and our cycle time, we obtain a value for the on-chip capacitance of 0.62 pF. To determine the noise spectrum we measured the noise of the CCD output MOSFET from 200 to 25 000 Hz with a calibrated noise-spectrum analyzer. Over this range, the noise power spectrum closely follows a $1/f$ law of the form F_0/ν, where F_0 is 230$(\mu\mathrm{V})^2$. The measured exponent of ν is 1.12 ± 0.05. Estimating, as in Sec. IV, that the white-noise power equals the $1/f$ noise power at 100 kHz, we calculate (Loh 1977; Hegyi and Burrows 1980) the readout noise to be 56.7 ± 6 electrons. The error represents an estimate of our overall measuring accuracy. The measured and calculated readout noise are in reasonable agreement.

By measuring the readout noise as a function of pixel location on the CCD, we find that the readout noise is position dependent. Figure 10 shows data taken in the upper-left corner of the chip, but in other locations the shape of the curve is no longer straight. For these areas, the measured total noise at intermediate level exposures (up to 200 000 electrons) is a few percent higher than the shot noise. This is most likely due to the greater contribution of transfer noise and trapping noise, both of which depend on the square root of the total number of transfers (Beynon and Lamb 1980). At higher levels (up to 350 000 electrons, a limit set by our ADC), we observe the same noise at all locations. Also, when no light is incident on the CCD, the noise is essentially the same everywhere.

A measure of the spatial noise as a function of wavelength and sample size was obtained by computing the standard deviation of a square area of a twilight-sky flat-field exposure as a function of the length of an edge of the square after removing the contribution due to readout and shot noise.

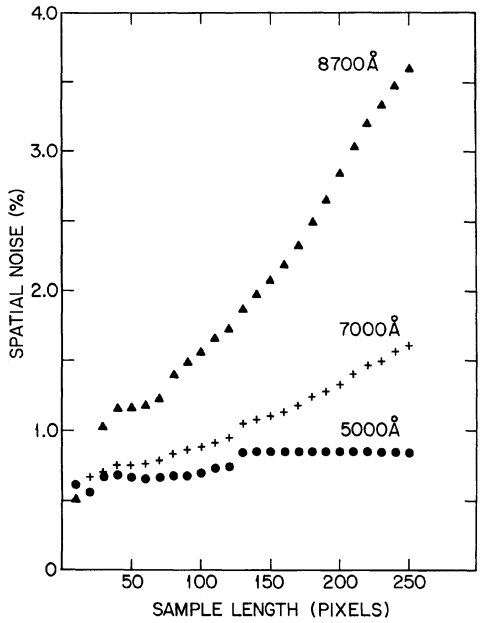

FIG. 11. CCD spatial noise as a function of wavelength and sample size.

From the plot of this data shown in Fig. 11, we see that the corrections for spatial noise become increasingly important at longer wavelengths and larger pixel sample sizes. Flat-field exposures taken of the dawn and twilight skies on different nights are consistent with each other on the large scale (the upper third compared with the lower third of the frame) from one part in 100 000 to a few parts in 10 000 peak to peak.

A complete correction of an exposure for spatial noise and zero level is carried out on a pixel-by-pixel basis using the expression

$$\frac{[(\text{exposure} - \text{bias}_E) - (\text{dark frame} - \text{bias}_D)]\text{scale}}{(\text{flat-field} - \text{bias}_F)}, \quad (5)$$

where the dark frames and flat fields are averages of many exposures and the scale is a normalization constant which could be set equal to the average intensity of a pixel for the average of the flat-field exposures. Each processed exposure must have a constant bias value for the exposure, dark frame, and flat field (bias_E, bias_E, and bias_F, respectively) subtracted; in our system, these are computed from an average of a 50×50-pixel area which is read out immediately following readout of the full frame. Dark frames for exposure times from zero to more than 45 min show structure on the large scale of about 40 electrons peak to peak. Because this can be as much as 1% of the sky level, it is important to subtract this structure from an exposure. At our operating temperature, the dark current contributes about 25 electrons per hour per pixel on average, with small pixel-to-pixel variations.

CCD's are quite sensitive to background low-frequency electromagnetic radiation. We found that a small ripple voltage on the shutter power supply filled pixels with charge at a rate several hundred times larger than the normal dark current. After eliminating the ripple so that the shutter voltage was purely dc, even under load, this problem disappeared.

VII. PICTURE PROCESSING

Our picture processing is done on a VAX/780 computer running under VMS. Images are displayed on a Grinnell GMR270. The Grinnell with joystick-controlled quad cursors and zoom and pan enable us to display images in 4096 pseudocolors. Two VT100 terminals with graphics, a Printronix, a Tektronix 4662 flat-bed plotter, and the Grinnel are each capable of displaying plots.

We have written most of our present software for the reduction of two-dimensional surface photometry of galaxies. The most notable subset of these routines can do a weighted least-squares fit of any chosen surface-brightness formula to an elliptical galaxy's light distribution, while taking into account overlapping stars and galaxies, including seeing effects, yielding the fitting parameters of the formula with their errors. We are currently writing programs to handle two-dimensional spectra. Some programs which were written at Mount Stromlo Observatory, Australia by Greg Quinn, Reet Valek, and Dennis Warne do fast data transfer to and from disk in a format called Standard Astronomical Data (SAD). A 320×512 array of REAL*4 data can be transfered in about 5 s. Copies of the Mount Stromlo programs were transmitted to us by Susan Simpkin. We have given the acronym MIIPS (Michigan Image Processing System) to our picture-processing software.

An example of the quality of our CCD images is shown in Fig. 12. It is a 30-min exposure of the field surrounding NGC 4874 in the Coma cluster taken through a 1000-Å bandpass filter centered at 5000 Å. This exposure, for which the entire frame is shown, has not been flat-fielded or edited in any way, yet there is no evidence of fringing or of charge trailing due to imperfect charge-transfer efficiency. There is one minor blemish in two adjacent columns extending over 34 rows. Although the bright star has saturated the ADC (more than about 350 000 electrons/pixel) in 43 pixels, the star image appears unbroadened with no bleeding of charge, and even diffraction spikes are visible.

In other exposures, we have found fringing at less than 2% of the sky level in the I band, centered at 8700 Å, but at 5500 Å, only one exposure out of 50 has shown fringing. It was observed to be at a level of less than 1% of the sky level.

We have tested the linearity of the system in three ways. By exposing on a constant uniform source of light and then plotting the observed signal in ADU against exposure time, we see no departure from linearity from our shortest exposures (less than 10 ms) to the highest level our ADC can record (16 383 ADU or 348 958 electrons). By measuring the total light in the images of calibration stars observed at a variety of exposure times and for a variety of magnitudes, we obtain a calibration curve with a standard deviation for a single observation of $0^m.0075$ (0.75%). It is interesting to note that we can calibrate on stars as bright as Vega. The exposure times must be kept below 1 s and the images defocused.

A more stringent test of linearity is defined on the scale of a pixel. That is, if at high levels of illumination, charge drains to neighboring pixels but the total charge is conserved, the above two tests will indicate linearity although it is violated on a per pixel basis. By measuring the FWHM of point-like images as the illumination was increased from very low levels up to saturation of our ADC, we observed no increase in the FWHM of the image.

As an example of image processing with our system, Fig. 13 shows the processed field of Fig. 12 with the brightest galaxy modeled by our programs and subtracted from the

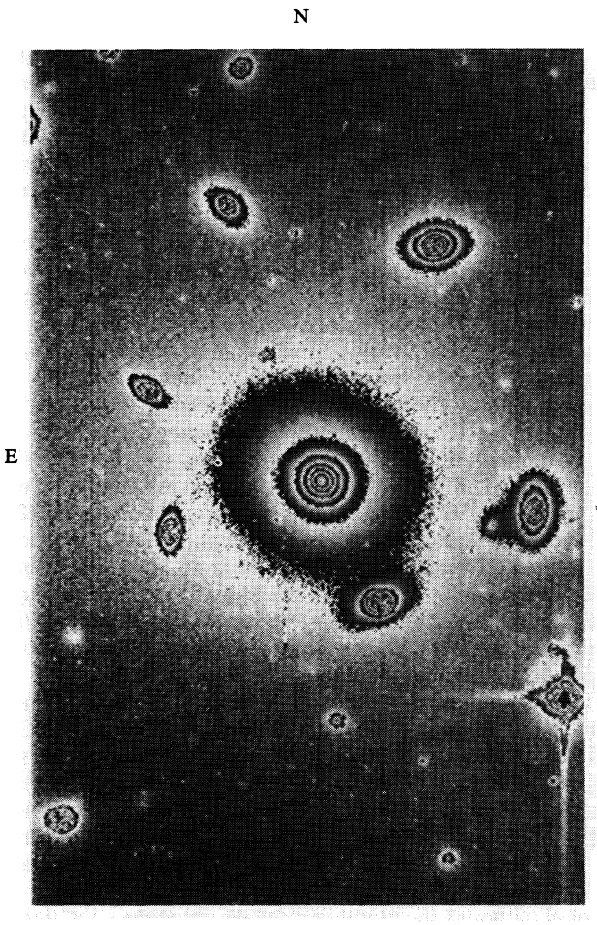

FIG. 12. Thirty minute CCD exposure of NGC 4874 (Coma cluster) at 5000 Å. No flat-fielding has been done.

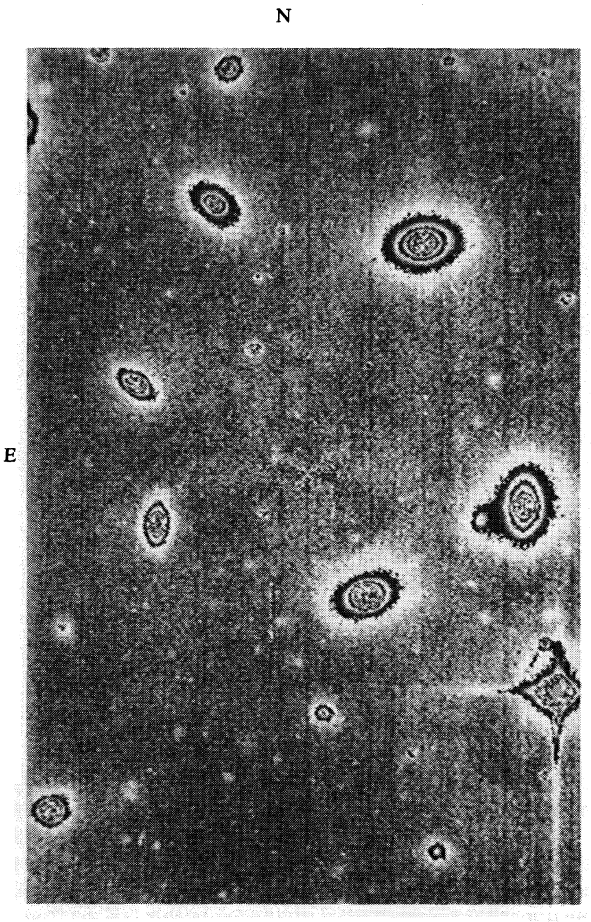

FIG. 13. Same as Fig. 11, but with NGC 4874 modeled and subtracted.

frame. The contrast of the Grinnel was increased slightly for Fig. 13 as compared to Fig. 12. For this figure, we use a logarithmic look-up table which equates a 1% change in intensity with the smallest visible change in grayscale level. Except for the central three pixels, the mean surface brightness of a circular contour has a maximum deviation from the sky level of only one part in 2300 when compared to the central surface brightness of NGC 4874. This deviation is typical of a blank region of sky. The faintest objects visible in Fig. 13 are about $23^{m}6$ (based on $m_{5000} = 0.0$ for Vega). A similar modeling and subtraction for an exposure of M87 reveals globular clusters to the very center of the galaxy.

We would like to acknowledge the help of students Tim Frank, John Haggis, Jim Haywood, Stephen Russell, Slava Sharkin, and Doran Smith, and the support of Air Force Grant No. AFOSR-80-0095.

REFERENCES

Amelio, G. F., Tompsett, M. F., and Smith, G. E. (1970). Bell Syst. Tech. J. **49**, 593.

Beynon, J. D. E., and Lamb, D. R., editors (1980). *Charge Coupled Devices and Their Applications* (McGraw Hill, London and New York).

Dewey, D., and Ricker, G. R. (1980). In *Proceedings of the SPIE Conference on Applications of Digital Image Processing to Astronomy, Pasadena, Cal.*, edited by D. Elliott (The International Society for Optical Engineering, Bellingham, Wash.), Vol. 264, pgs. 42–47.

Hegyi, D. J., and Burrows, A. (1980). Astron. J. **85**, 1421.

Hegyi, D. J., and Burrows, A. (1984). Astron. J. **89**, 1628(E).

Leach, R. W., Schild, R. E., Gursky, H., Madejski, G. M., Schwartz, D. A., and Weekes, T. C. (1980). Publ. Astron. Soc. Pac. **92**, 233.

Loh, E. D. (1977). Ph.D. thesis, Princeton University.

Loh, E. D. (1981). In *Proceedings of the SPIE Conference on Solid State Imagers for Astronomy, Cambridge, Mass.*, edited by J. C. Geary and D. W. Latham (The International Society for Optical Engineering, Bellingham, Wash.), Vol. 290, pgs. 150–154.

Marcus, S., Nelson, R., and Lynds, R. (1979). In *Proceedings of the SPIE Conference on Instrumentation in Astronomy III, Tuscon, Ariz.*, edited by D. L. Crawford (The International Society for Optical Engineering, Bellingham, Wash.), Vol. 172, pgs. 207–231.

Meyer, S. S., and Ricker, G. R. (1980). In *Proceedings of the SPIE Conference on Applications of Digital Image Processing to Astronomy, Pasadena, Cal.*, edited by D. Elliott (The International Society for Optical Engineering, Bellingham, Wash.), Vol. 264, pgs. 38–41.

Robinson, L. B. (1981). In *Proceedings of the SPIE Conference on Solid State Imagers for Astronomy, Cambridge, Mass.*, edited by J. C.Geary and D. W. Latham (The International Society for Optical Engineering, Bellingham, Wash.), Vol. 290, pgs. 124–129.

Wright, J. F., and Mackay, C. D. (1981). In *Proceedings of the SPIE Conference on Solid State Imagers for Astronomy, Cambridge, Mass.*, edited by J. C. Geary and D. W. Latham (The International Society for Optical Engineering, Bellingham, Wash.), Vol. 290, pgs. 160–164.