Chapter 3 = 10148 Words

1. High Resolution Imaging

2. Infrared Detector Technology

3. Camera

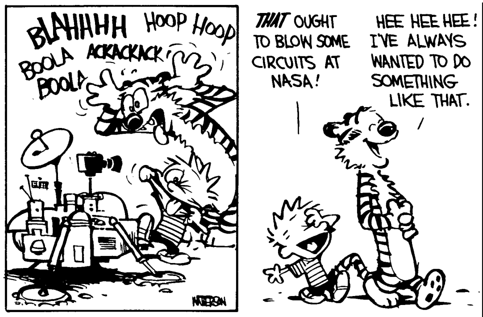

Outline

This chapter describes the near infrared camera built for COAST. It examines the properties required from the camera system, the range of devices available and the factors involved in the choice of detector. The design, construction and performance of the camera system is covered in detail.

3.1 Choice of detector

The choice of detector for the COAST camera depends on the wavelength range over which the system will operate and the format of the image that it must record.

3.1.1 Wavelength Range

As discussed in chapter one, the signal received by a COAST type interferometer increases strongly with wavelength, although for the same baseline the resolution decreases. This implies that the system should operate at the longest possible wavelength. In practice, the longest usable infrared wavelength is limited by the thermal background emission from the telescope. The analysis in chapter one assumed that the size of the telescope mirrors could be increased at longer wavelengths. In reality the fixed small size of the individual telescopes will prevent the system reaching its ultimate theoretical performance by mid-infrared wavelengths. In addition, the telescope is constrained to work in the atmospheric windows at 1.3μm (J band), 1.6μm (H band), 2.2μm (K band), 3.6μm (L band) and 4.8μm (M band).

3.1.2 Detector Format

At the time of choosing the detector for the infrared camera, the form of the COAST beam combiner had not been decided and so both single element and array detectors were considered. The visible system at COAST uses silicon Avalanche Photo-Diode (APD) detectors, Nightingale(1991). These are single element photon counting detectors with large internal gain and high quantum efficiency. Although these silicon devices are only sensitive to light with a wavelength, $\lambda < 1$ µm, it is possible to obtain APDs based on indium gallium arsenide (InGaAs) which are sensitive up to 1.6 µm. Unfortunately, the small band gap needed to detect low energy infrared photons means that a much larger dark signal is produced. This, together with the higher flux expected from many astronomical infrared sources, means that a much higher bandwidth would be needed than is possible with available devices, RCA(1988).

At near infrared wavelengths, array detectors have a large advantage over single element devices. Unless the single element detector contains some internal gain mechanism, as is the case in APDs, then the extra capacitance between the detector junction and the external circuitry needed to detect the signal generally leads to a higher noise level. With large scale array devices manufactured on an integrated circuit, the photon generated signal can be integrated directly on the detector. Since amplifiers and signal processing electronics are then much closer to the detector, the noise is correspondingly lower. This means that even if the beam combiner design could be achieved with single element detectors, because of these advantages it may be better to use sections of a large array detector.

3.1.3 Detector Material

The processes of infrared photon detection were outlined in the previous chapter. Of the materials described in Figure 2.1, the most commonly used in astronomical detectors are platinum silicide, indium antimonide and mercury cadmium telluride. Each of these has their own characteristic advantages and disadvantages which are discussed below.

3.1.3.1 Platinum Silicide (PtSi)

These devices are manufactured from silicon with a surface implant of platinum to form a Schottky barrier diode. Being based on silicon, manufacturing is simple and large format devices with good cosmetic quality and low noise are available. Unfortunately these detectors have a quantum efficiency of only a few percent and are rarely used for astronomical applications now that materials with better light gathering ability have become available. They are however still common in thermal imaging video systems where the high uniformity avoids the need for on-line image correction. Weckler(1987).

3.1.3.2 Indium Antimonide (InSb)

This is the main detector material for cameras in the mid-infrared region over the wavelength range λ = 1 - 5 µm. It has a high quantum efficiency of 80% and large format arrays are available, up to 1024 x 1024 at the time of writing. Fowler(1994).

3.1.3.3 Mercury Cadmium Telluride (HgCdTe - MCT)

This material is especially interesting because the bandgap, and so the long wavelength cut off, can be tuned by altering the relative concentration of mercury and cadmium, Rode(1987). Most of the devices intended for astronomy are designed to be sensitive only to the J, H and K bands, λ = 1 - 2.5µm. This avoids the thermal background and so is especially useful for telescopes which are not optimised for infrared operation.

3.2 Available Devices

Interferometers such as COAST generally operate in a regime where images from small telescopes, in a narrow bandwidth, have to be read quickly enough to freeze the atmosphere. Therefore high efficiency, low read noise and high speed operation are of major importance.

In the years since COAST was originally designed, the performance of infrared arrays has increased enormously and a wide range of devices is now available. The aim of this project was not only to build an infrared camera for COAST but to produce a general purpose instrument which could be used on other telescopes for both high resolution and conventional imaging. This camera would also form the prototype for future infrared astronomical instruments. For these applications an array based camera would be more useful.

Having decided on a large format hybrid array, a range of devices from a number of manufacturers were considered. The main candidates were 256 x 256 pixel arrays from Hughes Santa Barbara Research Centre and Rockwell International Science Centre.

3.2.1 SBRC

The Santa Barbara Research Centre (SBRC) device is an InSb array available in 256 x 256 pixel format. It is the successor to the 58 x 62 pixel array used extensively in astronomical instruments such as IRCAM, McLean(1987). The principle advantage of the InSb based device is its high quantum efficiency, reported as greater than 80%. The indium antimonide detector is sensitive to a wavelength range, $\lambda = 1 - 5.5\mu m$, and so can be used in all the near and mid infrared bands. However if the system is to operate only in the near infrared, then this response to longer wavelengths is a disadvantage as the camera will be more sensitive to the thermal background. Operating at these longer wavelengths also means that the device requires cooling to liquid helium temperatures and careful design of dewar and optics to shield the detector from the instrument background thermal radiation.

3.2.2 NICMOS

The NICMOS device built by Rockwell Internal Science Centre is also a 256 x 256 pixel hybrid array but designed for low background astronomical applications, particularly in space. The term NICMOS is from Near Infrared Camera Multi-Object Spectrograph, a proposed replacement instrument for the Hubble Space Telescope.

The material used in the NICMOS array is $Hg_{0.554}Cd_{0.446}Te$, with an infrared response from $\lambda = 1 - 2.4\mu m$ and so is limited to the J, H and K windows. This is ideal for our application since COAST is unlikely to operate at longer wavelengths due to the high thermal background signal and reduced resolution. The dark signal at 77K is reported to be less than 10 electrons/second/pixel, Vural(1990). This allows operation at liquid nitrogen temperatures which would considerably simplify the dewar design. Although the quantum efficiency is lower than that of InSb, 80% rather than 50%, the shorter cut off wavelength means that less thermal noise signal from the sky and telescope is detected.

Both these devices are similar in structure and would require comparable support hardware and software to operate. Read noise figures of 30 - 50 electrons were reported for these arrays by the manufacturers and a number of astronomical groups, Rieke(1993b).
The NICMOS array was chosen largely because of the higher operating temperature and reduced sensitivity to thermal infrared which simplifies the dewar design.

3.3 NICMOS Array

This section will describe the construction and operation of each of the components of the NICMOS array. The camera system and its operation are described in subsequent sections. The name NICMOS3 is generally used for the 3rd generation 256 x 256 device. Throughout this work the term NICMOS will be used to refer to the device used at COAST.

The NICMOS detector is an array of 256 x 256 pixels, each 40µm square, whose elements are arranged as four independent quadrants of 128 x 128 pixels. The division into quadrants is purely a function of the readout multiplexor, the detector itself is a single piece of semiconductor and there is no dead space between quadrants.
In a hybrid focal plane array each pixel must be regarded as a separate detector, with individual control and readout circuitry associated with it. This circuitry which is repeated in each pixel is called the Unit Cell. A multiplexor contains the matrix of switches which connects the output of each pixel to the output amplifier and so to the external signal processing system. The structure and function of each of these components is described below.

3.3.1 Detector layer

The NICMOS detector material is produced by a proprietary Rockwell process called PACE (Producable Alternative to Cadmium Epitaxy), Kozlowski(1990). The detector structure is built on a wafer of synthetic sapphire, which is sliced and polished to a thickness of 200µm to produce a stable substrate on which to grow the infrared detector material. A layer of cadmium telluride (CdTe) is chemically deposited on the surface and then the molten mercury cadmium telluride is poured over this to form a thin uniform layer. After the infrared sensitive layer is deposited, individual detector diodes are formed by an ion implant and then insulating barriers produced to define the pixel boundaries. The wafer is then diced into array sized squares and indium columns, called bump-bonds, are deposited on each pixel. These will form the electrical connection to the readout circuit.

The rest of the device containing the amplifiers and multiplexor is built from conventional CMOS silicon processes on a single integrated circuit. The final hybrid device is produced by depositing indium columns on each pixel in the readout circuit and pressing the sapphire wafer onto it. The indium cold-welds to form an electrical connection and the gap is back-filled with epoxy to provide a mechanically stable structure. A cross section diagram of part of the array is shown Figure 3.1, note that this is not to scale.

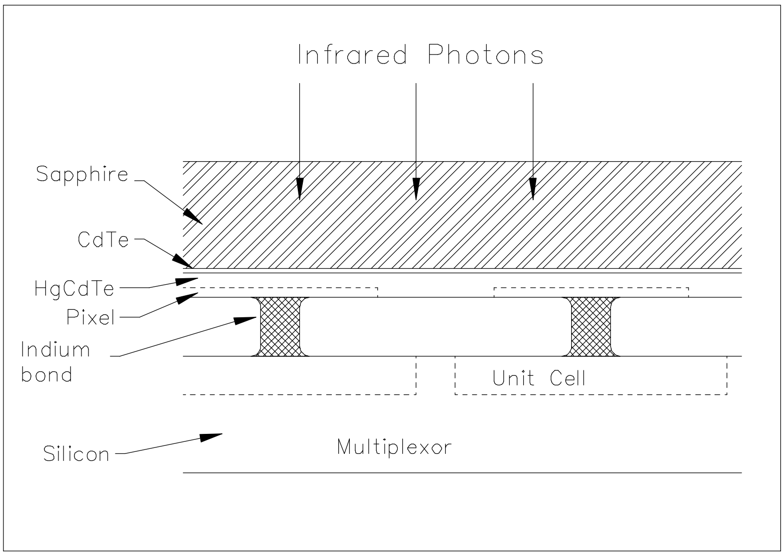

Figure 3.1: Cross section of NICMOS array

3.3.2 Unit Cell

The schematic diagram Figure 3.2 shows the components of the unit cell and the equivalent circuit in a simplified form. The signal arriving at each pixel is stored in the pixel until it is read at the end of an exposure. The signal charge is integrated on a capacitor formed from a diode junction. At the start of an exposure this junction is charged to a fixed voltage by momentarily closing the reset switch. During the integration the photons arriving in the detector layer are absorbed and produce free electrons. The current generated flows down the indium bond and onto the junction where it discharges the capacitor. At the end of the exposure the voltage remaining on the capacitor is measured by an insulated gate FET contained in the pixel circuit. To read a pixel, the multiplexor closes the select switch which connects the output of this FET to the output bus and so to the output amplifier, where the signal can be measured. Since the CMOS FET has a very high gate resistance, the signal charge on the capacitor does not leak away and so the process of reading a pixel can be performed repeatedly without affecting the stored signal. This is called non-destructive readout.

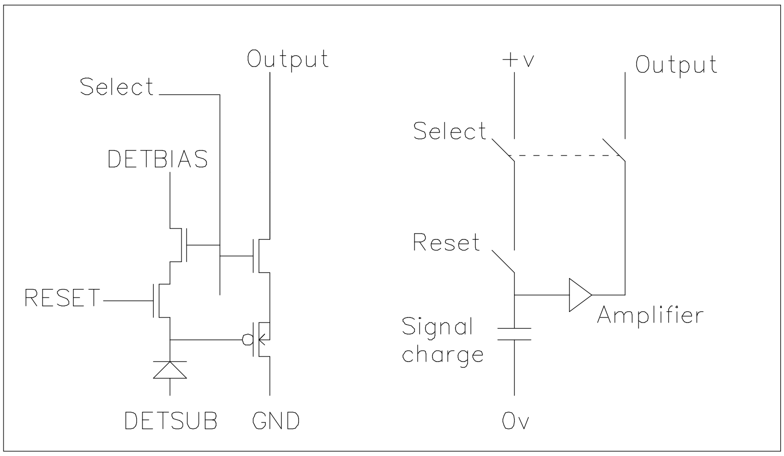

Figure 3.2: Unit cell and equivalent circuit

3.3.3 Multiplexor

The multiplexor circuit contains the array of switches which connect each pixel in turn to the output. A diagram of the multiplexor circuit is shown below in Figure 3.3, this is a simplified schematic and does not show all the power supply connections. The NICMOS array is organised as 128 rows running horizontally and 128 vertical columns. As described above, the output of each pixel is connected to a bus by a switch. To read a pixel the corresponding row and column select switches are closed to connect that pixel to the column bus and in turn connect that column to the output. The reset switch is also connected in series with this select switch and so the reset operation only applies to the currently selected pixel.

The address of the pixel to be accessed is selected by setting a bit in two shift registers. The PIXEL and LINE clocks increment the value in the shift registers, advancing them to the next column or row and so scan through the array. There is also an asynchronous clear function which resets both shift registers. The detailed operation and timing of this system is described in section 3.5.

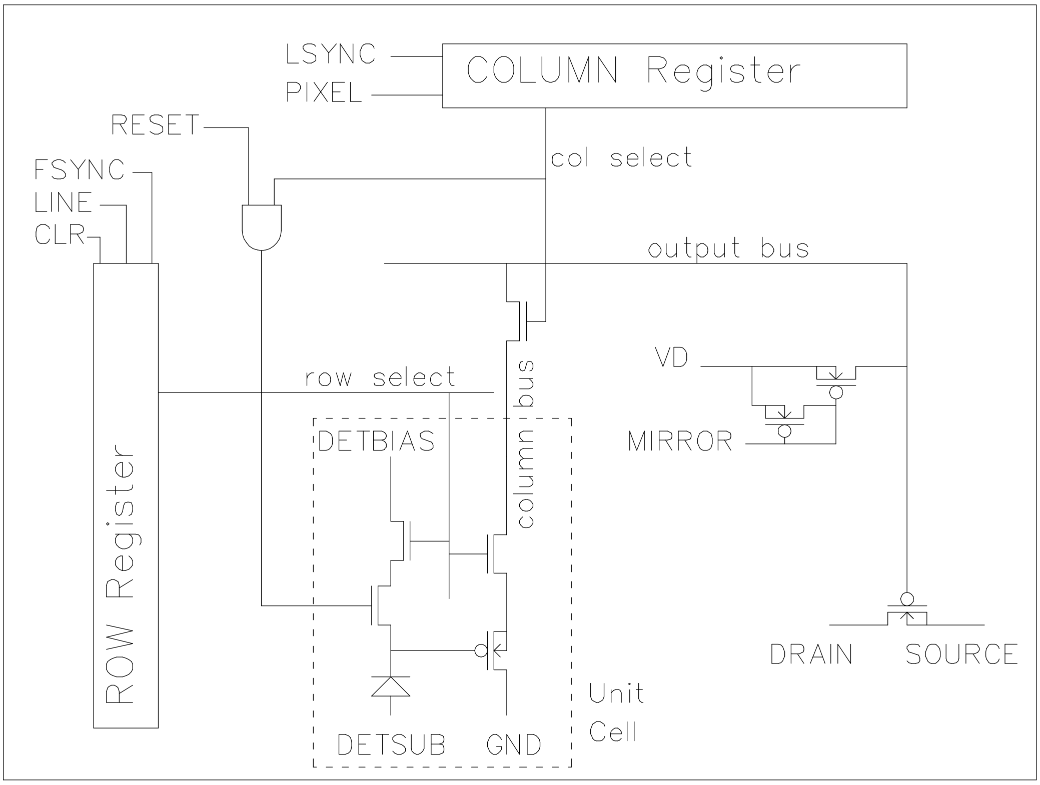

Figure 3.3: Multiplexor schematic

3.3.4 Output amplifier

Each pixel on the array has an individual amplifier which buffers the signal integrated on the detector junction, this allows the charge to be read without destroying the signal. When a pixel is selected the output of this transistor is connected to the output amplifier, which is a single P-channel FET. The VD and MIRROR lines provide a constant current bias for the output bus.

3.4 Hardware

The camera developed for COAST was designed not just to operate with the interferometer but to be a general purpose instrument. Time and budget constraints, together with a lack of previous infrared camera experience, mean that the design was based on visible CCD cameras using an existing CCD controller and a simple dewar. The design and construction of each of the main components of the camera system are described below.

3.4.1 Dewar - Cryogenics

The camera is housed in an Oxford instruments dewar normally used for CCD cameras. This is a single stage liquid nitrogen cooled system with a downward looking detector. To simplify the design there are no optics inside the dewar other than a single cooled filter which can only be changed by partially dismantling the camera. This eliminates the need for any vacuum feed-through to drive the mechanism. A filter wheel was designed, but all the schemes for the COAST beam combiner would need some realignment to work at different wavelengths and so quick changes of filter were not considered essential. However, this arrangement would limit observations on a conventional telescope to a single infrared band for the entire night.

The diagram in Figure 3.4 shows a cross-section of the dewar used for the COAST camera, the right hand side of the diagram has been exploded to show the components more clearly. Since the dewar is designed for visible CCDs with simple cooling requirements there is limited space available. The components in the dewar must fit inside a volume 130 mm in diameter and 75 mm high.

The NICMOS detector should ideally operate at a temperature of 60 K with a dark current of around 1 e^-/s/pixel. At the temperature of liquid nitrogen (77 K) this rises to around 10 e^-/s/pixel, Vural(1990). The aim of the dewar design is to cool the detector to as close to 77K as possible. There is no active temperature control since this would increase the minimum operating temperature, instead the design relies on the thermal inertia of large masses of metal to maintain a constant temperature. Tests show that the temperature of the components is stable to within the 0.1K resolution of the sensors over periods of a few hours. After filling, the dewar takes 30 minutes to reach a stable temperature and then has a hold time of around 24 hours.

The thermal contact between two surfaces inside a dewar is generally poor since the layer of grease and oil which normally coats metal surfaces is removed by the vacuum and the two surfaces, however well machined, only actually touch at a few raised points. The heat conductive pastes normally used in electronics cannot be used at cryogenic temperatures. A common solution is to place a thin sheet of indium foil in the joint. This soft metal fills the surface irregularities and improves the heat transfer. Most of the components inside the dewar are made from aluminium alloy, which has a poorer thermal conductivity than copper but is considerably easier to work with.

The camera components are built on a copper cold-face which forms the front part of the liquid nitrogen reservoir. This is held in position by a fibre-glass spider which keeps this surface fixed relative to the dewar while it is being cooled. The position of the detector is largely determined by the optical system used to feed the image from the telescope. A *fast* beam (high f number) would mean that the chip was close to the front face of the dewar and make efficient cooling difficult. A *slow* beam (low f number) allows the focal plane, and hence the detector, to be further back in the dewar but it is difficult to produce a wide field of view. If individual pixels of the array are to be used as single element detectors then diffraction effects will also limit the maximum focal ratio of the beam. The final camera design for COAST allows an f/6 beam of 25 mm diameter, although there is some vignetting.

To cool the detector efficiently the thermal conduction between the device and the cryogen reservoir must be maximised and the heat load into the detector minimised. The main cooling effect comes from a metal cold-finger which is in contact with both the cold face and the NICMOS chip. The design of the chip package itself is not ideal. The electrical contacts to the detector are on the top, sides and bottom. Only an area 16 mm square, about half the surface area, on the back of the package is available. The cold finger on the mount block has a large surface area in contact with the cold face and cools the back of the detector package. This is machined from a single piece of metal to improve the conductivity. The cylindrical mount also forms the side walls visible in the diagram. These increase the shielding of the cooled components from the warm outer surface of the dewar.

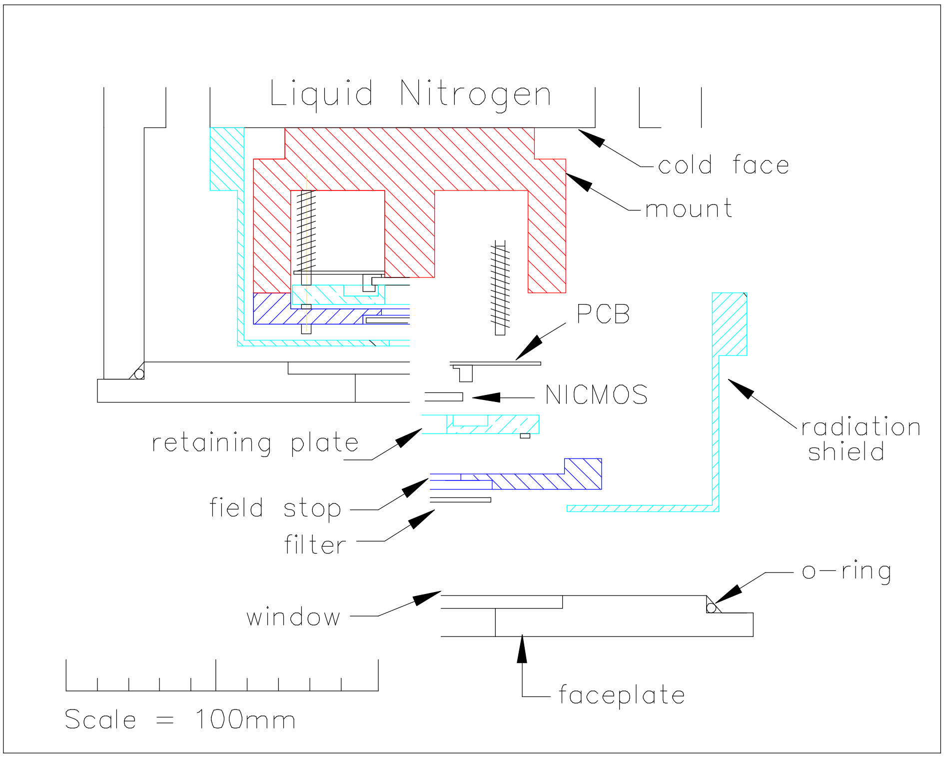

Figure 3.4: Dewar cross-section

The detector chip fits in a PLCC socket on a printed circuit board. A hole cut through the PCB and chip socket allows the cold finger to make contact with the back of the device. The circuit board connects the NICMOS package contacts to the dewar electrical connector and bridges the control lines to the four quadrants as described in section 3.4.3. The electrical connections to the outside of the dewar are made with low thermal conductivity, Be:Cu wire to reduce the heat load into the device.

Immediately in front of the NICMOS package is a retaining plate which is sprung to press the device onto the cold finger. The thermal conductivity of a junction between two components is strongly dependant on the force between them. The PCB is fitted with springs which push the socket against the device to maintain a good electrical contact. As the camera is cooled from room temperature to 77K, the metallic components contract. All the materials used are chosen so that this effect does not increase the mechanical force on the detector package. The retaining plate also provides a second source of cooling through copper straps attached to the cold-face.

A circular cover plate attaches to the top of the mount block, completely surrounding the detector except for the field stop which is described below. This plate also holds the single waveband selecting filter. The detector is sensitive to a longest wavelength of 2.5μm, at which there is a large signal from a room temperature source. Since the filter is opaque at wavelengths outside its pass-band, it must also emit at those wavelengths. This implies that the filter must be cooled.

The major heat load into the cooled device is radiation from the walls of the dewar which are at room temperature. To minimise this, a radiation shield completely surrounds the internal components, except for a 13 mm diameter hole to let the image in! This shield and the inside of the dewar are then covered in multiple layers of highly reflective aluminised Mylar to lower their emissivity.

The dewar of an infrared camera does not simply insulate the cold detector. At near infrared wavelengths the thermal signal emitted from a room temperature object becomes significant. The dewar must be designed so the field of view of the detector is restricted to only the cold sky seen by the telescope and the cold components inside the dewar. To work efficiently most infrared dewars operate with a Lyot stop. This is a small hole in a cold cover placed at the focal plane which allows only light from the astronomical target to reach the detector and blocks any thermal emission from the telescope structure. A cold metal shield would completely surround the detector, except for this hole. To re-image this focal plane on the detector then requires further cold optics inside the dewar.

The final design of the COAST beam combiner means that the camera operates as if it were looking at four separate telescopes simultaneously and so the design of cold stop and internal optics is more complicated than for a conventional astronomical camera. A compromise was reached which uses an oversized cold stop in the converging beams. This allows a simple dewar design at the expense of a higher background signal. This is not a problem for the short exposure used in interferometry but limits the usefulness of the camera for faint object astronomical imaging.

3.4.2 Controller

The array controller generates the sequence of clock signals to operate the NICMOS array. It also measures the output signal from each pixel and provides a number of stable power supplies to drive the array.

The COAST system uses an Astromed 3200 CCD controller which consists of a RISC microprocessor on an interface card contained in the host computer together with a separate Digital Drive Electronics (DDE) unit which contains the electronics to read the array. These components are connected by a bi-directional RS-422 bus. The RISC controller runs a sequence program which operates all the functions of the NICMOS array as well as the amplifiers and Analogue to Digital Converter (ADC) in the signal processing chain. In addition to this, the controller also provides a number of stable programmable power supplies which are used to supply the bias levels on the NICMOS device. The flexibility of having the array control signals generated by a microprocessor controlled sequencer rather than hard-wired logic allows a wide range of readout patterns, including sub arrays and even reading individual pixels. The design of the sequencer code is described in section 3.5 .

3.4.3 Interface

The CCD controller was not modified to operate the NCMOS array. Instead an interface was built which is housed in a small box attached directly to the dewar. This unit allows any of the four quadrants to be connected to the controller, adapts the output of the NICMOS array to appear more like a CCD and contains circuits to protect the NICMOS device from static shock or incorrectly programmed supply voltages.

The NICMOS device can be regarded as four independent 128 x 128 arrays, each with their own separate control and readout electronics. The controller is only able to measure a single data channel and so the interface can connect any of the quadrants to the controller. The control signals and bias supplies to all quadrants are wired in parallel in order to reduce the number of connections inside the dewar. Only the output signals from each quadrant are kept separate. The signal from each quadrant is connected through a relay to the controller which allows each quadrant to be read in turn. Relays were chosen since they give very high isolation and are noise free. The switching time of the relay (≈1ms) is negligible since it occurs only once at the end of each quadrant read.

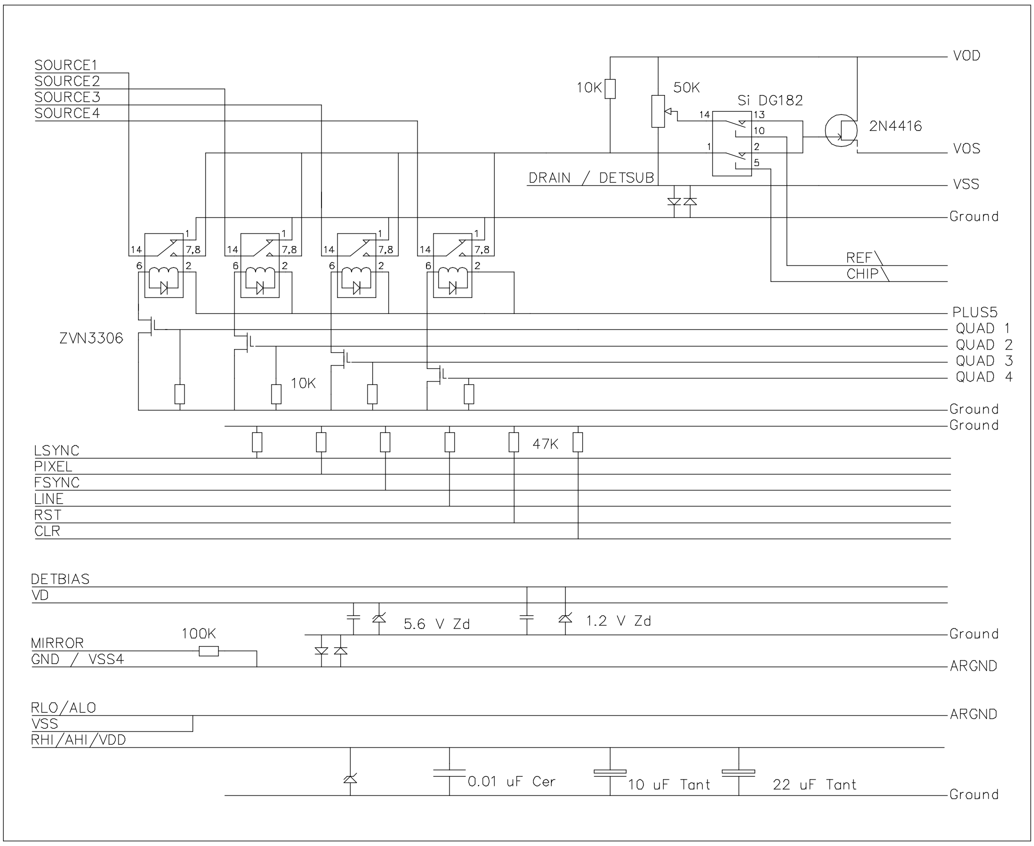

Figure 3.5: NICMOS-Astromed Interface schematic

The Astromed CCD controller provides all the bias voltage levels for the NICMOS array from a series of programmable power supplies which can be set under software control. To prevent an incorrectly set bias level from damaging the array, all the lines are connected to a safety earth through back-to-back diodes. This arrangement allows the analogue and digital levels to remain independent but prevents them moving out of a -0.6 to 0.6 volt range. The 5v bias supplies are similarly connected through 5.6v Zener diodes. This system also protects the CMOS circuitry in the multiplexor from electro-static shocks. Care was taken to avoid earth loops, the connections marked ground are used for safety and isolation only. The ARGND is analogue reference ground and refers to separate connections to a single star point earth in the controller.

The NICMOS output amplifier is constructed from a P channel FET. Most silicon CCDs use an N channel device and the Astromed controller is designed for this arrangement. An extra N channel FET (2N4418) in the interface adapts the NICMOS output to be compatible with the controller. The silicon switch DG182 is used to allow differential readings of the NICMOS output and a reference level, this is described in section 3.5.4 .

3.4.4 Control Computer

The RISC processor is housed in an IBM-PC and connected to a normal expansion slot. The host computer downloads a sequence program into the RISC system which then executes it independently of the host machine. During the readout of the array the image data is read back into a buffer memory on the RISC card. When the program ends the data in the buffer is available to be read. The pixel values are then transferred to the host computer in single byte units by programmed I/O. The controlling software for COAST consists of a simple interactive MS-DOS based graphical program for use by astronomers at the telescope and a more flexible UNIX based programmable system for engineering design and testing.

3.5 Operation

The hardware and software support needed to operate the NICMOS device is very simple. Instead of the complex waveforms needed to transfer charge between pixels in a CCD, only TTL digital logic levels are needed to select the array functions. All the functions of accessing and reading a pixel are controlled through six input lines to the multiplexor. There are separate lines for each quadrant, but in this camera they are connected in parallel to reduce the number of connections into the dewar.

The clock signals and their functions are:

CLR	Resets both shift registers, deselects all pixels
FSYNC	Data bit for row select, select first row
LINE	Clock for row select, negative edge triggered

LSYNC	Data bit for column select, selects first pixel in row
PIXEL	Clock for column select, advances on negative and positive edge
RESET	Clears signal charge in currently selected pixel

The instructions to read the array fall naturally into two parts. Firstly, the sequence of instructions to the multiplexor to select a particular pixel, and secondly, the operation of the controller to read and digitise the output signal from that pixel.

The current pixel is selected by addressing the two shift registers. A pixel can only be read or reset while it is being accessed and only one pixel can be selected at any time. The shift registers each have two inputs, a data bit which loads a value into the first cell when selected and a clock bit which advances the set bit into the next cell. A cell with a bit set selects that row or column. There is also an asynchronous clear function which resets both shift registers. An extra complexity arises because the column shift register which selects the pixel in the current row is dual edge triggered, selecting a new pixel on each edge of the input clock. This is designed to allow a higher pixel rate for a given controller clock speed. Unfortunately the design of the RISC controller in the Astromed system means that it is not possible to change the PIXEL clock independently of the other lines and so a separate code segment is needed to read positive and negative edge pixels.

3.5.1 RISC microcode

The RISC controller is programmed in a very simple assembly language with all the functions of the controller and the clock lines to the array mapped onto an address space. The assembly language has commands to write a bit pattern to these memory locations and read the value of the data returned from the ADC. There is also limited support for loops but no decision or arithmetic operations. To simplify the programming, a series of higher level macros are used to perform repeated operations so that a final program to read the array can be written in only a few lines.

The program files, referred to as microcode, are written as a plain text file. The macro definitions are handled by a pre-processor and the instructions are then assembled to a byte code. This program is then downloaded into the RISC processor's program memory and executed to operate the array. The RISC memory is limited to only 256 program steps and this together with the lack of any function or decision support in the language makes it difficult to achieve some of the more complicated readout schemes which could be devised.

3.5.2 Readout scheme

A typical sequence for a non-destructive read of a quadrant is shown below, the timing diagram is in Figure 3.6 .

```
Start new frame
Loop over 128 rows
    Start new row
    Loop over 64 column pairs
        PIXEL  high
               Read odd pixel
        PIXEL  low
               Read even pixel
    end loop
end loop
```

A reset sequence to clear each pixel would just replace the pixel readout section:

```
Start new frame
Loop over rows to skip at start
    Start new row
    Loop over columns
        PIXEL high
               Reset pixel
```

PIXEL low

Reset pixel

end loop

end loop

To read the entire array this sequence is repeated four times, with a different quadrant output connected to the signal processing electronics. The main steps in the readout scheme are described below.

3.5.2.1 Framestart function

This function is called once at the start of each read of a frame. It loads the shift registers with the correct pattern to begin reading the first row of a quadrant.

1. CLR is set and cleared to reset both shift registers.
2. FSYNC and LINE are set high.
3. The falling edge of LINE clocks the FSYNC bit into the first cell of the row register
4. Clearing FSYNC leaves the first row only selected.

3.5.2.2 Rowstart function

A similar sequence is used to start reading a new row.

1. LSYNC and PIXEL are set high
2. The falling edge of PIXEL clocks the bit into the first column
3. Clearing LSYNC leaves the first pixel in a row selected

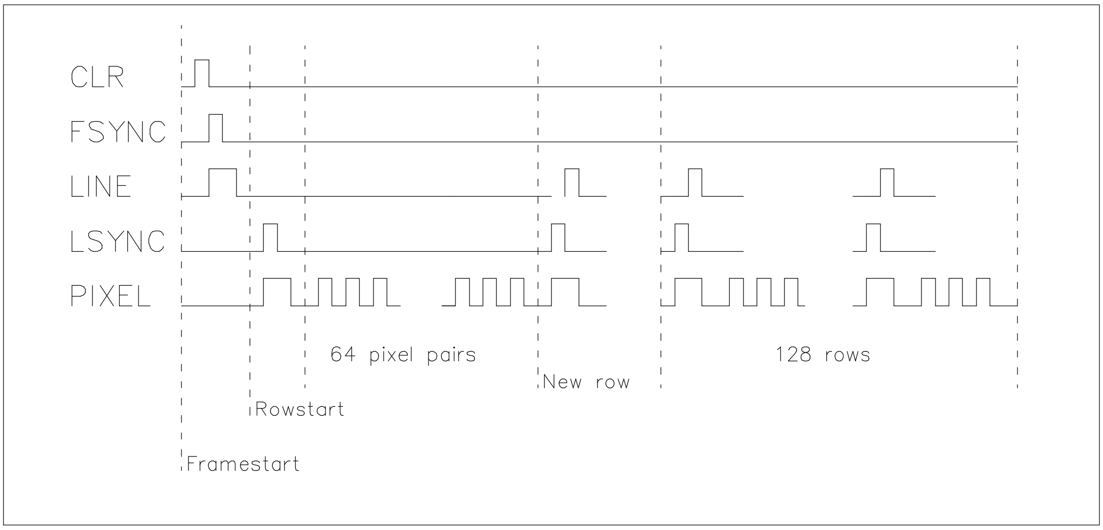

Figure 3.6: Timing diagram for full frame readout

3.5.3 Reading sub-arrays

The multiplexor only allows pixels to be accessed through the sequence of shift registers, as described above. It is not possible to randomly access an arbitrary pixel directly. The process of reading out sub-sections of the array or individual pixels is achieved by quickly skipping over other parts of the array without reading the pixel values. For example to read a 32 x 32 pixel region in the middle of the array, the program sequence would be:

```
Start new frame
Loop over 48 rows at top of the array  i.e. (128-32)/2 rows
    Start new row
end loop
Loop over 32 rows to read
    Start new row
    Loop over 48 columns before region, i.e. (128-32)/2
        Clock pixel without reading
    end loop
    Loop over 32 pixels to read
```

```
                    Read pixel
          end loop
          Loop over 48 columns after region
                    Clock pixel without reading
          end loop
end loop
```

This scheme can be extended to any number of sub-arrays, limited only by the program memory in the RISC controller.

3.5.4 Pixel Readout

The ultimate aim of the camera is to measure the signal charge arriving in each pixel during an exposure. Before explaining the sequence of operations in detail it is useful to examine some of the problems in accurately measuring this quantity.

At the start of an integration the pixel is reset by closing the RESET switch and connecting the charge integration node to the DETBIAS voltage, see Figure 3.2. When the switch is released, the voltage across the capacitor is fixed with an uncertainty of $\sqrt{kTC}$ where k is Boltzman's constant, T is the absolute temperature and C is the capacitance of the integration node. This is due to the finite resistance of the switch and the statistics of the electron flow. In the NICMOS device at normal operating temperature this noise is about 100 electrons rms and would be the dominant noise source if not corrected. The obvious way of removing this uncertainty is to read each pixel immediately after resetting it, but before the shutter is opened, and then store the value of the black level in software. At the end of the exposure the pixel is read again and the black level subtracted.

In a CCD the technique for removing KTC noise is slightly different. The charge from each pixel is measured by transferring it to a single node capacitor near the output amplifier. This capacitor is reset and a measurement taken of the black level immediately after the reset switch is released. The charge is then transferred onto the node capacitor and another reading taken. Since the two measurements are made consecutively, the difference can be measured directly and the CCD controller contains dedicated hardware to do this. In the Astromed controller used here, the black level is fed through a unity gain amplifier to an integrating circuit. After the charge is transferred, the signal level is measured again, but through an inverting amplifier which then integrates in the opposite direction. The voltage remaining on the integration capacitor at the end of the process is amplified and digitised. This process is known as double correlated sampling (DCS). The DCS integration capacitor is reset between pixels by closing a switch which short circuits it. Provided the gains of the two (+1 and -1) amplifiers match and the integration times are identical the device reset noise can be largely eliminated.

In the NICMOS array this system cannot be used to remove KTC noise directly since consecutive measurements cannot be made of the black level and the final signal level of every pixel. Nevertheless, the DCS procedure has the advantage of reducing noise from other sources. While the signal on the pixel is being integrated in the DCS circuit any high frequency noise components are averaged out. Similarly since a differential reading is made between the empty and full levels on the detector output, any drifts in bias supplies or low frequency noise is also rejected.

To achieve the lowest possible read noise, the NICMOS camera uses the principles of the double-correlated readout. The interface between the camera and controller contains a dummy NICMOS output amplifier, which is formed from a potential divider driven from the same bias supplies driving the NICMOS output, see Figure 3.7 below. The potential divider is adjusted to give approximately the same level as the NICMOS output amplifier immediately after a reset, effectively forming a dummy NICMOS output. There is also a low noise FET switch which can connect the controller to either the NICMOS output or to the reference. During a read of a pixel, the first half of the DCS cycle measures the reference level. The switch then transfers to the NICMOS output and a measurement of the pixel level is made for the second half. The difference is then digitised, giving the signal from the pixel. This sequence also allows the sense of the output to be inverted so that the signal recorded for the pixel increases with integrated flux even though the potential on the pixel actually decreases.

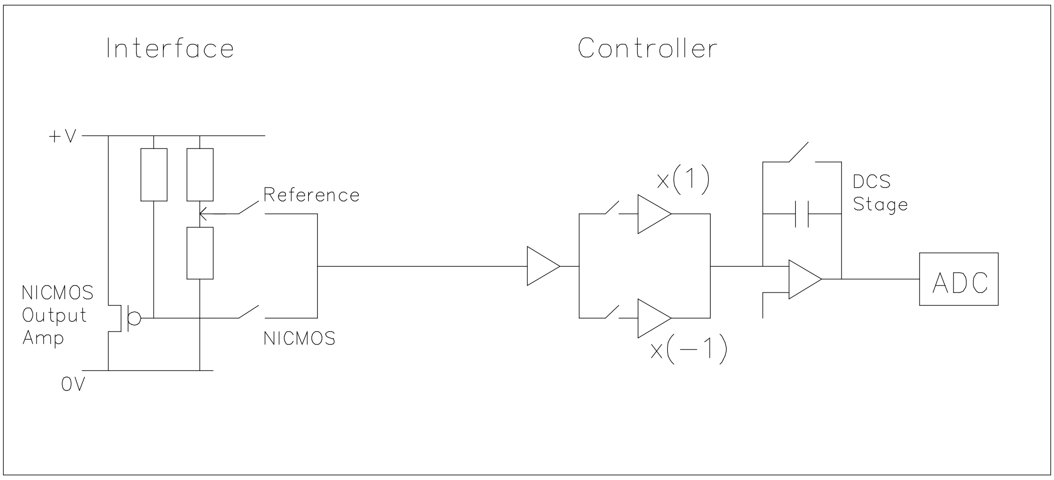

Figure 3.7: Double Correlated Sampling schematic

The timing for the microcode to operate this sequence is shown in Figure 3.8, below. The integration on each half of the DCS cycle is 20µs giving a total pixel time of approximately 45µs..

The REF line connects the controller to the dummy reference and the CHIP line connects it to the NICMOS output. The INT+ and INT- select the DCS stage positive and negative amplifier routes respectively. The RST line discharges the DCS stage capacitor before each measurement is made. Finally the RDADC triggers the analogue to digital converter which reads and digitises the current signal level.

It is this system of rejecting noise on each readout and using two complete readouts to remove KTC that gives the very low noise performance. A plot showing the link between the length of each half of the DCS integration and read noise is shown in Figure 3.11 .

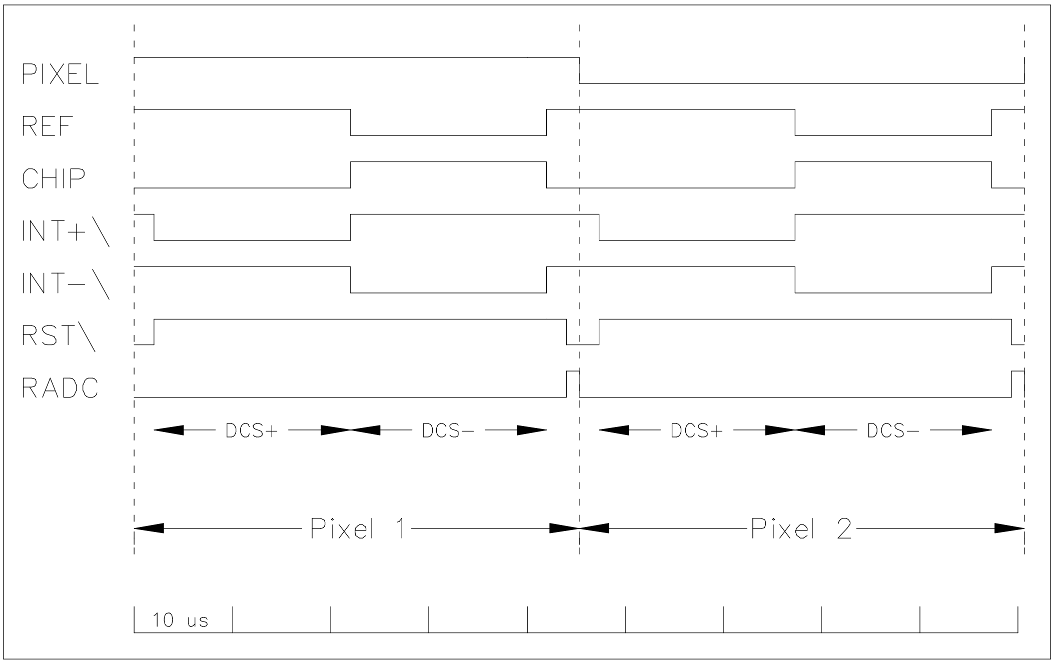

Figure 3.8: Timing diagram for pixel readout

3.6 Performance

There are many areas of the NICMOS device's performance which can be studied usefully. This work concentrates on those features which are important to the operation of COAST, principally read noise and flexible operation. The high background signal due to the dewar design means that it has been impossible to measure some effects involving non-very low signal levels, particularly image persistence and amplifier luminescence.
image persistence.

3.6.1 Quantum Efficiency

Accurate measurements of quantum efficiency rely on having a standard flux source with a stable and well-known signal. In the infrared these measurements are usually made by imaging a temperature controlled black-body. We do not have the facilities to do this and so no detailed measurements of quantum efficiency have been made. However, to provide a general check on the device operation and gain measurements, some experiments were carried out comparing the signal level from a daylight illuminated scene with a visible CCD camera. These confirmed the manufacturer's quantum efficiency estimates to within a factor of two. The device specifications are for a quantum efficiency of 60% in the K band, Rockwell(1993).

3.6.2 Gain

The data received back from the camera are in the form of data numbers (dn) received from the Analogue-Digital Converter (ADC). These values are related to the voltage on the detector integration node by a multiplicative gain factor which is made up from each of the gain components in the signal path. It is relatively simple to measure the gain from the pixel output to the analogue-digital controller by adjusting DETBIAS and measuring the reset level, as described in section 3.6.5. However, to convert these values to a number of electrons requires the capacitance of the integration node, which is almost impossible to measure. To determine the relationship between electrons integrated in the pixel and voltage present at the output amplifier a less direct, statistical technique must be used. The gain was measured by a modification of the variance plot commonly used with CCDs, Mackay(1986).

This scheme uses the property that the arrival of individual photons at the detector follow Poisson statistics. This means that a relationship can be derived between the variation in the signal recorded at different signal levels and the number of photons actually present. A series of exposures are made of the same uniformly illuminated scene at a range of signal levels. At each signal level an image is taken and the variance in the signal recorded. This variance will contain a fixed component from the read noise and a signal dependant component from shot noise. These components add in quadrature to give the measured variance.

Plotting the variance in the measured signal against the mean signal shows two main regions shown in the diagram below.

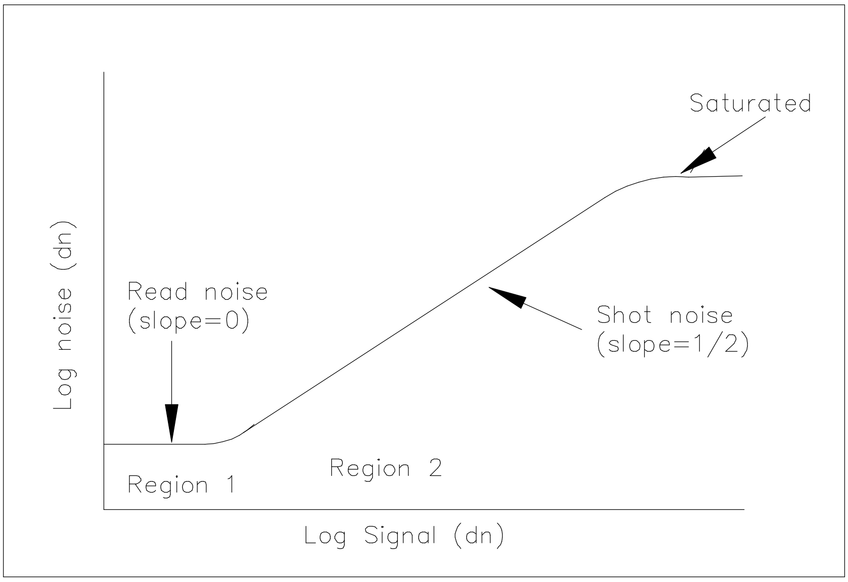

Figure 3.9: Variance curve schematic

1, At very low signal levels the variance is due to the fixed readout noise of the system. This is independent of signal level. The readout noise is a lower limit to the accuracy with which a signal can be measured and is set by the internal noise sources in the components of the device.

2, At higher signal levels the shot noise becomes larger and eventually the readout noise becomes negligible. The variance in the recorded signal is now due to the Poisson statistics of the photon arrival rate. This shot noise has the useful property that the variance in the number of events recorded in any time interval is equal to the square root of the number of photons arriving. As the signal level and variance are measured in data numbers but the shot noise depends only on the number of the photons, the conversion factor between electrons and the signal chain can be determined. This technique is especially useful because it does not depend on any assumptions about the signal processing chain parameters.

Consider measuring a signal S in data numbers, with variance σ_s (in dn).
In the absence of any signal the camera records a black level B with a noise level σ_B.
The system has a gain of E (e^-/dn).

At any signal level the variance is made up from the read noise σ_B (independent of signal level) and a contribution from the shot noise. The signal received (in electrons) plus the read noise squared, must equal the square of the variance measured (also in electrons).

$$S.E + \sigma_B^2.E^2 = \sigma_S^2.E^2$$

or

$$S = E\left(\sigma_S^2 - \sigma_B^2\right)$$

At higher signal levels read noise becomes negligible and so $\sigma_B \to 0$, then:

$$S = E\sigma_S^2$$

and so the gain becomes $$E = \frac{S}{\sigma_S^2} \quad (e^-/\text{dn})$$

The gain can then be simply determined from the signal and the square of the variance at a range of different signal levels

Practical considerations

With CCD devices, the variance plot is generally produced by making exposures of a uniformly illuminated source at a range of exposure times to produce a range of signal levels. For each exposure, the image is read out and a mean and standard deviation of all the pixels in the image are calculated. These two values provide a single point on the variance curve. This is possible because the CCD has a single output amplifier which has the same gain and read noise for every pixel.

The case of a hybrid FPA based on a direct readout multiplexor is more complicated since each pixel has a separate amplifier associated with it. To produce a variance plot with the NICMOS device, a number of separate exposures (typically 100) were made at each signal level. These exposures were then processed to give a mean and standard deviation separately for each pixel. A separate variance plot could then be generated for each pixel. This technique, however, is very time consuming since it requires 100 times as many exposures as the same test for a CCD. It does have the advantage that any gain variation between pixels in the device do not bias the result. This scheme may be useful even when calculating the gain of a conventional CCD, particularly when large pixel to pixel non-uniformity is a problem.

Result

At the highest controller gain setting (1μV/dn) a value of $E = 2.2 \pm 0.4$ e$^-$/dn was obtained. This is equivalent to 420 ± 70 e$^-$/mV from the NICMOS device.

The variance plot in Figure 3.10 was obtained at a higher controller gain of 2μV/dn and so is twice the gain calculated above.

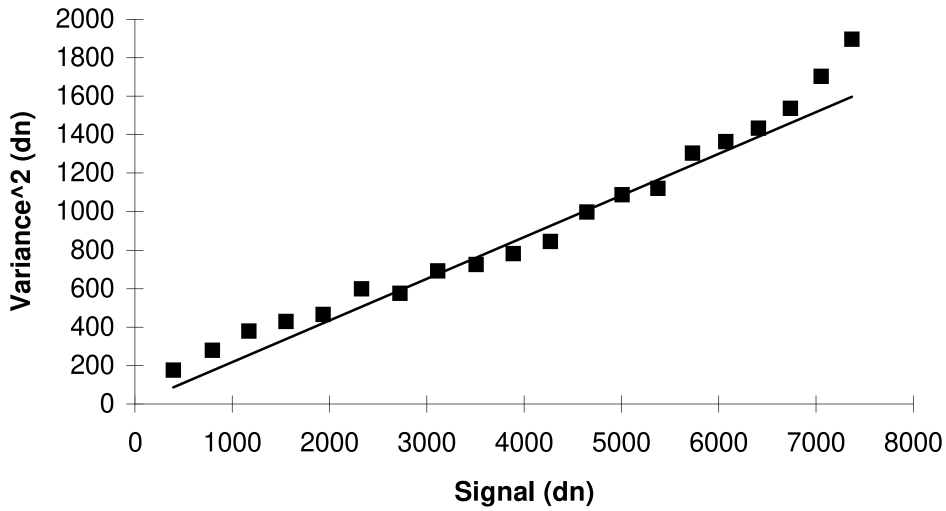

Figure 3.10: Variance plot for NICMOS array

3.6.3 Read Noise

The read noise is the uncertainty in the measurement of a fixed signal due to noise sources in the components of the signal chain. A common definition of this is the standard deviation of a large number of reads of the same signal level. The read noise can be conveniently measured directly from the variance curve. At the lowest signal part of the curve the contribution from shot noise is negligible and the noise level is independent of signal. The variance in this region gives the read noise.

To produce an image with the NICMOS device two reads of the array are needed. After initially resetting each pixel, the array is read out. This first read contains no signal and is called the black level. At the end of the exposure the array is read again. The difference between the two values for each pixel gives the signal which arrived during the integration. Since each of these reads have the same magnitude of read noise superimposed on them and the noise from each read is uncorrelated, the overall noise on the final image is a factor of $\sqrt{2}$ times the read noise.

Using the non-destructive readout property of the array this technique can be extended. If the array is read N times without the signal level changing then the uncertainty in the measurement of the true signal contained on the array can theoretically be reduced by a factor of $\sqrt{N}$ compared to a single read. Clearly this technique would only be useful when the signal level and background are extremely low and the read noise of this system dominates the signal or background shot noise. Although the read noise dominates in the case of COAST there is not enough time to perform more than one read.

In general this technique is less useful. The process of reading the array causes light to be emitted from the output amplifier and shift registers due to LED effects in the silicon switches. The shot noise on this extra signal quickly cancels the advantage of performing multiple reads.

Infrared arrays have traditionally been operated at very high pixel rates. At longer wavelengths, high source and background fluxes mean that it is a challenge to readout the device before saturation occurs. Under these conditions the shot noise on the signal dominates and read noise is negligible. Using devices with a very low dark current such as the NICMOS array and in applications such as interferometry or high resolution spectroscopy with very low signal levels the read noise becomes important. The techniques to reduce the read noise, used in the COAST camera are borrowed from CCD cameras where the intrinsic read noise of the device approaches a few electrons.

The details of the readout scheme were described above in section 3.5.2. The important features are, the use of a double-double correlated sampling arrangement which filters interference and drifts in supply levels, together with the use of very low pixel rates. The effect of pixel rate on read noise is shown by the plot below. The horizontal axis shows the time for each half of the double-correlated sample procedure, the total time to read a pixel is twice this plus approximately 5µs for the ADC conversion and data transfer. The COAST camera operates at 45 µs /pixel (22kHz pixel rate). At slower rates than this there is no improvement in read noise, probably due to 1/f noise in the output transistor on the NICMOS array.

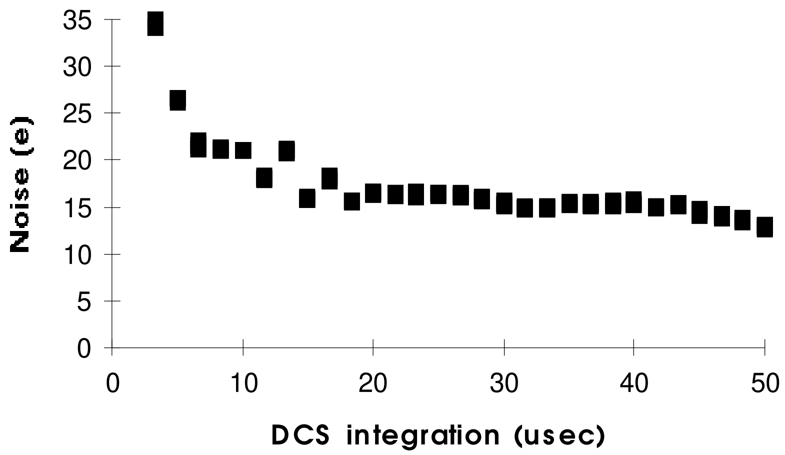

Figure 3.11: Effect on read noise of DCS integration time

Results

The read noise of the COAST camera at 20 kHz pixel rate, assuming a gain of 450 e^-/mV, is 16 electrons. The plot below shows a histogram of read noises from single quadrant for a single non-averaged read at this rate. Approximately 95% of the pixel values are included.

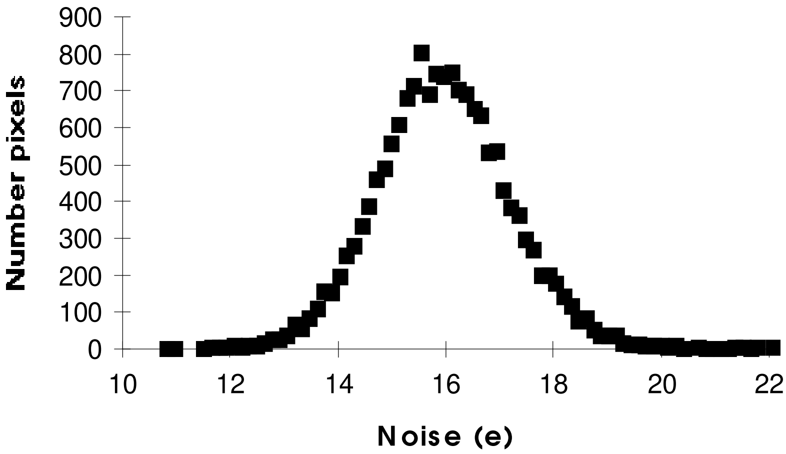

Figure 3.12: Histogram of noise values from one quadrant

Practical considerations

The read noise could be measured by taking a large number of exposures with the shutter closed and calculating the standard deviation on the signal recorded by each pixel. This follows from the definition of read noise given above. In the case of the COAST camera operating at slow pixel rates, the dark current is large enough that in the time needed to make two reads of the entire array the shot noise on the dark current is larger than the read noise. The system is actually shot noise limited even for exposures of zero seconds!

An alternative technique is used which makes exposures using only one line of the array at a time, this is shown in the diagram below.

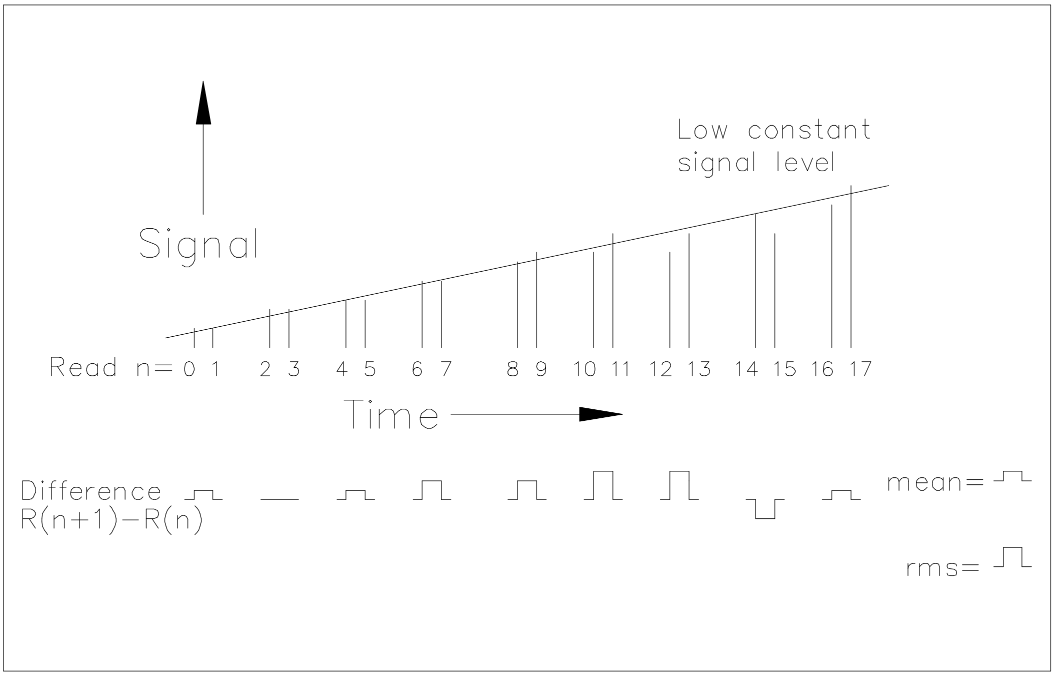

Figure 3.13: Diagram of single line noise analysis

A single row of the array is reset and the same line is then read non-destructively a large number of times. Pairs of these reads are then subtracted to give individual exposures. The signal rate is low enough that the signal arriving between two reads is negligible. The mean and standard deviation of these differences from the same pixel in each exposure is then calculated. The mean is a small positive number giving the overall average signal rate during the exposure. The variation in individual readings are due mainly to the read noise if the signal rate is low enough that shot noise can be ignored. The standard deviation of these pairs of measurements is the read noise multiplied by $\sqrt{2}$, since two reads contribute to each difference. This process can be repeated for each row in the array to produce a noise map of the entire device.

3.6.4 Dark Current

The camera operates in a single stage liquid nitrogen cooled dewar at a temperature of approximately 80K, the details of the cryogenics are described in section 3.4.1 .

The dark current is measured by replacing the cooled filter with a metal blank and taking a long exposure. The results in Figure 3.14 show a mean signal rate of 90 e^-/second. The manufacturers data suggest that the dark current at 77K should be 10 e^-/second, Vural(1990). The extra signal is due to the inefficient cold stop allowing the detector to 'see' warm parts of the inside of the dewar.

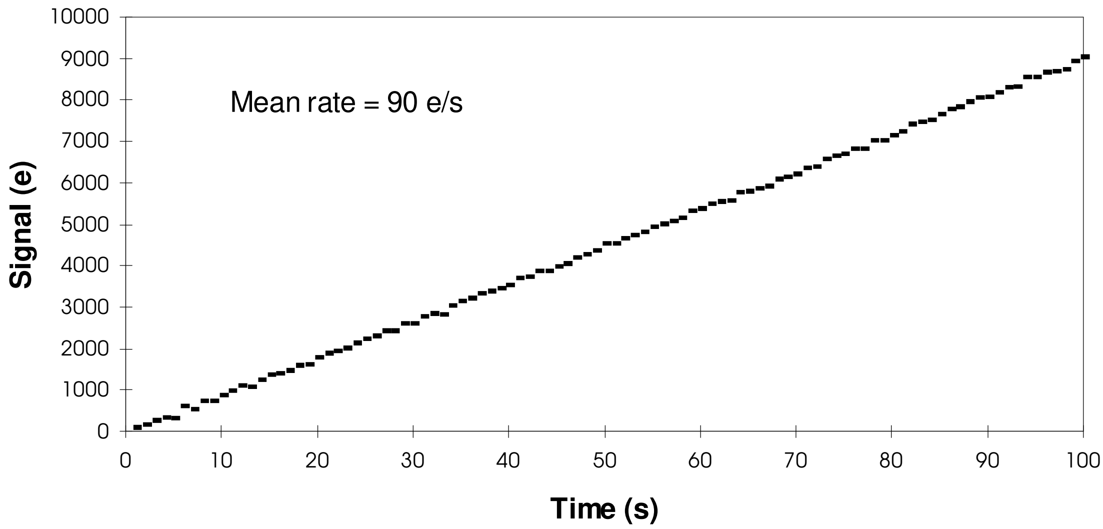

Figure 3.14: Dark-Count rate

3.6.5 Linearity

An ideal detector is linear, the output signal is directly proportional to the photon arrival rate during the exposure. Infrared arrays using a direct readout multiplexor are not linear. The arriving photon generated signal acts as a constant current discharging a capacitor, the signal at the end of the integration is measured from the voltage remaining on the capacitor. Since the detector diode is connected to the capacitor, the detector bias voltage decreases as the voltage remaining on the capacitor falls. This effect reduces the collecting efficiency of the detector and so should produce a system where the signal rate decreases with integrated signal. In tests, the relationship is slightly different as shown by the plot in Figure 3.15.

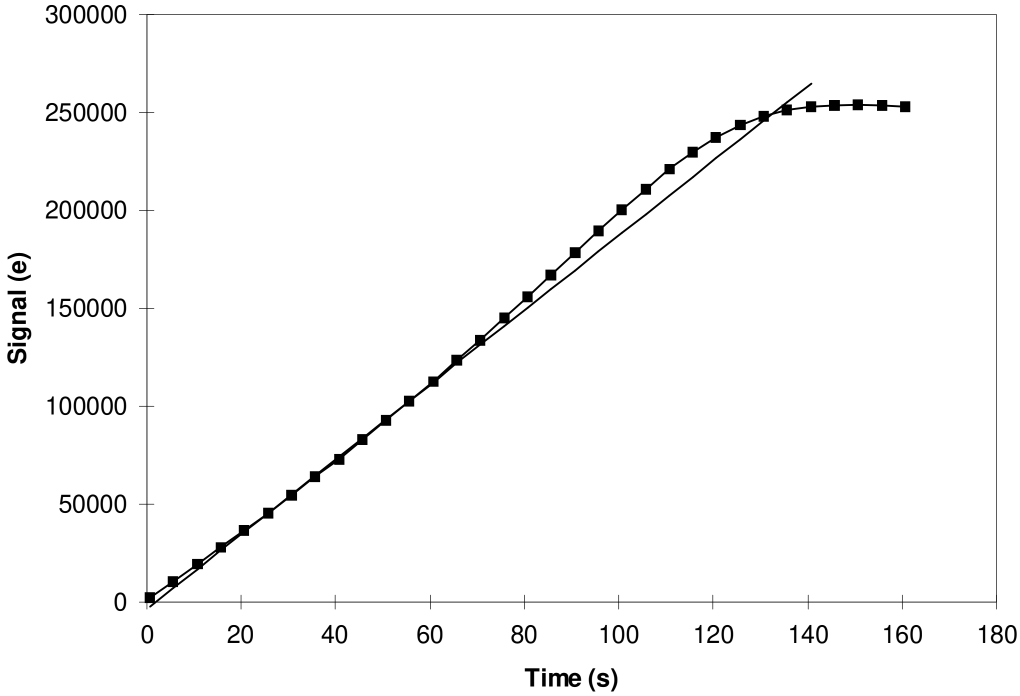

Figure 3.15: Linearity of integrated signal and output

As the pixel becomes approximately half full, the measured signal rate increases. There are two possible explanations for this, either the gain of the pixel amplifier increases or the capacitance of the detector node decreases.

The linearity of the pixel amplifier can be tested directly. When the RESET switch is on, the integration capacitor is connected directly to DETBIAS and so the pixel amplifier reads this voltage. If the DETBIAS is adjusted and the output reset level measured at the gain of the components inside the NICMOS array can be measured. The data in Figure 3.16 show that the device is linear over a large range of detector bias levels.

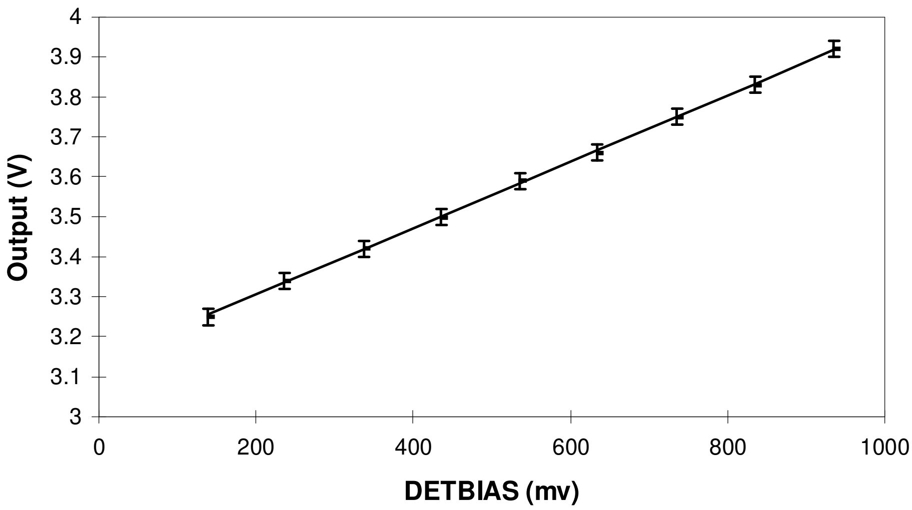

Figure 3.16: NICMOS amplifier linearity

The alternative suggestion is that the capacitance of the pixel decreases as the voltage across it decreases. This may be possible and would also explain the reduced full well capacity at low detector bias levels seen in Figure 3.17 .

3.6.6 Detector bias and well capacity

The well capacity is the maximum number of pixels which can be stored on the capacitor in the pixel unit cell. At the start of an integration the pixel capacitor is charged to a fixed bias voltage DETBIAS and then the photon current discharges the capacitor during the integration. The maximum signal range depends on this bias voltage and the capacitance of the detector. The full well capacity can be measured by making an exposure with the detector exposed to a strong signal and measuring the signal recorded by each pixel after no further change is seen.

The device is normally operated at a DETBIAS of 0.5V giving a well capacity of 250,000 electrons. The well capacity was measured over a range of bias levels from 0.2 to 0.9

volts. At the higher bias voltages the number of 'hot' pixels increased rapidly and the experiment was stopped at 0.9V to avoid any permanent damage to the array.

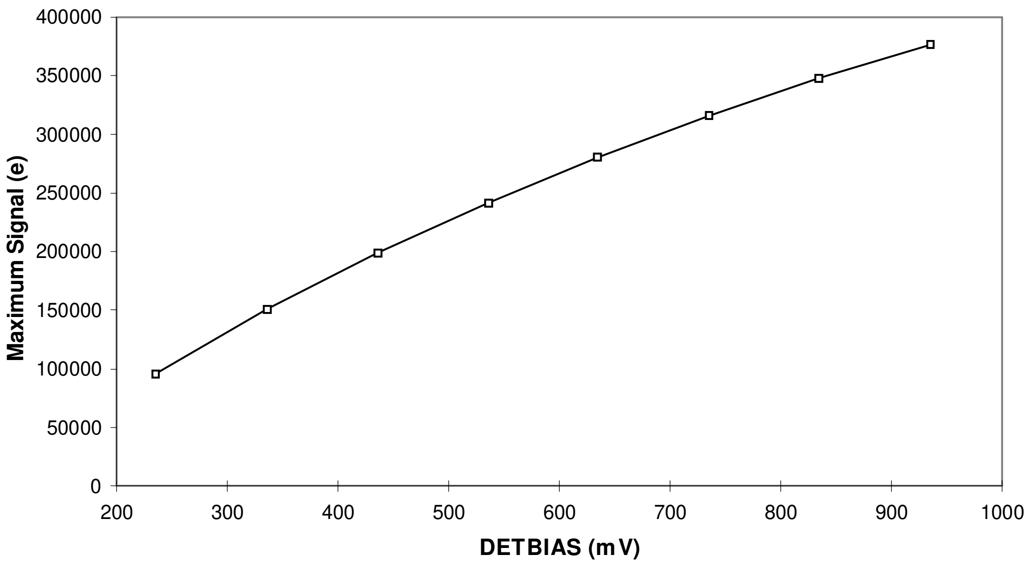

Figure 3.17: Well capacity as function of detector bias

The curve evident in the results implies that the well capacity, and so the capacitance of the integration node, decreases at low detector bias levels.

3.6.7 Persistence

The NICMOS device has a feature reported by a number of groups. When a device is illuminated with a bright source, the pixels which received a high signal have an increased dark current which continues after a number of resets. I tried to measure this effect but any increased dark current in these pixels was hidden by the large background signal in the COAST camera.

3.6.8 Amplifier luminescence

The amplifiers and multiplexor circuitry on the NICMOS device are made of silicon and emit light as a side effect of their operation. During a long exposure the output amplifier is turned off by de-selecting all the quadrants and so removing the bias supply to the output source follower. The shift registers can be turned off by sending the CLR instruction which clears all the cells. However, the amplifiers must obviously be used to read the image at the end of the exposure. In this case the effect is seen as an extra signal which is injected into the array with each read, but is independent of the exposure time.

This effect limits the advantages of the multiple read scheme discussed in section 2.3.4.1. Each non-destructive read of the array introduces an extra signal into the pixel. This extra signal , N electrons, has a shot noise with an rms value of $\sqrt{N}$ associated with it. This extra noise quickly destroys the reduction in read noise obtained by multiple reads.

The data in Figure 3.18 implies that the extra signal comes mainly from the output amplifier and partially from the shift register. There is an increase in the dark rate seen across the entire array. It is difficult to say whether any of this is due to emission in the unit cell circuitry or if it is all light from the output amplifier transmitted through the silicon. The next generation of 1024 x 1024 pixel arrays have been designed so that the output amplifier can be disconnected, and a separate off-chip amplifier be used to avoid this problem.

Result

The camera was fitted with the dark slide used to measure the dark current. The array was reset and then read non-destructively a number of times over a period of 200 seconds. The number of readouts in each exposure was varied from 2 to 200 with the time between them adjusted to keep the same total exposure time. The signal due to real dark current is proportional to the total exposure time but independent of the number of reads. This was removed by fitting a constant rate to a plot of signal vs. number-of-reads, and extrapolating to the case of no reads. The plot in Figure 3.8 shows the extra signal detected as a function of the number of reads. The extra signal was measured over 16x16 pixel regions placed near the output amplifier, at the furthest corner from the amplifier, and alongside the edge of the chip nearest the shift register. The average over the entire quadrant was also measured.

A summary of the results is shown in Table 3.1. These results were obtained from a single quadrant at a pixel rate of 22kHz.

Region	Whole Array	Near Output	Away from Output	Shift Register
Signal/pixel/read (e^-)	25	100	15	40

Table 3.1: Signal due to amplifier emission

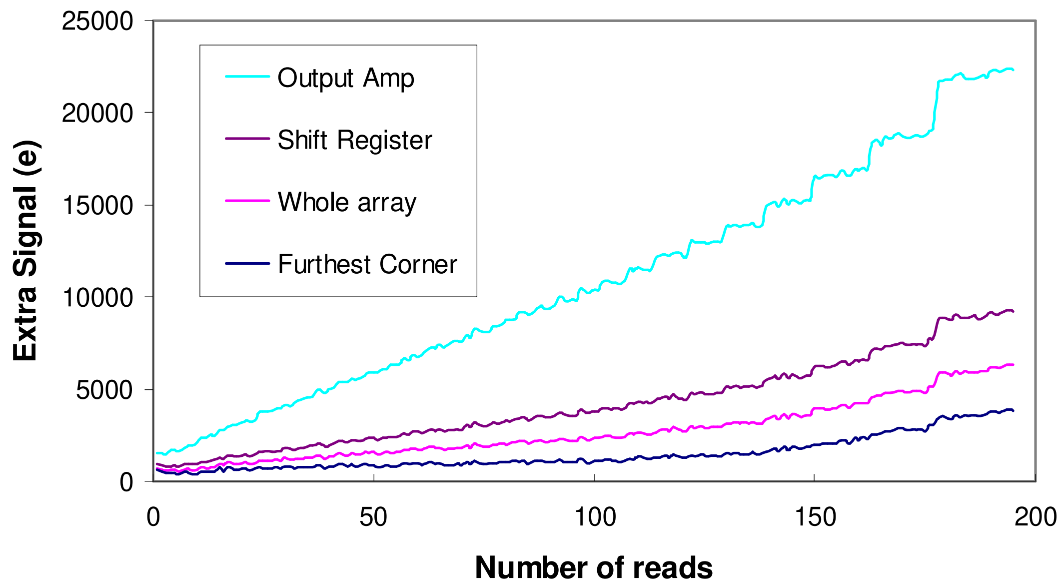

Figure 3.18: Extra signal from amplifier glow

3.7 Readout mode

The COAST system has extreme requirements of the infrared camera. The signal on the four outputs of the beam combiner must be measured with an exposure time of less than 1 millisecond. The data must then be transferred to another computer before the next exposure is finished. Finally the dead-time, where new photons arriving are lost, between exposures must be minimised.

The optical system described in section 4.3.5 light from the four output beams on to four pixels in the corners of the array. These pixels are arranged in a square with the first two beams focused on two pixels in the same row near the top of the array and the second two beams on two pixels near the bottom, to simplify the software the pairs 1,3 and 2,4 are on the same columns. This is shown in the diagram below.

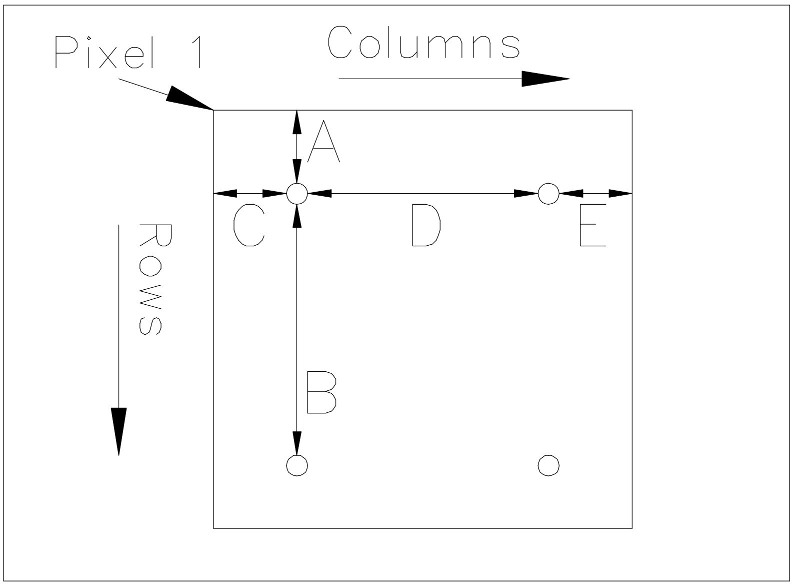

Figure 3.19: Position of COAST output images

The signal from these four pixels is read repeatedly by the camera control system. The microcode to do this is described by the code below. Each time the sequence is run the camera first reads the final signal level on the pixel. That pixel is then reset and the black level measured for the start of the next exposure. This process is repeated for each of the four pixels. The process of reading a pixel transfers the data from the ADC back to a buffer in the RISC card. This memory is only available to be read by the host computer once the program finishes. This is done while the next iteration of the program is running. The signal integrated on each pixel during the exposure is determined by subtracting the black level from the previous read from the signal level immediately before the reset.

The extra read at the end of the program is necessary due to the pipeline architecture of the ADC which transfers a value to memory when it starts to process the next signal, the dummy read acts to flush the last value into memory. The letters in ( ) are the distances on the array shown in Figure 3.20.

The microcode sequence to read four output beams from COAST is shown below:

1. FrameStart	
2. Skip the unused lines at top of the frame (A)	
3. Skip the unused pixels to first beam (C)	
4. Read pixel 1	Value 0
5. Reset pixel 1	
6. Read pixel 1 again	Value 1
7. Skip the unused pixels to second beam on same row (D)	
8. Read pixel 2	Value 2
9. Reset pixel 2	
10. Read pixel 2 again	Value 3
11. Skip the unused rows to beams 3 and 4 (B)	
12. Skip the unused pixels to first beam (C)	
13. Read pixel 3	Value 4
14. Reset pixel 3	
15. Read pixel 3 again	Value 5
16. Skip the unused pixels to second beam on same row (D)	
17. Read pixel 4	Value 6

18. Reset pixel 4
19. Read pixel 4 again Value 7
20. Skip the unused rows to end of frame (E)
21. Read a dummy pixel
22. Wait for approximately 500 microsecond

The photon signal is arriving in the pixel continually as the shutter is kept open indefinitely. The exposure is defined by the read immediately after the pixel is reset and the read just before it is reset again the next time the program runs. The only dead time occurs while the pixel is actually being reset.

$$Signal_{pixel=1,run=1} = Value_{run=1,value=0} - Value_{run=0,value=1}$$

The values 0,2,4,6 on the first run of a sequence obviously contain no useful data and the values 1,3,5,7 of the last run are never used. The actual algorithm is more complicated due to byte swapping and order swapping in buffers but the above outlines the principle.

The diagram in Figure 3.20 shows the timing of this sequence. The signal is shown increasing in the pixels upto the point at which they are reset.

Since the time taken to read a pixel (45μs) is comparable to the time taken to execute the code it is clear that although the pixels all receive the same exposure time these exposures are not synchronised. This can be accounted for in the system which combines the recorded signal with the path trolley position to determine the fringe phase.

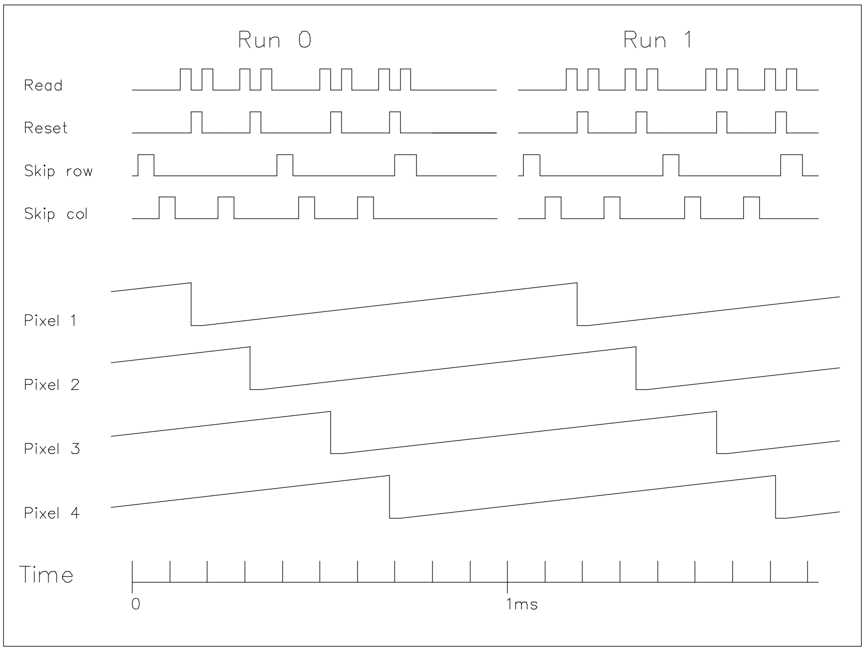

Figure 3.20: Timing of COAST four beam read

3.8 Conclusion

This chapter has described the infrared camera built for COAST. The NICMOS detector is one of a new generation of infrared devices which offer array sizes and noise levels which lag only a little way behind visible CCDs. This project has also shown that not only do these new devices allow infrared cameras to be built easily by groups with only visible camera experience, but using techniques common in CCD electronics has allowed a far better performance to be achieved with the infrared array.