Infra-red Imaging and Spectroscopy with HAWAII and PICNIC Arrays.

C D Mackay, M G Beckett,R G McMahon, I R Parry, F Piche, K A Ennico,
M Kenworthy, R S Ellis, A Aragon-Salamanca.
Institute of Astronomy, University of Cambridge, Madingley Road,
Cambridge CB3 0HA, U.K.

ABSTRACT

This paper describes the results of a test program to evaluate four Rockwell HAWAII (1024 x 1024 pixel) and two PICNIC (256 x 256 pixel) near IR array detectors with a view to their application in imaging, spectroscopy and in fast telescope tracking and interferometer fringe detection. Results of the laboratory tests of the arrays are presented, together with a guide for their general operation.

1. INTRODUCTION

A new generation of infra-red imaging detectors is now being manufactured by the Rockwell International Science Center in California. These devices use a mercury-cadmium-telluride (MCT) layer as the detector on a sapphire substrate which is then indium bump bonded to a silicon CMOS multiplexer which in turn provides pixel addressing and so access to each of the pixels on the array. The Institute of Astronomy, University of Cambridge, UK, has been very fortunate to be able to purchase four Science grade HAWAII detectors, each of 1024 by 1024 pixels (18.5 micron pixels, 18.94 mm square) as well as two of the PICNIC arrays each of 256 by 256 pixels (40 micron pixels, 10.24 mm square). The four HAWAII detectors are in use in the CIRSI infra-red survey camera described elsewhere in these proceedings[2], and the two PICNIC detectors are used as part of the COHSI instrument which has also described in another paper elsewhere in these proceedings[3].

This paper describes the experiences that we have had in driving these detectors and the results the we have obtained. In this paper we will try to give some general guidance to other workers who are the process of making systems based on these devices.

There is considerable experience in the astronomical community of using CCDs as scientific imaging detectors. The new Rockwell MCT arrays are generally much easier to use and to get working well than are CCDs and this paper will emphasize the differences between CCD and MCT X-Y addressable technology.

Details of the internal structure and basic properties of these new Rockwell devices may be found in references [1].

2. DEVICE ARCHITECTURE

Charge Coupled Devices (CCDs) are structured so that the light sensitive region is also where the accumulated charges are transported across the CCD as part of the readout process. A CCD is a continuous sheet of silicon covered in electrodes. The charge collected under each electrode during the exposure is progressively passed through adjacent collections sites to the output register as part of the parallel transfer process. The output register is then clocked in a similar way so that the charge is progressively passed from element to element of the output register and then to the output amplifier.

Part of the SPIE Conference on Infrared Astronomical Instrumentation • Kona, Hawaii • March 1998
SPIE Vol. 3354 • 0277-786X/98/$10.00

With X-Y addressable arrays such as these made by Rockwell, the structure is very different. Each pixel is entirely independent, and the signal from that pixel goes nowhere near the components that make up any other pixel. Each pixel of the X-Y array consists of a capacitor which is charged before the exposure to a fixed level (the reset level). Photons are converted to carriers which then progressively discharge the capacitor throughout the exposure. At the end of the exposure the voltage that remains across the capacitor is measured by connecting the individual buffer amplifier that is part of each pixel to a common video line through two transistor switches, one of which is enabled as part of the row address electronics and the other which is enabled as part of column address electronics.

In the case of a CCD it is generally true that the act of reading out the charge destroys it. With an X-Y addressable array the reset of the pixel is completely independent from its measurement and so reading the array is intrinsically non-destructive.

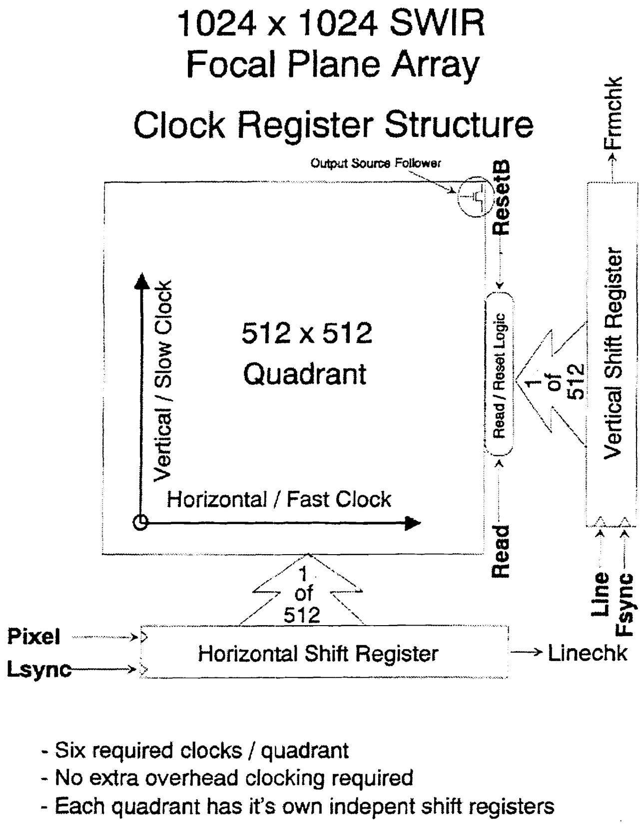

Figure 1: The layout of one quadrant of a HAWAII device. The device consists of four quadrants to give a total of 1024 x 1024 pixels, each of 18.5 microns square. (Courtesy Rockwell International Science Center)

The general structure of these devices shown in figure 1. Digital electronic circuits (made with the same complementary metal oxide on silicon (CMOS) technology as are standard memory and microprocessor integrated circuits) are provided at the edge of the device. They are arranged so as to enable each of the row and column enable lines in sequence as they are clocked and so allow the full 2_D array of pixels to be read in turn.

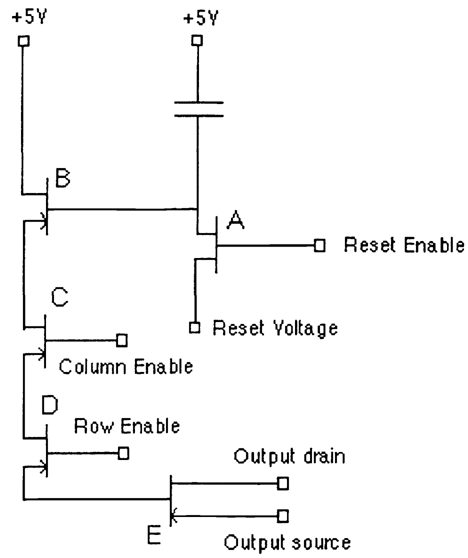

Figure 2: The internal structure of the HAWAII and PICNIC detectors.

The structure of the basic unit cell of these X-Y addressable arrays is shown in figure 2 as the capacitor and transistors A, B and C. The detector element is represented by the capacitor at the top of figure 2. It is charged to a fixed voltage at the beginning of the exposure by enabling the reset line and so connecting the storage capacitor to a fixed voltage, the reset voltage, through transistor switch A. Because of the finite capacitance C of the detector element there will always be a reset noise of $\sqrt{(kTC)}$, the uncertainty in the voltage across a capacitor with absolute temperature T. This noise is generally quite large, of the order of a few hundred electrons rms, and so the precise voltage to which each detector element is reset must be measured before the exposure. This is done by reading (non-destructively) the whole array and recording the values from each pixel in turn. The voltage on each pixel is measured by connecting one pixel at a time to the common output system. The pixel is addressed by column and row enable lines which drive enable transistor switches C and D, connecting the addressed pixel via the buffer transistor within each pixel (transistor B) to the common video output line. This is connected to a single transistor at the edge of the device (transistor E) which allows the impedance of the device to be matched to the outside world. Transistor E may be bypassed in these Rockwell devices and an external buffer amplifier used. In these tests we have only used the on-chip amplifier transistor E. At the end of the exposure the array is again read in exactly the same way and the values measured are recorded after subtracting the measured values at the beginning of the exposure. This removes the reset noise in exactly the same way that double correlated sampling is used to remove the reset noise with a CCD. The principal difference is that with a CCD the double correlated sampling process happens within the time that each pixel is read-out and so is good at removing a wide range of frequencies in the noise spectrum of the CCD output amplifier. With X-Y addressable array, however, there is no such suppression because of the read and reset measurements are generally separated by a long period of time, typically one second.

In order to try and achieve a much better suppression of the various noise contributions that there are in the system we have used a rather different approach to the readout of these devices. In practice the noise performance of any imaging system at the telescope is significantly affected by things such as electronic pickup of interference from motors, computer monitors and ground returns. The double correlated sampling systems used in CCD cameras can give quite good suppression of many of these noise sources. What we have done is to create a equivalent dummy output which is supplied with the same analog voltage and ground wires that are connected to the analog circuitry of the X-Y array, and then to arrange to switch between this reference value and the signal value produced by the X-Y array within the time each pixel is read. This has the additional advantage of making the output of the chip look very much like the CCD output so that we can then use a CCD controller with its double correlated sampling architecture to give us an additional level of noise suppression. In this way we have two forms of double correlated sampling, one which suppress the noise on frequencies which are comparable with the pixel frequencies, and the other that provides reset noise suppression by subtracting the images derived at the beginning and the end of the exposure.

This method of quadruple correlated sampling (QCS) is a very effective in creating a system that is fairly immune the main sources of interference at the telescope. We used a fast commercial CCD controller made by LSR-AstroCam Limited, Cambridge, England (the Capella model 4100), which has a PCI interface to a computer running Windows 95. The controller was 12-bit (4096 grey levels) for these tests, though we will shortly upgrade it to 14-bits (16384 grey levels). We operate the camera system with the PixCel software package also produced by LSR-AstroCam, modified provide automatic reset image subtraction as part of the data taking procedures as well as providing integration with the local telescope control system (TCS)[2].

The controller we used has only a single input channel. Each HAWAII detector and each PICNIC detector is structured with four relatively independent quadrants. In order to allow us to read out (in the case of CIRSI) the 16 quadrants of the four HAWAII devices consecutively we have constructed an interface box which allows us to switch (with miniature relays) between each of the 16 outputs of the four HAWAII devices. This interface box also contains the electronics that provide the QCS operation described above. It buffers the output signals before passing them back to the 4100 controller over low impedance signal lines.

We have only ever used the on-chip buffer amplifier transistor. Other workers have used an external buffer amplifier but we feel that the added capacitance is likely to produce an overall noise performance that is less particularly at the higher pixel rates.

3. RAW DATA COSMETICS

The data that are obtained from a completely dark device after it has been fully reset are still very non-uniform. With a CCD the non-uniformities in response are principally due to pixel to pixel non-uniformities associated with the exact location of the covering electrodes. With an X-Y addressable array it is important to remember that each detector element has its own buffer transistor to connect it to the common video line. Despite the very high standards of manufacture of the silicon multiplexer manufactured by Rockwell for their devices it is inevitable that the various thresholds of the transistors in the unit cell will vary significantly between cells giving quite marked pixel to pixel non-uniformities in the final image. In practice the magnitude of the non-uniformities is typically a few thousand electrons. This means that the images that are taken without reset correction can look much poorer than they are once the reset image has been subtracted.

It is interesting to note that unlike CCDs the X-Y addressable array is an extremely complicated device. Each pixel is made up of three transistors so that the whole device contains in excess of three million transistors, giving them a complexity comparable to that of the latest Pentium class processors. The transistors marked A,B and C in Figure 2 are present in each and every pixel together with the charge storage capacitor.

4. DARK CURRENT NON-UNIFORMITIES

The mean dark current of both HAWAII and PICNIC devices is very low indeed. The HAWAII devices that we have tested show a mean dark current level of about 0.1 electrons per pixel for three devices and about two electrons per pixel for a fourth (somewhat older) device. For the PICNIC devices the mean dark currents are about 0.6 electrons per pixel per second, significantly higher than the HAWAII devices mainly because the size of the pixels in the PICNIC devices is 40 microns square rather than 18.5 micron square for the HAWAII devices.

These figures relate to the mean dark current at a temperature of 80K. There are, however, a significant number of hot pixels that have a dark current which is very much higher than this mean. The devices the we have show that approximately 10 percent of all pixels have a dark current that is 5 times the mean or above, 4 percent have a dark current of 10 times the mean or above, and one percent have a dark current of 20 times the mean or above. These hot pixels behave in an entirely predictable and repeatable way and so do not cause a great problem as they may be corrected fairly easily.

The tests reported here were done with the PICNIC devices at a temperature close to 80K and the HAWAII devices operating at temperatures between 90 and 110 K.

5. AMPLIFIER AND SHIFT REGISTER GLOW

The HAWAII and PICNIC devices are based on an earlier technology developed originally as part of the NICMOS project for the Hubble Space Telescope. The NICMOS devices showed quite significant glow from the output amplifier and shift register structures on the edge of the device. This glow is thought to be caused by hot electrons in the silicon multiplexer which emit radiation in the near infrared and which is detected efficiently within the silicon multiplexer itself. Although we expected to see more evidence of this we have never seen any significant glow from the HAWAII devices which we have. We have, however, seen significant levels of glow in the PICNIC devices. In order to minimize this it is essential that the output of the on-chip buffer amplifier is connected only during readout. In addition the ability to vary slightly the analog and digital voltages provided to the detector can be very helpful since a small reduction in the voltages does seem to reduce greatly the amount of glow experienced.

6. COSMIC RAY RATES

We have made careful measurement of the cosmic ray detection rates which we see with the HAWAII devices. The rates which we have found are very similar to those seen with CCDs. They are typically one event per square centimeter per minute. The events have energies in the range of about 3000 to 35,000 electrons total with most of the energy confined to a single pixel. The way these measurements were done required the events to be recognized visually and counted manually. It is possible that a significant number of lower energy events were missed by these tests. There is no evidence that the energies of the events which we did see are anything other than randomly distributed over this range, with no sign of a concentration of the energies into particular peaks.

7. PIXEL READ-OUT AND CLOCKING RATES

Rockwell suggest that a suitable readout rate for the HAWAII devices is approximately 300 KHz but that it is possible to run these devices as fast as 1 MHz. The CMOS registers may be clocked at a much higher rate, and we normally operate these at a rate of 12 MHz. We understand from Rockwell that the CMOS circuitry will operate at 20 MHz quite happily. This ability to clock the CMOS registers at high speed is a considerable advantage over CCDs. In applications where it is important to be able to read out a small sub-array quickly the time it takes to clock through the parallel and serial registers of a CCD in order to get to that sub-array can be a major part of the readout time of a CCD even of modest size. In the case of an X-Y addressable array this access time can be made very small indeed, allowing devices to be read out at very high frame rates. We have in fact used the system described here to read out an 8 by 8 subarray at a frame rate of 16,000 per second with a read-out noise of about 15 electrons. With a frame rate of 2000 per second the read noise is approximately 8 electrons rms.

The main reason to run these devices at a lower readout rate is that the readout noise will be better but also that the gain of the on chip amplifier goes down markedly as the readout rate increases. This is likely to be worse if an external buffer transistor is used rather than the one provided on the chip. The results of these tests for a PICNIC array are shown in figure WWWW. Results for the HAWAII arrays are very similar indeed. The gain of the output amplifier transistor is approximately halved by operating the device at 1 MHz pixel rate compared to the gain at 100 KHz pixel rate. The attenuation is even stronger at higher rates. We do not have measurements to present here even though we have operated the device at pixel rates as high as 2 MHz.

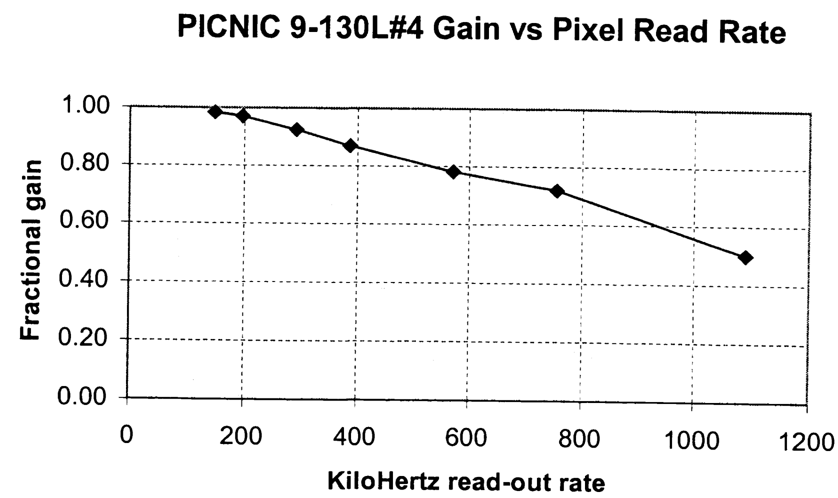

Figure 3: The reduction of the system gain as a function of pixel read out rate for a PICNIC device. Most of this attenuation is due to response roll-off in the on-chip output transistor.

8. READOUT NOISE

All the measurement we have made of read-out noise have been made with a single read. It is clearly possible to take multiple reads as these are intrinsically non-destructively and then to average them in order to produce a lower readout noise. There is nothing we have seen to suggest this will not work perfectly well with these devices but we have not used this ourselves. We also have only used the on-chip amplifier and have never used an external buffer transistor in its place.

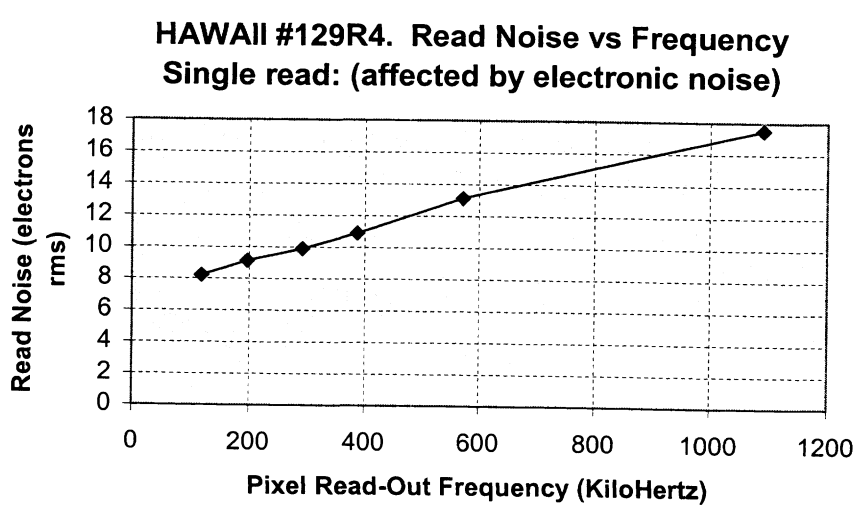

Figure 4: The variation in read-noise measured with the variance method[1] as a function of pixel read-out rate.

Results of the measurement of read-out noise using the well-known variance technique[5] as a function of frequency are shown in figure4. At present our systems suffer from low-level interference that makes our measurements poorer than they should be. Under certain circumstances we have been able to achieve a noise floor of less than about four electrons rms. The readout noise for the PICNIC chips is typically a factor of two worse than that of the HAWAII chips because of the higher capacitance on the node of the on-chip buffer transistor.

9. DEVICE LINEARITY

In order to understand what aspect of driving these arrays affects linearity it is important to understand the underlying structure of the unit cell. The voltage applied to the reset line is connected to the detector capacitor when the reset transistor is enabled. Charge created by incident photons progressively discharges this capacitor. Clearly the higher the voltage applied across the detector the more charge can be contained in it and therefore the full well capacity of the pixel is increased. When the voltage on the detector is connected to the on-chip amplifier transistor the bias level on the detector is going to be a significant part of the bias on the gate of the output transistor. If that capacitor is fully charged then the voltage applied to the buffer transistor can be such as to reduce its gain quite markedly because it will not be properly biased. This may be compensated by biasing the transistor to a higher voltage so as to ensure that it will work properly over the whole signal range. The manifestation of this being incorrectly set is that if the bias voltage is too low with a high reset voltage then the linearity of the response at the lowest signal levels becomes very poor. As the signal integrates on the detector the buffer transistor is eventually turned on properly and so the performance becomes very good. It is important to check that this low-level linearity is acceptably good. This is demonstrated in Figure 5.

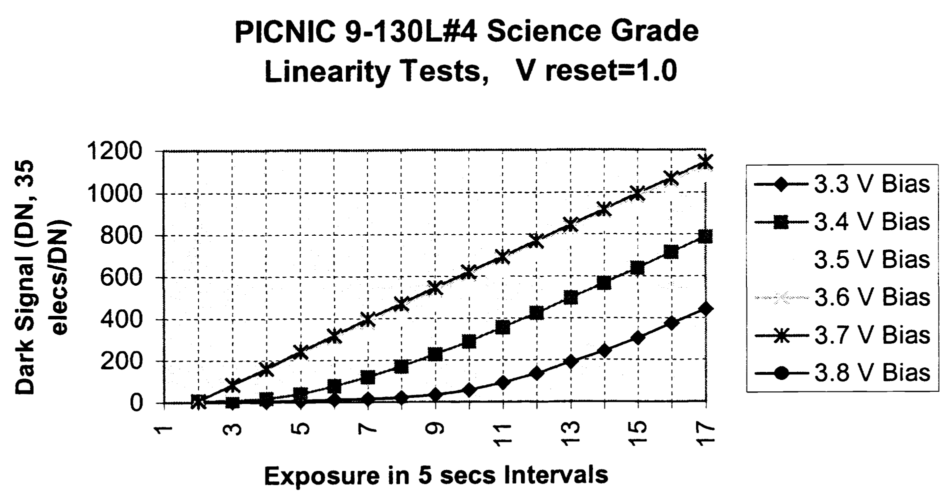

Figure 5: The effect of different output transistor bias levels on the low level non-linearity of a PICNIC device.

The evidence is that careful selection of reset and bias voltages is much more of a problem with the PICNIC devices than with the HAWAII devices. Nevertheless it is important that this aspect of calibration of the devices is always looked at very carefully.

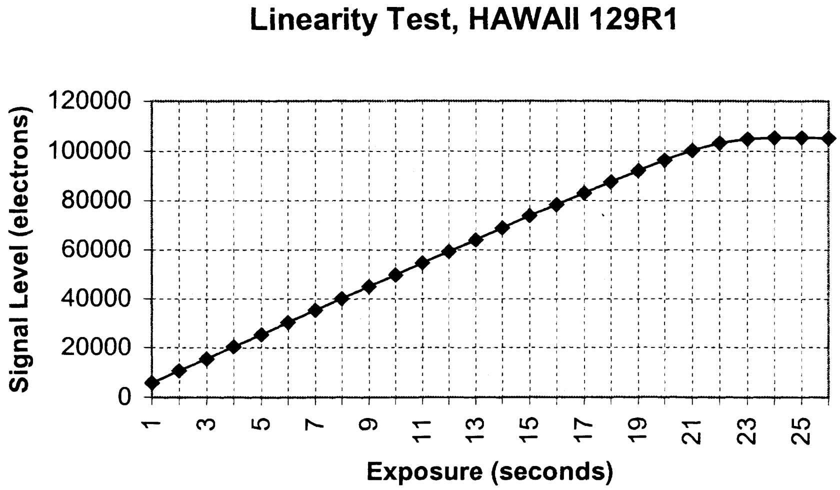

Figure 6: A plot of the linearity of a HAWAII device as a function of time.

At high signal levels the linearity of the device begins to fall away. This happens once the detector capacitor is almost completely discharged. In all devices at the highest level a gradual roll-off is seen. In figure 6 a typical linearity plot is shown for a HAWAII device and in figure 7 the differential linearity is shown for the same device.

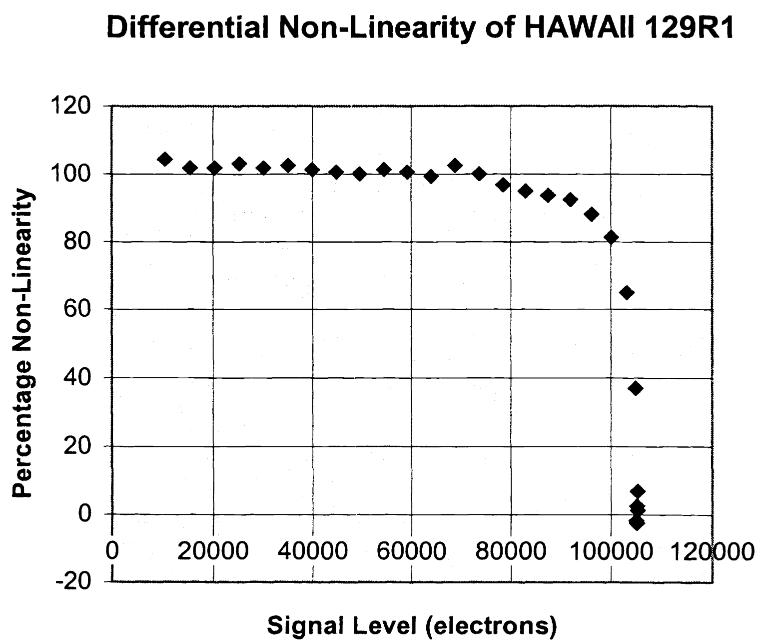

Figure 7: The differential non-linearity as a function of mean signal level for one of the HAWAII devices.

It is clear that linearity is extremely good indeed. Should it be required to use these devices for precision photometry then, depending on the accuracy needed, it may be necessary to do a proper linearity calibration on a pixel by pixel basis. The figures given here are averaged over a large number of pixels, but the description given above of how the basic pixel unit cell works means that the threshold of each pixel is likely to be slightly different and this in turn means that the linearity of each pixel may also be slightly different.

10. RESET PROPERTIES

One of the critical aspect of getting good performance from these X-Y addressable arrays is the whole business of their reset. The architecture of the device is slightly different from the earlier NICMOS devices where each pixel may be reset individually. With the HAWAII and PICNIC devices the reset is done one whole row at a time. With high speed clocking it is possible to reset the whole array very quickly and, indeed, to reset all quadrants of the chip in parallel.

It is important, therefore, to think about the current demands that such a reset operation will make on the electronics which provide the reset supply voltage since this is the voltage to which each of the pixels is being reset. The full well capacity of each pixel (with a reset level of about one volt) is approximately 100,000 electrons. With a million pixels and a reset time of perhaps one microsecond this corresponds to a current that could be as high as many milliamps. It is very difficult indeed to design the supply for the reset voltage not to drop slightly as a consequence of providing this current over a very short period of time. We have used reset periods of from about 0.3 microseconds to 20 microseconds no significant difference in results. Even a small drop in the voltage of the reset supply will give a reset level which is different from that normally used. In practice it is much easier not to try to design an exotic supply for the reset line than rather to accept that multiple resets are necessary and that these resets should be extended over a period that is long enough to allow the power supply to recover and stabilize. Inadequate resetting is not in itself harmful since the slightly incorrect reset level that might be established still subtracts correctly from the image taken at the end of the exposure. However it can produce effects on the raw data such as apparent ramps over the first few rows which look rather odd. It is more important however when working at very low light levels where the knowledge of the performance of the device at these levels, and linearity that can be achieved, may be important.

11. TRAILED PIXELS

Under some circumstances an unusually effect may be seen on the data taken. A small number of bright pixels (usually ones with a very high dark current) seem to be much harder to reset properly than other pixels. If a single reset operation is used these pixels appear to be able to affect subsequent pixels in the same column (see Figure 8), possibly by not allowing the pixel enable gate to be able to turn off properly and thereby leak into the common video line when that column is been addressed, giving slightly different pixel values. It may be that it is a defect in the column addressing transistor of that pixel that is the source of a high dark current. Again even though it looks cosmetically rather unattractive it is not a fundamental problem as all comes out once the reset image subtraction occurs.

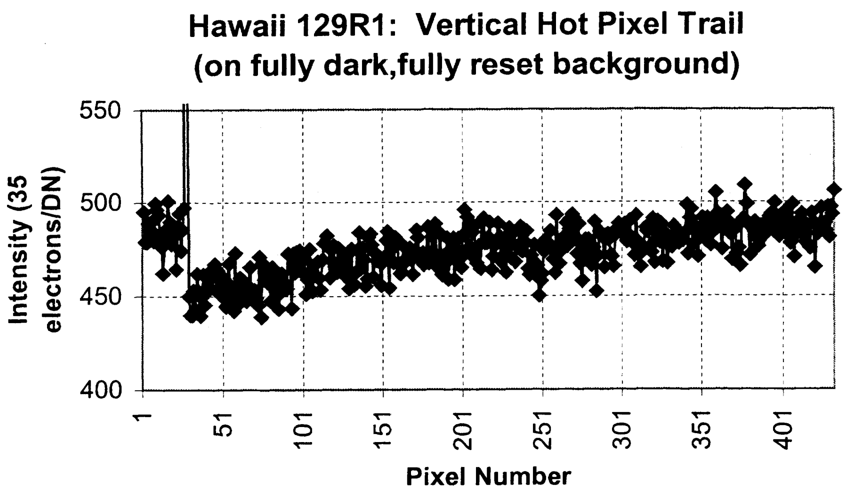

Figure 8: A vertical plot through a bright pixel showing the effect of inadequate resetting in that the corresponding pixel in subsequent rows is reset to a different level from the mean level. Adjacent columns are unaffected by this.

12. CONCLUSIONS

Both the HAWAII and PICNIC devices have been found to be very satisfactory indeed. Many of features described above are really of relatively minor importance and provided they are known about it is easy to manage them as part of overall observing system. With a high detective quantum efficiency, low read-out noise and good full well capacity they offer exceptional performance in the wavelength range from about 750 nm to 2.5 microns. It is remarkable that the quality of performance of these infra-red arrays has advanced as rapidly as it has. In many respects they are equivalent to the performance that is now routinely expected from CCDs and it is clear that their availability is now opening up the near infrared wavelength band for astronomy in a way that has not been possible before.

13. ACKNOWLEDGEMENTS

Many people have contributed to this work. In particular we would like to acknowledge many useful discussions with Kadri Vural, Craig and Scott Cabelli from the Rockwell International Science Center. We would also like to thank Jim Pritchard, Steve Medlen, David Palmer and Peter Weeks for the help with project. Above all however we would like to acknowledge the very generous support and encouragement that we have received from Dr Raymond and Beverly Sackler without whose help none of this work could have been done.

14. REFERENCES

1. L J. Kozlowski, K Vural, D. E. Cooper, R. B. ,D. Q. Bui, D.M. Stephenson, “Status and direction of PACE-I HgCdTe FPAs for astronomy”, SPIE Vol. 1946,pp.148-160 (1993).
2. M. G. Beckett, C. D. Mackay, R. G. McMahon, I. R. Parry, R. S. Ellis, ,”CIRSI: progress with the Cambridge infrared survey instrument”,SPIE Vol. 3354,(1998).
3. K. A. Ennico, I. R. Parry, M., Kenworthy, R. S. Ellis, C. D. Mackay, M. G. Beckett, A. Aragon-Salamanca, K. Glazebrook, J. Brinchmann, J.M. Pritchard, S. Medlen, F. Piche, R. G. McMahon, F. Cortecchia, “The Cambridge OH Suppression Instrument (COHSI)- Status afterfirst commissioning run”, SPIE Vol. 3354,(1998).
4. Mackay, C.D., “CCDs in Astronomy”, Ann. Rev. Ast. And Astrophys., Vol. 24, 255, 1986