

Chapter 3

Hybrid Imager Features

In contrast with CCDs, the nondestructive nature of the HxRG readout makes it possible to do *Up-The-Ramp sampling* (UTR) of the detector, and the multiplexer in the ROIC makes it possible to randomly select and read pixels. With these features we can define at least three distinct ways of operating the detector.

The first, and most conventional way of using the detector, is taking *full frame* exposures – ones in which we clock through every pixel in the detector one or more times. The second is taking *window mode* exposures – ones in which we clock through only the pixels in a sub-region of the detector one or more times. And the third is *guide mode* exposures – ones in which we alternate between full frame reads for the purpose of collecting science data and window mode reads that allow us to reset or read a small window on the detector. The full-frame and guide modes are well suited for long exposures of faint objects and the *window mode* is more appropriate for high speed photometric measurements or telescope guiding. Each will be discussed in turn.

3.1 Full Frame Mode

A diagram that illustrates a *full frame* Up-the-Ramp (UTR) sequence is shown in Figure 3.1. The *ramp*, synonymous with the *exposure*, consists of a series of N *frames*, the basic unit in which the entire array of pixels is clocked through. The clocking scheme is similar to the one employed in CCDs, with a fast direction along the rows controlled by a horizontal clock (HCLK) and a slow direction along the columns controlled by a vertical clock (VCLK). The duration t_f of the *frame* is determined by the pixel time t_p —which corresponds to the time between edges of the HCLK signal and can be as low as 2 microseconds for the native *slow mode* of the HxRG detectors—and the number of outputs N_o used on the multiplexer. Following the nomenclature used for astronomical images, if NAXIS1 is the number of columns and NAXIS2 is the number of rows in the frame, then $t_f = t_p * \text{NAXIS1} * \text{NAXIS2} / N_o + \text{NAXIS2} * t_{rs}$, where t_{rs} is a time associated with moving to the next

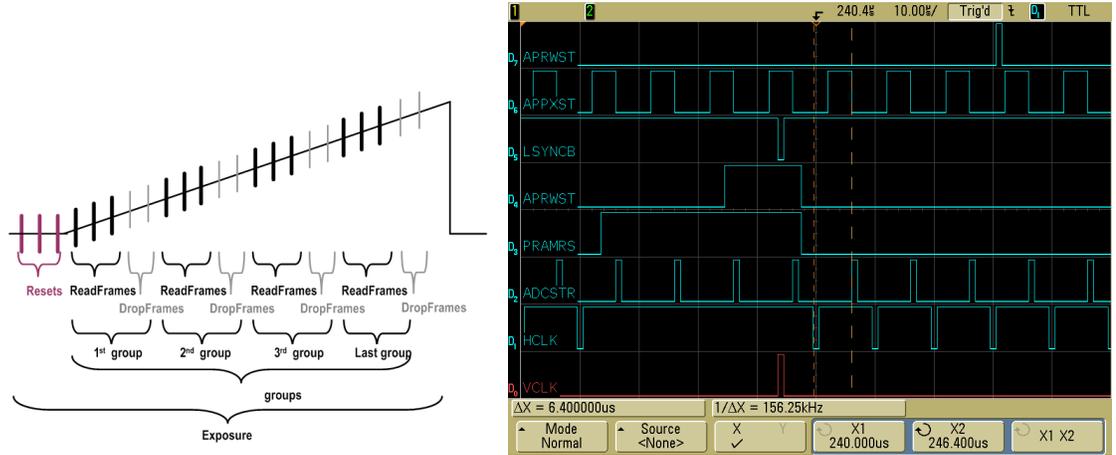


Figure 3.1: (Left) An Illustration of the ramp sequence, courtesy of Markus Loose. The number of resets, drops, groups, and reads are programmable in software, allowing control over the exposure time. (Right) An oscilloscope trace of several clock lines on the SIDECAR ASIC. The delay time t_{rs} is visible.

row of N_{AXIS2}/N_0 pixels in the array. t_f is also the time between reading a pixel in one *frame* and reading that same pixel again in the following *frame*. As a representative example, the SIDECAR ASIC we used had a maximum of 36 inputs, so we chose $N_o = 32$ when using the H4RG HyViSI. We used a $10 \mu s$ pixel time, giving $t_f = 5.24288s + 4096 * t_{rs}$. Eventually, we implemented a counter on the SIDECAR that allowed us to precisely measure a time of $t_f = 5.453 s$ and $t_{rs} = 52 \mu s$. This was also verified on an oscilloscope, the trace of which is shown in Figure 3.1. It is worth noting that t_{rs} was limited by the assembly code used to operate the SIDECAR and not the multiplexer itself. The time for one VCLK and a slight delay before and after to allow the output voltage to settle will place a smaller lower limit on t_{rs} , but this was not tested.

The illustration in Figure 3.1 shows three different types of *frames*: *reset frames*, *drop frames* and *read frames*. During the *reset frames*, the array is clocked so that the gate of the reset FET in the ROIC pixel is enabled and the pixel is held at the voltage V_{RESET} supplied by the control electronics. This allows the charge on the capacitive node in the silicon to dissipate so that it is ready for a subsequent integration. In the *read frames*, the column select and row enable signals for all of the pixels are enabled in turn, bringing the voltage at a pixel to its corresponding output. This voltage is converted to a digital number and sent out to the Data Acquisition System (DAQ) for recording. The unique feature of the *drop frames* is that no data is output from the control electronics to the computer. This allows the user to take a very long exposure without dealing with overwhelming amounts of data. It should be mentioned that data from the *reset frames* can be output to the external DAQ and the pixels can be clocked during the *drop frames*. In fact, as will be discussed later, clocking the pixels during *drop frames* is absolutely essential if one wishes to avoid

temporal drifts and instabilities in the pixel voltages.

3.1.1 Up the Ramp Terminology

In a fashion similar to the scheme used in the NICMOS and JWST specifications [59], we refer to a ramp having a **cadence** of N_{rs} *reset frames*, N_d *drop frames*, N_g *group frames*, and N_{rd} *read frames*, as a N_{rs} - N_{rd} - N_d - N_g ramp. Such a ramp has a total of $N_g \times N_{rd}$ data frames. With these definitions, the exposure time of a ramp in *full frame* is given by

$$t_e = N_g * (N_d + N_{rd}) * t_f, \quad (3.1)$$

with t_f being the frame time given above. To avoid confusion in the following sections, we will use this syntax and reserve the terms *ramp* or *exposure* for such a sequence. We will use the terms *read* or *frame* to describe a *read frame* in these ramps. And the term *image* will be used to refer to the data produced after reduction of these *ramps*, which is described Section 5.2.

3.2 Window Mode

The HxRGs contains a serial register that can be configured to allow random access to a contiguous $X \times Y$ subset of pixels, or window, through one output of the detector. X and Y , the number of columns and rows in the window, respectively, must be greater than 1 pixel and less than NAXIS1 pixels. Programming the addresses of the window limits, X_{start} , X_{stop} , Y_{start} , and Y_{stop} , can be done on the fly. This allows multiple windowed regions of the detector to be read out in a ping-pong like fashion and also the possibility of feedback control loops in which the window coordinates are continually adjusted to track an object. Other applications of window mode, such as telescope focusing, are presented in Simms et al. [39] and more will be discussed in Chapter 5.

Our basic unit in a *window mode* observing sequence is a correlated double sample (CDS) in which we clocked through all $X * Y$ pixels in a window 3 times, resetting each pixel once and then reading the pixel voltages twice.¹ The clocking is similar to full frame mode, with the fast clocking along the rows and the slow clock along the columns. With a pixel time of t_p , one row takes $X * t_p$ seconds with an additional t_{rs} seconds of overhead in shifting to the next row. To complete all rows thus requires $t_w = Y * (X * t_p + t_{rs})$ seconds and this is the total time for a reset or a read in the window sequence. An effective exposure time for a CDS in *window mode* is given by

$$t_e = 2t_w + 2t_{rs}, \quad (3.2)$$

¹Note that there is nothing preventing one from obtaining UTR data in window mode. UTR windows were used for a variety of purposes as well, but the CDS sequence is more relevant for applications like telescope guiding where the pixels are *shot-noise* limited (see Section 4.4.2.2).

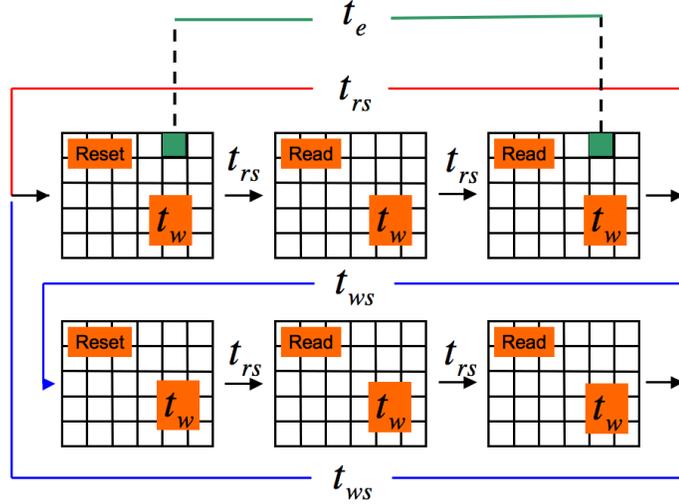


Figure 3.2: An illustration showing the sequencing and times (described in text) involved with *window mode*. The red line at top shows the path taken in *single window mode* and the blue line shows that taken in *multiple window mode* when $N = 2$ windows are used.

which is the time between the reset of a pixel and sampling it in the second read. For our observations, we used a pixel time of $t_p = 10 \mu\text{s}$ and minimized the overhead to attain $t_{rs} = 18 \mu\text{s}$. t_p can be decreased in order to decrease the time to complete one CDS sequence and up the sampling rate with a penalty in noise.

After a CDS sequence, either another CDS is repeated on the same window or the serial register is programmed with a new set of window coordinates. We refer to the former as *single window mode* and the latter as *multiple window mode*. With the microcode used on the SIDECAR ASIC, the operation of writing new window coordinates for X_{start} , X_{stop} , Y_{start} , and Y_{stop} takes $t_{ws} = 150 \mu\text{s}$; adjusting only two of these four would take $\sim 75 \mu\text{s}$. A diagram showing the times associated with window mode is shown in Figure 3.2. It should be apparent from the diagram that if N windows are used in *multiple window mode*, the time it takes to complete a full cycle and return to the first window is given by:

$$t_c^{mw} = 3Nt_w + 2Nt_{rs} + Nt_{ws} \quad (3.3)$$

And if *single window mode* is used this time becomes:

$$t_c^{sw} = 3(t_w + t_{rs}) \quad (3.4)$$

It is also important to know the time between exposures in different windows for the purposes of measuring temporal correlations. In *multiple window mode*, the time it takes between a CDS in

window n and a CDS in window m is given by

$$t_{nw} = 3(m - n)t_w + 2(m - n)t_{rs} + (m - n)t_{ws} \quad (3.5)$$

The full cycle is repeated M times in one full observing sequence, giving a total time of $M * t_c^{mw}$ in *multiple window* mode and $M * t_c^{sw}$ in *single window* mode.

3.2.1 Variation of Multiple Window Readout Sequences

In the simple implementation just described, the flux measurements of individual windows are made serially. A given window is reset and a series of reads is performed before moving onto the next window and repeating. We refer to this method as a **Staggered Reset-Read (SRR)**.² Another type of sequencing is to reset each window in sequence and then read them in sequence a number of times afterwards so that the exposure time is given by an expression similar to Equation 3.3 instead of Equation 3.2. This method we call **Staggered Reset-Staggered Read (SRSR)**. The SRR sequence yields a serial set of flux measurements in time, but offers the benefit that very bright stars are read immediately after reset to prevent saturation of the pixels. The SRSR sequence offers a more parallel set of flux measurements in time—the windows are integrating the same wavefront for some portion of the sequence—and is useful for dim stars where the signal to noise can be improved with the longer integration. There are, of course, other possibilities, one of which is interleaving the SRR and SRSR sequences. For our measurements, only these two were implemented.

3.3 Guide Mode

The HxRG multiplexers can be operated in a special “guide mode” in which the *full frame mode* and *window mode* sequences are interleaved. In this mode, a subwindow of the array that contains a guide star or set of saturated pixels can be continuously read or reset while the rest of pixels in the array integrate charge. If the former is the case, the centroid of the guide star can be calculated in the control electronics or DAQ for values of Δx and Δy that can be fed to a Telescope Control System (TCS) in order to make adjustments in the telescope pointing. At any time during the integration the full array can be read as well, allowing for UTR sampling. Essentially, the guide mode is just an alternation between full frame and window mode. Further discussion of guide mode will be saved for Section 5.5.

²The name **Staggered Reset-Read** implies only one read of the window after it is reset. However, many reads of the window can be performed before moving to the next window. In fact, at least two should be taken to allow for a correlated double sample. Using one **read** is only for nomenclatural convenience.

3.4 Reference Pixels

The H1RG, H2RG, and H4RG all contain a ring of reference pixels that surround the science pixels. There are 4 columns of reference pixels on the left and right hand side of the array and 4 rows of reference pixels on the top and bottom sides of the array. Ideally, the bottom and top reference rows can be used to remove offsets common to an entire output of the detector and the left and right reference columns can be used to remove noise common to a particular row since they are being read only fractions of a millisecond before the science pixels in that same row.

Figure 3.3 shows the average signal of 3 reference pixels in one row and the signal of a science pixel in that same row recorded during a dark exposure with H1RG-022. Both signals show an initial droop during the first 40 reads or so: an effect that is due to some type of electrical settling immediately following a reset of the pixels. This effect is similar to the reset anomaly, seen elsewhere in H2RG hybrid detectors [59]. The size of the drop can be alleviated with tuning of the biases at a given temperature, but if it is present in data that was collected, the reference pixels provide a means of removing it.

As seen in Figure 3.3, the peak to peak amplitude of the reference pixels is about 4 times greater than that of the science pixel. This is most likely due to a difference in capacitance between the two types of pixels. It is noted in the specifications that parasitic capacitance is present on the reference pixels, increasing their capacitance, C_{ref} , by about 50% relative to the capacitance of the science pixels, C_{sci} . From the plot, it appears that the ratio of the two might be greater for this particular device. Figure 3.4 shows a subtraction of the average of 3 reference pixels in a row from two different science pixels in that same row. The signal in the initial reads is seen to increase quickly due to the

Figure 3.3: Signal vs. time in an UTR exposure for the average of 3 reference pixels in row 220 (black) and 1 science pixel (red) in the same row. The first sample has been subtracted so both ramps start from zero. Note the initial dip in signal for the first 40 reads. Since this exposure was taken in the dark, after this drop, the pixels are integrating only dark current at a very slow rate.

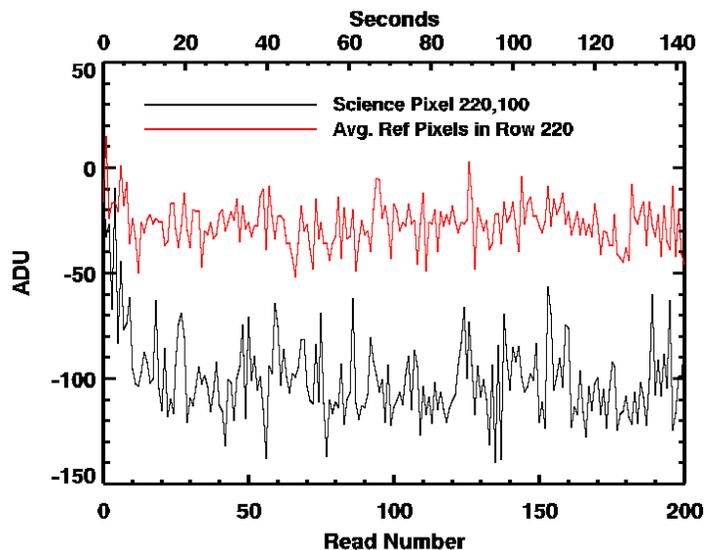
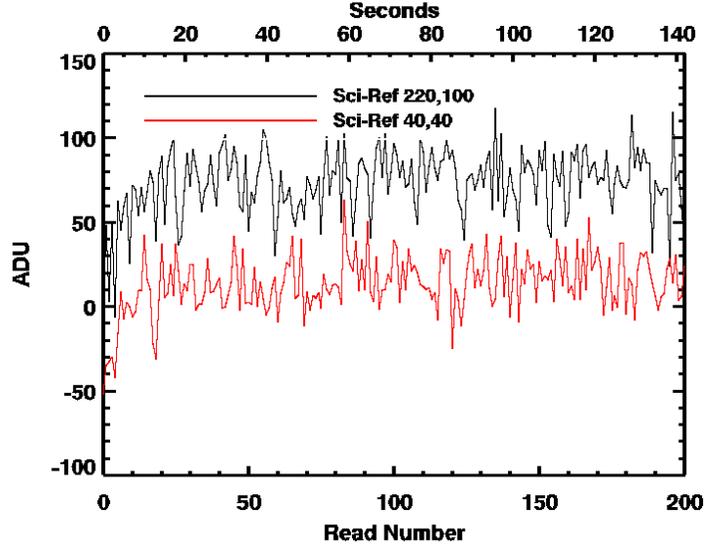


Figure 3.4: Difference in signal between science pixels and the average of 3 of the reference pixels in that same row vs. time in a UTR exposure. Because of a mismatch in capacitance, the signals do not subtract evenly. Instead of an initial droop, the subtraction makes the science pixels appear to rise in signal.



difference in amplitude between the science pixels and reference pixels.

To account for the discrepancy between the capacitance of the science and reference pixels, the coefficient C_{FAC} is introduced in the relation

$$C_{sci} = C_{FAC} * C_{ref}. \quad (3.6)$$

With the assumption that the sum of the dark and photo-generated signal in the science pixels, $S_{sci}(t)$, should be linear with respect to time, we seek to find the value of C_{FAC} that maximizes the linearity of the ideal signal, $S(t)$, with the following relation:

$$S(t, i, j) = S_{sci}(t, i, j) - C_{FAC} * S_{ref}(t, j), \quad (3.7)$$

where i and j are the column and row of the pixel, respectively, and $S_{ref}(t, j)$ is some estimate of the signal in the reference pixels in the row j . $S(t)$ is the signal we would expect to see in the absence of electrical noise in the detector and control electronics. It is the sum of the dark-generated signal, $D(t)$, and the photo-generated signal, $I(t)$.

To find C_{FAC} , we use exposures with low illumination so that the electrical noise is not dominated by shot noise and choose a region of pixels on the detector that contains a minimal amount of defects. A value of C_{FAC} is chosen and for that value a line is fit to the points given by Equation 3.7 for all of the pixels in the region. For $S_{ref}(t, j)$ the average of the 4 reference pixels on the left side of the detector and the average of the 4 reference pixels on the right side are taken separately across the rows and then smoothed with a Savitsky-Golay filter to eliminate spikes in the noise. The left

average is used for the pixels on the left half of the detector and the right average for the right half. The error of the fits (see section 5.2.4 for how this was calculated) for all of the pixels is then averaged, and this process is repeated for a range of C_{FAC} values until the minimum is found.

The results for the HyViSIs H1RG-022 and H2RG-32-147 are shown in Figure 3.5 and listed in Table 3.1. Measurements made on other regions of these detectors show that C_{FAC} does not vary significantly with pixel location, but does depend on the bias voltages V_{RESET} and D_{SUB} . The value was stable over multiple nights of using the detectors with these voltages held constant.

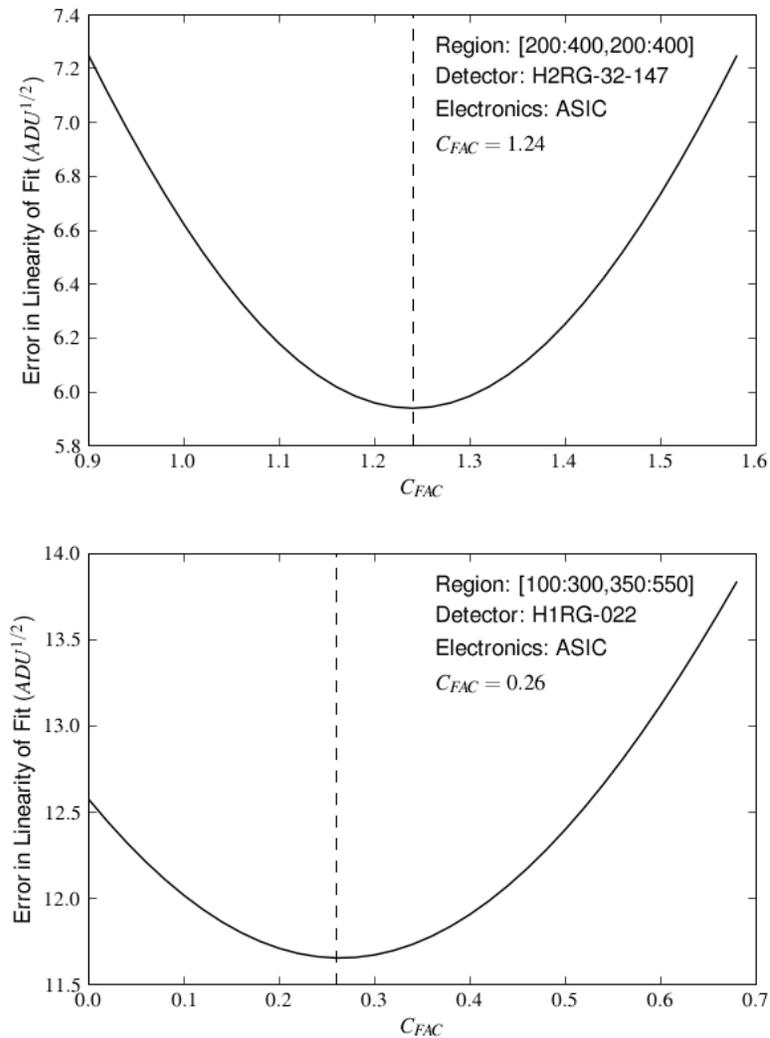


Figure 3.5: The error in the fit of a straight line to the data points obtained with Equation 3.7 vs. the capacitance factor C_{FAC} . The pixels, detector, and electronics used are indicated on the figures.

Table 3.1: C_{FAC} values along with the detector biases V_{RESET} and D_{SUB} . The average error in the linear fit of Equation 3.7 is also indicated.

Detector	Electronics	D_{SUB} (V)	V_{RESET} (V)	C_{FAC}	Error (\sqrt{ADU})
H1RG-022	SIDECAR	0.375	0.094	0.26	5.97
H2RG-32-147	SIDECAR	0.305	0.300	1.24	11.65

Figure 3.6 shows the same subtraction as the one in Figure 3.4 except that the reference pixel average has been multiplied by the factor $C_{FAC} = 0.27$. No initial dip or rise is seen with the factor taken into account. Rather, the slope of the expected dark current signal is well matched throughout the ramps indicating that an adequate value of C_{FAC} has been applied.

This treatment was not applied to H4RG-007 because the dark current shot noise was dominant over the electrical noise. Also, it should be noted that this method was not applied to exposures taken with the ARC electronics. The reason for this is that crosstalk on the video boards causes a large coupling between the signal measured in the reference pixels and the science pixels.

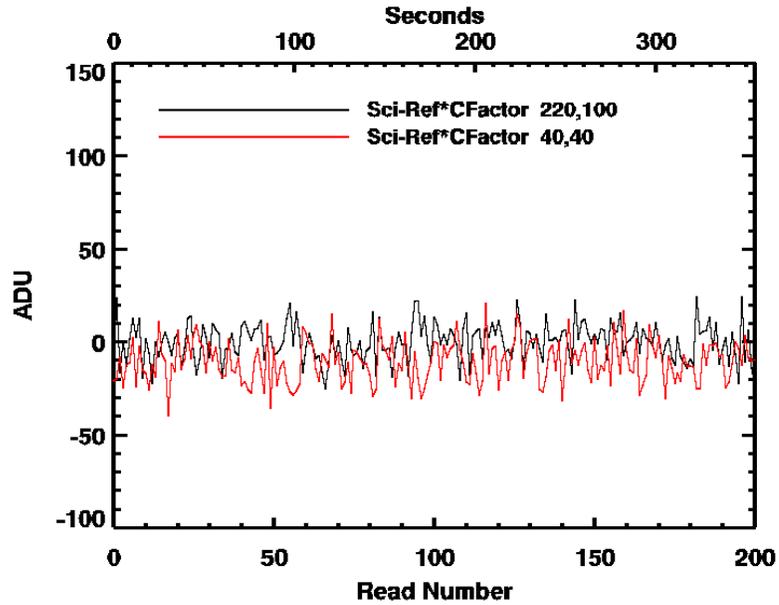


Figure 3.6: Difference in signal between science pixels and the average of 3 of the reference pixels in that same row multiplied by a factor C_{FAC} that takes into account the difference in capacitance between the two types of pixels.

3.5 Readout Electronics: SIDECAR ASIC

Interfacing focal plane arrays (FPAs) to a data acquisition system (DAQ) in astronomy usually requires a middle-man set of electronics that are referred to as the “readout” or “control” electronics. The readout electronics handle tasks such as converting the analog video signals to digital numbers, filtering noise from the signals, providing power or bias voltages to the FPA, and generating the clocking signals necessary to take an exposure with the array. In many cases this set of electronics is equipped with a microprocessor or microcontroller that can store a set of instructions in its internal memory. Different sets of instructions can be loaded based upon the desired mode of operation for the FPA (i.e. binning pixels, reading a subset of the array, etc.). And with a given set of instructions loaded, writing individual registers allows for fine tuning of parameters such as the frame rate and exposure time.

In most cases these readout electronic systems are rather bulky. Large racks are often needed to hold them and the power supplies that they need to operate. In some cases the readout electronics draw enough electrical current from the supplies to necessitate a fan to prevent overheating. They typically consist of multiple circuit boards with discrete chips for each function, i.e. voltage regulators, DACs, ADCs, memory etc. Often times, to make slight adjustments, additional components such as resistors or capacitors must be inserted or soldered to the boards. Since these electronics are usually bolted to a telescope or flying on a satellite in astronomical applications, weight and size can be an issue, and certainly one would rather avoid swapping out components.

To step away from the bulkiness and large power requirement of traditional electronics, Teledyne Scientific has produced a multi-purpose control Application-Specific Integrated Circuit (ASIC) called the SIDECAR (**S**ystem for **I**mage **D**igitization, **E**nhancement, **C**ontrol **A**nd **R**etrieval). A block diagram of the SIDECAR ASIC and photographs of the 22 x 14.5 mm² die mounted in two different packages is shown in Figure 3.7. As the diagram indicates, the chip contains all of the functionality needed to control and readout a detector: clocks, biases, ADCs, etc. And in addition to the functions shown, the chip provides pre-amplification and amplification stages as well as array processors that permit data processing function such as co-adding channels or subtracting offsets.

In addition to the compact science mission packaging schemes shown in Figure 3.7, Teledyne offers the SIDECAR in a development kit intended for laboratory use. The development kit, which was used for most measurements included in this thesis, is shown in Figure 3.8. In this configuration, the ASIC is placed on a board that has numerous test points that allow one to probe its various output signals, and LEDs that show clock activity. This board connects to a smaller board called the JADE card that handles communication with a Windows PC via a USB interface. The development kit as a whole is essentially a plug and play device. The user only needs to write the assembly code instructions for the microcontroller in the ASIC and the software that the PC will use to extract the data through the USB bus. For the latter, a library of Microsoft COM functions is provided so that typical astronomy applications such as IDL can be used for data retrieval and configuration.

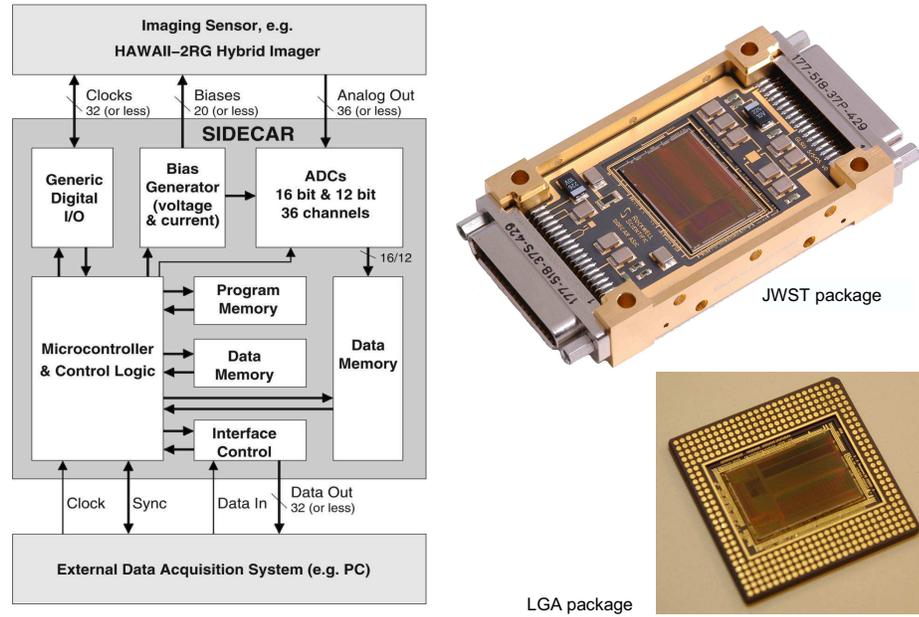


Figure 3.7: A block diagram of the SIDECAR ASIC along with two packaging options available for the chip (taken from Ref. [60]).

At the time of writing, Teledyne offers the SIDECAR Development Kit in two flavors: warm and cold. The warm kit operates at room temperature, and so it must be placed outside the cryogenic system enclosing the detector to which it is attached. This implies that there is a significant length of cable over which the analog signals must travel between the ASIC and the detector. The cold, or *cryogenic*, kit can operate at cryogenic temperatures, which means that it can be placed directly beside the detector so that only digital signals going to and from the DAQ must travel over long distances. This has the benefit of reducing noise pickup during the analog transmission and the lower temperature of the chip inherently reduces the noise and leakage currents in the ASIC itself. Both warm and cryogenic kits were tested for this thesis work, but in this section only the warm kit is discussed. Results for the cryogenic kit are included in later chapters.

In the following subsections we will not provide a comprehensive description of the chip; only a brief review of some of the features and its performance in the various modes of operation with the HyViSI devices. For a detailed description of the SIDECAR, the reader is referred to reference [60].

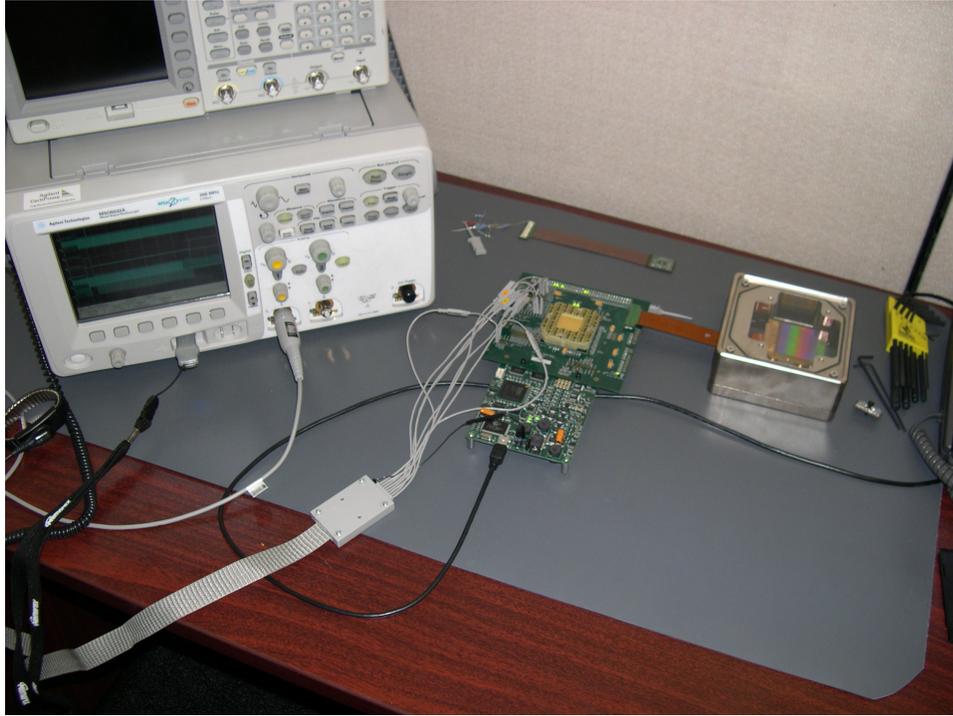


Figure 3.8: A photograph of the SIDECAR development kit attached to an H4RG multiplexer and a USB cable going to a laptop computer not shown. Multiple clocks and analog signals are being traced on the oscilloscope.

3.5.1 Pre-Amplification Stage

Before analog-to-digital conversion takes place, the video signals from the detector are first fed into the SIDECAR pre-amps. In the pre-amp stage, signals can be routed to different channels or shared among them via an internal mux, offsets can be added to them with a DAC, and they can be amplified and filtered. The amplification sub-block of the pre-amps, displayed in Figure 3.9, shows that capacitive feedback is used for gain selection with the capacitors C_{FB} and the inputs are capacitively coupled to the amplifier through C_{IN} . Capacitive feedback has the advantage that it does not lower gain since it does not trade gain for bandwidth [61].

During the amplification stage, the switches $S3$ and $S6$ are open. This leaves the nodes of the capacitors C_{FB} floating at a potential set by a reset transistor switch. As mentioned in Ref. [60], inevitable leakage currents in the silicon will cause these nodes to drift with a time constant that depends highly on temperature. When the SIDECAR chip is placed inside the dewar and cooled along with the detector, the leakage currents are very small and this does not present a problem. However, when the SIDECAR is held at room temperature, the pre-amp drifts are very noticeable after fractions of a second. This drift and the noise associated with resetting the capacitors take

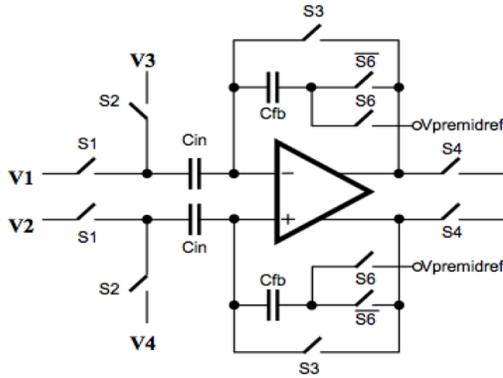


Figure 3.9: A schematic of the amplification stage for a given channel in the SIDECAR. Note that capacitors are used for feedback rather than resistors.

different forms based upon the implementation of the signals available in the pre-amp stage.

Pre-amp reset once per frame From the standpoint of assembly coding, an easy-to-implement conversion sequence is to reset C_{FB} once per frame. However, for frame times as short as 0.25 seconds, the leakage currents can cause the output voltage going into the ADCs to swing by as much as 25 mV. This drift over the frame, which will continue for longer frame times, is shown in the top of Figure 3.10. It is evident then that more frequent resets of the pre-amp capacitors are needed at room temperature.

Pre-amp reset once per row: No kTC Removal An alternative to resetting once per frame is to reset before every row conversion. The row may consist of the number of pixels in a row per output of the detector or the number of pixels in the row of a sub-window. The problem with this is that after each reset, a random amount of charge will be left on C_{FB} . The noise associated with this random charge should be proportional to kTC_{FB} , where k is Boltzmann's constant and T is the temperature. The pattern associated with this noise, shown in the middle frame of Figure 3.10, consists of horizontal bands across the frame. And since it is uncorrelated from frame to frame, CDS subtraction of two consecutive frames will boost this noise by $\sqrt{2}$.

Pre-amp reset once per row: With kTC Removal In order to deal with the horizontal banding caused by kTC noise, an intrinsic analog CDS scheme is used in the pre-amp and ADC blocks. In each ADC conversion, both the video signal from the pre-amp and the voltage on the capacitors are sampled. The ADC digitizes the sampled signal using the latter as a reference so that the kTC offset subtracts out. The last frame in Figure 3.10 shows that the noise is essentially white when this method is used. The implementation of the kTC removal scheme involves toggling one signal in the pre-amp and slightly changing the internal bias voltages from those used in the normal reset schemes; it does not require additional time for conversions.

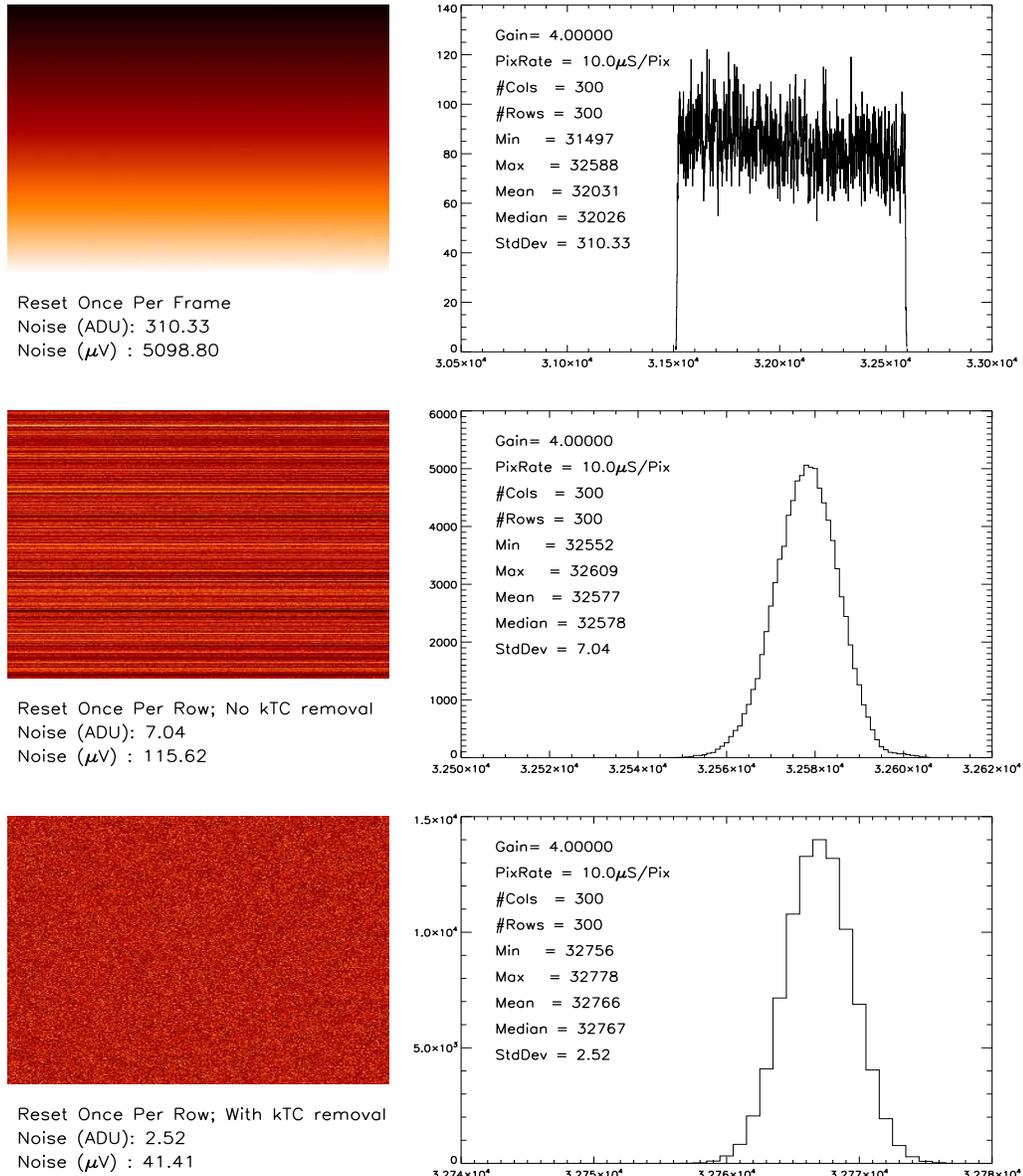


Figure 3.10: Different reset schemes for the SIDECAR pre-amplifiers. The methods used to obtain each are described in the text. All measurements were made with the input voltages $V1-V4$ connected to the internal ground.

3.5.2 Conversion Gain

In all measurements described in this thesis, the 16 bit Successive Approximation (SAR) ADCs were used. For the SAR ADC, the conversion from microvolts to ADU is given approximately as

$$\text{Counts(ADU)} = 32768 * \frac{V_{In} - V_{Ref}}{VRP - VRN} * G + 32768, \quad (3.8)$$

where V_{In} is the input voltage to the pre-amp, V_{Ref} is one of the selectable reference voltages provided on the chip, VRP and VRN are ADC reference voltages, and G is the gain of the pre-amp. G is configurable in 3 dB increments from -3 dB to 27 dB. V_{In} and V_{Ref} correspond to the input voltages $V1$ and $V2$ in Figure 3.9. The selection $V_{REF} = V_{REFMAIN}$ was used for all measurements in this section aside from the ones where both inputs were set to ground. As will be discussed, when the detector outputs are being sampled, $V_{REFMAIN}$ results in a much higher noise than that obtained by using the signal V_{REFOUT} from the detector. Routing V_{REFOUT} from the detector to $InPCommon$ or $InNCommon$ on the SIDECAR is absolutely essential for low noise performance.

To obtain a conversion from voltage to ADU and measure the actual values of G for our configuration, a voltage was supplied to the pre-amp inputs with an Agilent E3647A Dual Output 30 volt power supply. This voltage was increased from 0-3.3 V by 0.1 V increments and at each increment a set of 20,000 digitizations were recorded. The mean of these digitizations were taken to yield an average ADU value at that voltage. The conversion from μV to ADU at a gain of 1 is shown in Figure 3.11. The conversion measured at other gains is listed in Table 3.2.

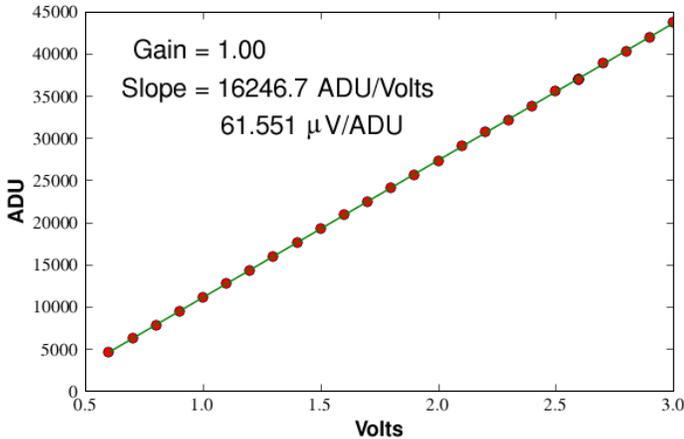


Figure 3.11: Average digital ADU values (shown as red dots) vs. the input voltage at which they were measured. The inverse slope of the best-fit line (shown in green) to the points yields the conversion in $\mu V/ADU$.

Table 3.2 shows that as the gain setting in the pre-amp is increased, the measured gain moves farther from the expected value, with the former being less than the latter. This is due to the use of low noise resistors and single-stage buffers in the system (low noise settings result in diminished

gain). However, the measured values are very repeatable, so with proper calibration this discrepancy can be taken into account and should not be a performance-limiting issue.

The last two columns in the table labeled *Ground Noise* show the digitization noise when the pre-amps are configured to measure the internal ground signal on the chip. In agreement with Figure 5 in Loose et al. [60], the noise is ADC limited below a $G = 4$ and above that it is limited by the pre-amp. The two columns labeled *Supply Noise* contain the digitization noise measured when reading the voltage from the Agilent power supply. The RMS output noise of the power supply was measured with an oscilloscope to be about $340 \mu\text{V}$, so this asymptotic behavior of the noise with increasing gain is to be expected.

Table 3.2: Measured pre-amp gains for the SIDECAR ASIC.

Gain	Gain	Measured	Measured	Supply	Supply	Ground	Ground
	(dB)	Gain	Gain	Noise	Noise	Noise	Noise
		($\mu\text{V}/\text{ADU}$)		(ADU)	(μV)	(ADU)	(μV)
0.71	-3.01	86.07	0.71	4.79	412.3	2.70	232.4
1.00	0.00	61.53	1.00	6.81	383.9	2.72	166.7
1.41	3.01	43.56	1.41	8.35	363.7	2.78	121.1
2.00	6.02	30.77	2.00	11.67	359.1	2.80	86.2
2.83	9.03	22.96	2.83	16.22	352.4	2.75	59.7
4.00	12.04	15.46	3.97	22.52	348.2	2.81	43.4
5.65	15.05	11.01	5.58	31.10	342.4	2.97	32.7
8.00	18.06	8.01	7.67	43.03	344.6	3.22	25.8
11.31	21.07	5.75	10.71	59.90	344.2	3.62	20.8
16.00	24.08	4.15	14.84	83.82	347.2	4.29	17.8
22.62	27.09	3.00	20.48	115.01	345.9	5.25	15.8

3.5.3 Averaging Multiple Channels

The input routing multiplexer of the pre-amp and math capabilities of the array processor (AP) allow one input signal to be shared and digitized on multiple channels and then averaged before it is written to the dual port memory and read out by the DAQ. This might be advantageous if buffer size or memory overflow is an issue in the readout system. And this feature is particularly useful for the HxRG multiplexers as the number of outputs is configurable. For instance, the detector can be run in four output mode, with each output being sampled and averaged between eight channels on the SIDECAR. And in window mode only one output of the detector is used, so there is no reason not to take advantage of multiple channels on the SIDECAR other than power consumption.

Each math operation and read/write in the array processors requires at least one clock cycle.

One would expect that doing math between A/D conversions would therefore result in an overall slower pixel conversion time. However, the array processor clock can be configured to run at a faster rate than the ADC clock so that no time is lost between successive A/D conversions.³ Example clock rates and conversion times are given in Table 3.3. The redundancy in the table is meant to illustrate that no decrease in pixel rate is suffered in the channel averaging process.

The basic process for averaging N channels is as follows: The first channel writes its value to dual port memory. After a certain delay, the second channel reads this value from memory, adds its own A/D value to it, and writes it back to the same address. This process continues with the N channels until all have been coadded. Then a bit shift (or a multiplication followed by a bit shift for the case of $N = 6$) is used to achieve the averaging. Finally, the value is written back to memory and stored until it is extracted by the DAQ.

Table 3.3: SIDECAR clocking scheme for averaging multiple channels.

Channels	System Clock (MHz)	AP Clock (MHz)	ADC Clock (MHz)	$\frac{\text{AP Cycles}}{\text{Pixel}}$ (Used/Total)	$\frac{\text{ADC Cycles}}{\text{Pixel}}$ (Total)	Pixel Time (μs)
1	10	1.00	1.00	2/10	10	10
2	10	1.00	1.00	5/10	10	10
4	10	5.00	1.00	13/50	10	10
6	10	5.00	1.00	17/50	10	10
8	10	5.00	1.00	24/50	10	10

Figure 3.12 shows the reduction in noise when multiple channels are averaged. The decrease goes nearly as the theoretically predicted $1/\sqrt{N}$ channels, indicating that the noise in the ADCs is uncorrelated.

Application to H1RG HyViSI Detector As an example of the usefulness of averaging multiple channels and selecting the proper conversion gain on the SIDECAR, temporal read noise measurements were made on the 1024x1024 Hybrid Visible Silicon (HyViSI) detector H1RG-022. To make these measurements we ran the detector in window mode and obtained a set of correlated double sample (CDS or *reset-read-read*) frames. The window selected included the first 100 rows and 100 columns of pixels. The first 4 rows and first 4 columns of the window were reference pixels that have a capacitance of roughly $C_{Ref} = 35 \text{ fF}$ and all of the others were science pixels that have a capacitance of $C_{Sci} = 14 \text{ fF}$. After obtaining a stack of CDS frames, we look at the temporal variation of each pixel and measure its standard deviation. Averaging the standard deviations yields the temporal read noise.

³The only caveat is that the ADC clock must be divided relative to the AP clock. Since both clocks are derived from the microcontroller system clock, the ADC clock must run more slowly than the system clock if more than 2 channels are averaged. See Table 3.3 for the relative rates.

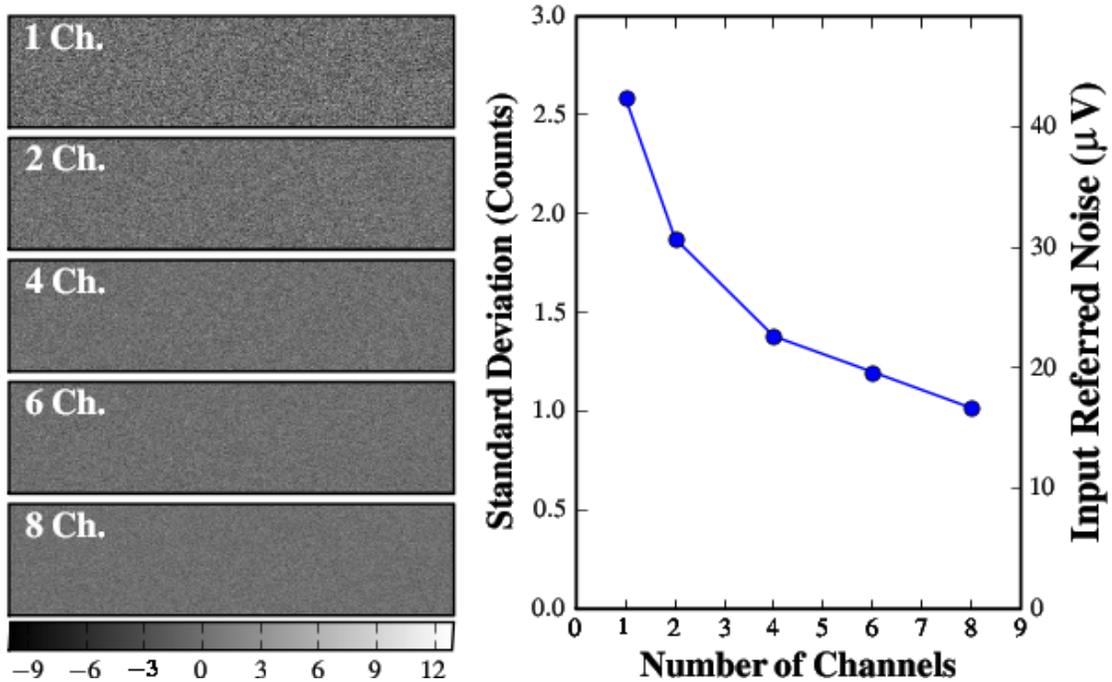


Figure 3.12: (left) Digitizations of noise in the SIDECAR for averaging 1-8 channels at a pre-amp gain of 4. The colorbar shows the ADU values after subtraction of the mean value of 32768. (right) The RMS value of the noise as a function of the number of channels averaged.

As illustrated in Figure 3.13, the difference in pixel capacitances makes a significant difference in the read noise, especially when measured in electrons. In astronomical applications, quantifying the read noise in electrons is the most useful since this is the quantity that can be directly used, along with quantum efficiency, to estimate the number of photons that hit a given pixel. We see that for small capacitances, the larger gain yields a much smaller read noise in electrons. For the eight channel average noise of the reference pixels, both measurements yield a noise of about 3 ADU, indicating that the ADC contribution to the noise might be dominating rather than the contribution from the detector noise.

The lower read noise at higher gain is very desirable, suggesting that the SIDECAR pre-amps should be run at a high gain with HyViSI detectors. However, one must also consider the trade-off between gain and the usable dynamic range of the detector. For instance, in the case of H1RG-018 the well depth is approximately $110,000 e^-$. If the SIDECAR pre-amps are set at a gain of 4 the conversion gain is about $1.4 e^-/\text{ADU}$ and the full ADU range needed to cover the entire well will be 78,500 ADU. However, the 16 bit ADC allows a maximum of 65,536 values, so some portion of the voltage range will be lost. If the entire well depth is to be used, a gain slightly below this must be chosen.

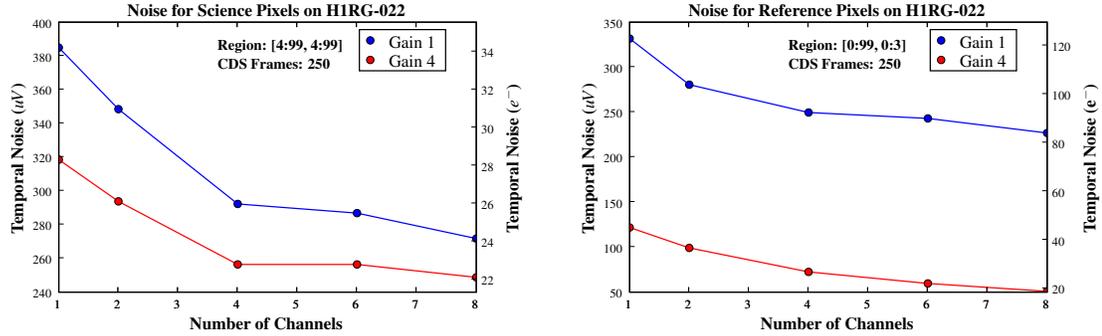


Figure 3.13: (Left) Read noise of HIRG-022 science pixels measured with multiple channels averaged. (Right) Read noise of HIRG-022 reference pixels measured with multiple channels averaged. The capacitance of the science pixels is roughly 2.5 times as large as that of the reference pixels. For this particular measurement, the SIDECAR was configured so that $G = 1$ yielded $64 \mu\text{V}/\text{ADU}$ and $G = 4$ yielded $16 \mu\text{V}/\text{ADU}$.

3.5.4 Noise Performance vs. Pixel Time

Among other projects, the SIDECAR is being considered for use in the guider cameras of the **L**arge **S**ynoptic **S**urvey **T**elescope (LSST) and the **S**uper**N**ova **A**cceleration **P**robe (SNAP). An important parameter for both of these systems will be the frame rate for imaging a guide star, which is dependent on the pixel conversion time. For 100 Hz adjustments of the telescope pointing using centroids from multiple guide stars, a pixel conversion time smaller than $10 \mu\text{s}$ may be required. And for science imaging applications, reducing the pixel time has the added benefit of increasing dynamic range since bright stars can be imaged before they saturate the pixels. The performance of the SIDECAR ADC at faster A/D conversion times will thus directly affect both of these applications.

Increasing the pixel conversion rate means that voltages in the system, i.e. the analog output of the detector, the digital clock signals, the DAC voltages used in the SAR ADC, etc., must change more quickly. The relationship $I = C dV/dt$ tells us that changing these voltages more quickly requires increases in electrical current. In general, increasing the speed at which signals are converted to digital output should be accommodated by an increase in drive currents in the system and an attempt to reduce source impedances and load capacitances to account for this. In the case of the SIDECAR, we need to increase the bias currents in the pre-amp and SAR ADC.

For most ground based astronomy missions, boosting the currents is perfectly fine since there is usually plenty of electrical power for the electrical and thermal systems. However, for space based missions it might be desirable to keep the currents and heat dissipation as low as possible since electrical power from solar panels is not an abundant resource. For this reason, the minimum ADC and pre-amp currents needed to operate the SIDECAR at each clock speed were sought. Values for which the noise histogram remains approximately Gaussian in shape up to a gain of $G = 4$ are listed

Table 3.4: Operating parameters and noise performance at different A/D sampling rates on the SIDECAR ASIC. Note that the noise can be kept at a nearly constant level below 167 kHz by increasing the bias currents.

Sampling Rate	100 kHz	125 kHz	167 kHz	250 kHz
$I_{PreAmpBias}$ (μA)	6.40	6.40	16.00	17.60
$I_{PreAmpCasc}$ (μA)	1.60	1.60	1.60	1.60
I_{NBIAS1} (μA)	11.30	14.50	17.70	48.00
I_{NBIAS2} (μA)	11.30	14.50	17.70	48.00
I_{NFB1} (μA)	5.00	5.00	13.00	34.00
I_{NFB2} (μA)	5.00	5.00	13.00	34.00
$I_{VRP DAC}$ (mA)	0.343	0.343	0.343	0.654
$I_{VRN DAC}$ (mA)	60.00	60.00	60.00	60.00
RMS Noise (ADU)	2.65	2.65	2.76	3.13
RMS Noise uV	41.0	41.0	42.7	48.4

in Table 3.4. It should be note that these currents minimized the noise when the SIDECAR inputs were connected to ground. They do not necessarily offer the best performance when reading out the detector.

If the drive currents fall significantly below these values, certain instabilities in the digital output of the SIDECAR will result. A description of these instabilities, along with the current that is most likely lacking is included in the following section. In general, increasing the currents I_{NBIAS1} , I_{NBIAS2} , I_{NFB1} , and I_{NFB2} will decrease the ADC noise (Markus Loose, private communication). However, if brought too high they can also cause instabilities. It should also be mentioned that the values listed in Table 3.4 work for gains less than or equal to $G = 4$. If a higher gain is used, one or more of the currents may need to be increased.

3.5.5 Reference Voltages for SIDECAR when Connected to an HxRG

The most important point to consider when hooking an HxRG detector up to the SIDECAR, by far, is that using the signal V_{REFOUT} from the detector as a reference in A/D conversions offers the best performance. In some versions of the cabling that we used, V_{REFOUT} was routed to one of the upper input channels (IN32) of the SIDECAR. In this configuration, it must be measured against a reference voltage and digitized separately, which only allows it to be used in digital subtraction. In other versions, V_{REFOUT} was routed to $InPCCommon$ or $InNCommon$, where it can be routed to all 36 channels and used as a reference for measurement of the detector outputs. This is preferred since it provides a truly differential measurement and eliminates common mode noise before digitization.

As a measure of the significance of using V_{REFOUT} versus an internal reference voltage on the

SIDECAR, in a particular configuration of the detector, using $V_{REFMAIN}$ will result in 40-50 e^- read noise. When V_{REFOUT} is used instead, with no other changes being made, the read noise reduces to 10 e^- . Care should therefore be taken when designing the cables from the detector to the SIDECAR to ensure that V_{REFOUT} connects to $InPCCommon$ or $InNCommon$.

Additional Notes

- Noise histogram is very sensitive to DAC buffer current of VRP and VRN.
- The most important settings in determining the RMS and the shape of the noise histogram seem to be $I_{PreAmpBias}$. If $I_{PreAmpBias}$ is too low, the noise will be high and at a certain threshold, the histogram will take a pitchfork shape like that in middle panel of Figure 3.14.
- When we increase the pre-amp gains, we should also increase I_{NBIAS1} , I_{NBIAS2} , I_{NFB1} and I_{NFB2} . Otherwise, the noise distribution becomes jagged, indicating a loss of the LSB.
- If I_{NBIAS1} and I_{NBIAS2} are too small, the histogram will appear jagged, with the odd ADU values having smaller values than their even neighbors as in the top panel of Figure 3.14.
- If I_{NFB1} and I_{NFB2} are too small, the distribution will be less strongly peaked. Increasing these currents will make the peak tighter. If these currents are too small, a gap might form in the distribution.
- Large capacitance loads will cause ringing somewhere in the ADC. The resultant histogram will have at least three separate peaks. Increasing the drive currents will bring these peaks closer together and eventually they will merge.
- The DAC buffer current for VRP should be adjusted for different gains in order to keep a nicely peaked histogram.
- If I_{NBIAS1} , I_{NBIAS2} are set too high, speckles will appear in the digitizations, indicating voltage spikes somewhere in the system, as in the bottom panel of Figure 3.14.

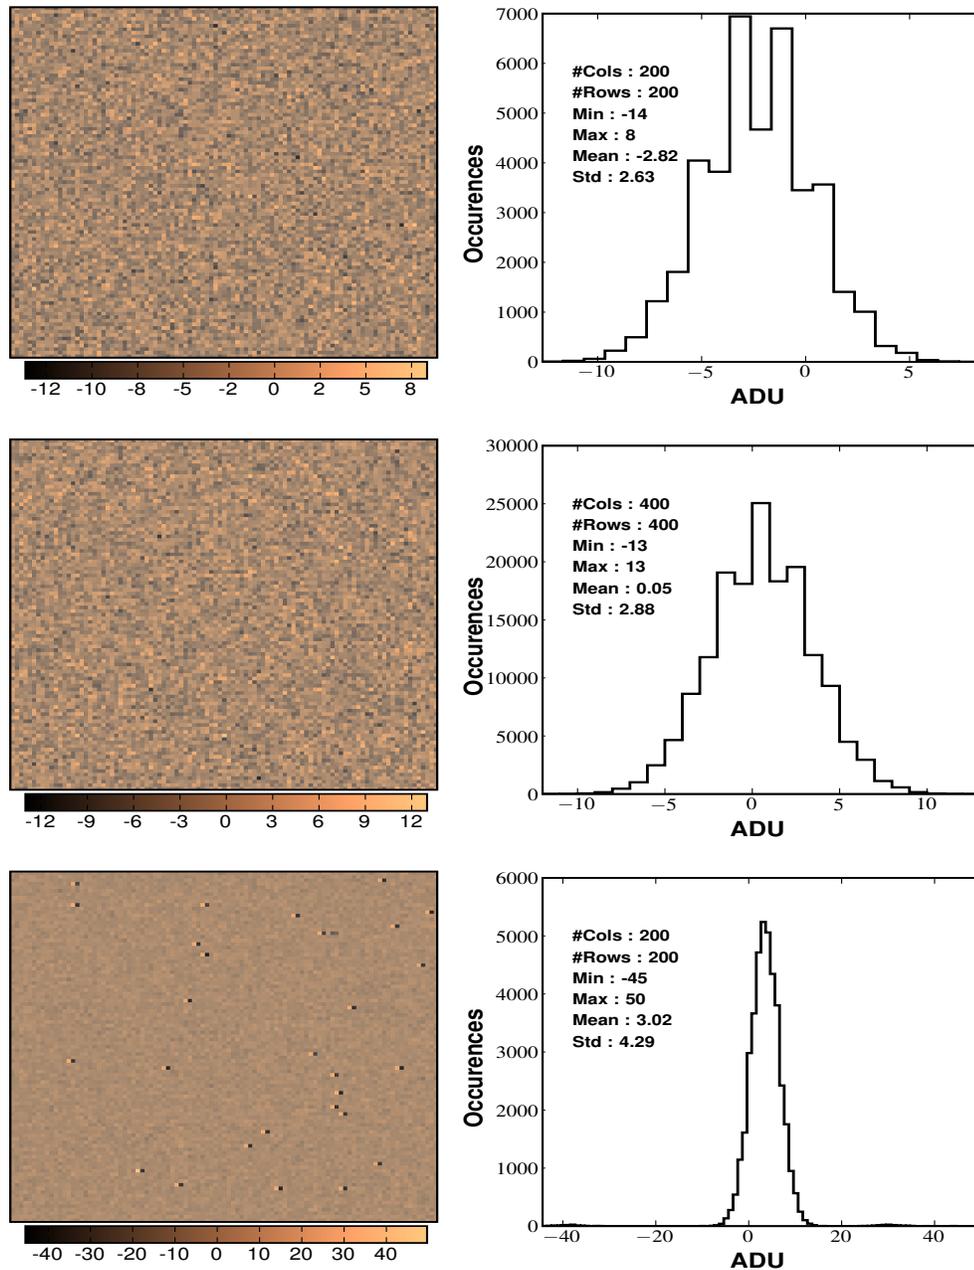


Figure 3.14: (Top) Digitizations taken at 8 μ s when $I_{NBIA1} = I_{NBIA2} = 8.1 \mu$ A were too low. (Middle) Digitizations taken at 6 μ s when $I_{PreAmpBias} = 1.6 \mu$ A was too low. (Bottom) Digitizations taken at 4 μ s when $I_{NBIA1} = I_{NBIA2} = 53 \mu$ A were too high. All digitizations were taken with a pre-amp gain of $G = 4$.