Journal of Astronomical Telescopes, Instruments, and Systems

AstronomicalTelescopes.SPIEDigitalLibrary.org

Overview of the SAPHIRA detector for adaptive optics applications

Sean B. Goebel
Donald N. B. Hall
Olivier Guyon
Eric Warmbier
Shane M. Jacobson

SPIE.

Sean B. Goebel, Donald N. B. Hall, Olivier Guyon, Eric Warmbier, Shane M. Jacobson, "Overview of the SAPHIRA detector for adaptive optics applications," *J. Astron. Telesc. Instrum. Syst.* **4**(2), 026001 (2018), doi: 10.1117/1.JATIS.4.2.026001.

Overview of the SAPHIRA detector for adaptive optics applications

Sean B. Goebel,^{a,b,*} Donald N. B. Hall,^a Olivier Guyon,^{b,c,d,e} Eric Warmbier,^a and Shane M. Jacobson^a

^aUniversity of Hawaii, Institute for Astronomy, Hilo, Hawaii, United States

^bNational Astronomical Observatory of Japan, Subaru Telescope, Hawaii, United States

^cUniversity of Arizona, Steward Observatory, Tucson, Arizona, United States

^dUniversity of Arizona, College of Optical Sciences, Tucson, Arizona, United States

^eAstrobiology Center of NINS, Tokyo, Japan

Abstract. We discuss some of the unique details of the operation and behavior of Leonardo Selex avalanche photodiode for HgCdTe infrared array (SAPHIRA) detectors, particularly in relation to their usage for adaptive optics wavefront sensing. SAPHIRA detectors are 320×256 at $24\text{-}\mu\text{m}$ pixel HgCdTe linear avalanche photodiode arrays and are sensitive to 0.8- to $2.5\text{-}\mu\text{m}$ light. SAPHIRA arrays permit global or line-by-line resets of the entire detector or just subarrays of it, and the order in which pixels are reset and read enables several readout schemes. We discuss three readout modes; the benefits, drawbacks, and noise sources of each; and the observational modes for which each is optimal. We describe the ability of the detector to read subarrays for increased frame rates and, finally, clarify the differences between the avalanche gain (which is user-adjustable) and the charge gain (which is not). © 2018 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: 10.1117/1.JATIS.4.2.026001]

Keywords: Selex avalanche photodiode for HgCdTe infrared array; detector; HgCdTe; APD; adaptive optics; wavefront sensing.

Paper 18007 received Feb. 21, 2018; accepted for publication May 22, 2018; published online Jun. 13, 2018.

1 Introduction

Due to their ability to detect individual photons with both high temporal and spatial resolutions, electron-multiplying CCDs (EMCCDs) have greatly improved the sensitivity limits of adaptive optics (AO) systems. However, EMCCDs are sensitive to optical wavelengths, and in most AO implementations, the science instrument operates at near-infrared wavelengths. This wavelength sensitivity difference between the wavefront sensor and science instrument occurs for two reasons. First, the path length differences introduced by atmospheric turbulence are approximately independent of wavelength, so the aberrations as a fraction of phase are reduced at longer wavelengths. Therefore, better image quality is obtained at longer wavelengths, so the science instruments are designed to take advantage of this. Second, until now high frame rate, low noise, reasonable cost infrared detector arrays have not existed, so it was not feasible to do high-order wavefront sensing at near-infrared wavelengths. The Selex avalanche photodiode for HgCdTe infrared array (SAPHIRA) detector is the first such technology that enables this.

Operating the wavefront sensor at near-infrared wavelengths provides several benefits.¹ It has the potential to expand the sky coverage available for natural guide star sensing because it enables the observation of targets that have minimal emission at optical wavelengths, such as late-type stars or obscured objects. In particular, there is significant interest in observing nearby M dwarfs because they have favorable contrasts for directly imaging reflected-light extrasolar planets located in their habitable zones. Second, if both the science module and wavefront sensor are sensitive to similar wavelengths, they can share a greater fraction of optical elements inside the instrument and thereby

minimize noncommon path errors. For these reasons, SAPHIRA detectors enable exciting new AO observing modes.

SAPHIRA arrays are well-suited for wavefront sensing because they provide two advantages over other infrared devices. First, they have a user-adjustable avalanche gain. This is discussed at length in Sec. 5, but, in short, it amplifies the signal from photons but not read noise sources. At high gains, the signal of an individual photon is greater than the noise. If the noise is divided by the avalanche gain, it can be $<1e^-$. The second advantage of SAPHIRA is that all outputs are still used when reading subarrays, so greatly increased frame rates can be achieved by reading less than the full detector (Sec. 3). This is in contrast to, for example, Teledyne HAWAII detectors in which each output reads a contiguous stripe of the detector and only a single output is used when in subarray mode.

SAPHIRA detectors are 320×256 pixel mercury cadmium telluride arrays and have a $24\text{-}\mu\text{m}$ pixel pitch.² They are manufactured by Leonardo (formerly called Selex) and have 32 outputs. The absorption layer (where photons are intended to be absorbed) has 0.8 to $2.5\text{ }\mu\text{m}$ sensitivity and is transparent at longer wavelengths. These longer wavelengths must be filtered out to avoid spurious signal from photons absorbed in the multiplication layer. The multiplication layer, which is where the avalanche multiplications occur, lies below the absorption layer and has sensitivity to $3.5\text{ }\mu\text{m}$.³ Because SAPHIRAs are a developmental program, they are assigned mark numbers, which refer to the generation of metalorganic vapor phase epitaxy layer architecture; higher numbers are newer designs. Mk. 3 SAPHIRA arrays were the first science-grade ones and were delivered in 2013. The Mk. 13 and 14 arrays (which are indistinguishable for the purposes of wavefront sensing) were delivered in 2015 and (as of the time of writing) remain the best arrays

*Address all correspondence to: Sean B. Goebel, E-mail: sgoebel@ifa.hawaii.edu

for telescope deployments due to their uniform cosmetic qualities. Later generations of arrays have improved in some areas (e.g., dark current) but had problems in other areas (e.g., dead pixels or operational reliability). There are plans to produce larger format SAPHIRA arrays for low-background imaging and high-order wavefront sensing on 30-m class telescopes.⁴ SAPHIRA devices produced until late 2015 utilized ME911 read out integrated circuits (ROICs), which only permitted global resets and, therefore, were limited to up-the-ramp readout mode (Sec. 2.1). The ME1000 ROIC enabled row-by-row resets, which made possible read–reset–read and read–reset modes (Secs. 2.2 and 2.3). The ME1001 ROIC is functionally identical to the ME1000 ROIC, but it features reduced glow during operation.

While the operation of the SAPHIRA avalanche photodiode (APD) array is very similar to that of conventional HgCdTe arrays, such as the Teledyne HAWAII series⁵ and the Raytheon Virgo,⁶ there are some important differences. The conventional arrays are normally operated at a bias of a few hundred mV, occasionally up to a volt, and the full bias voltage between the substrate and the node is within the dynamic range of the readout circuit. After reset, the array discharges to the substrate voltage. The depletion is modest and varies with bias voltage and that results in a highly bias-dependent diode capacitance. In contrast, the SAPHIRA typically operated at bias voltages of 2.5 to 18 V, way beyond the dynamic range of the readout circuit, and the node voltage is referenced to ground in the ROIC. At these bias voltages, the photodiode is largely or completely depleted, and the response is highly linear over the dynamic range. As a result of the depletion, there is no trapping of charge when the bias voltage is changed.

SAPHIRA arrays are beginning to enter widespread deployments. The University of Hawaii Institute for Astronomy has extensively tested the detectors in the lab and at telescopes. Over the last several years, we have continuously deployed a SAPHIRA system to the Subaru Coronagraphic Extreme-AO (SCEXAO) extreme AO instrument at the Subaru Telescope, where it is primarily used for focal-plane wavefront sensing.⁷ It has also been deployed⁸ short-term to the NASA Infrared Telescope Facility (IRTF), where it obtained diffraction-limited images of binary star systems using the lucky imaging⁹ technique. Also, a SAPHIRA system is currently being integrated into the Keck II AO system to enable near-infrared pyramid wavefront sensing.¹⁰ SAPHIRA is being used for tip/tilt sensing on the Robo-AO instrument, which was initially deployed to the 1.5-m telescope at Palomar¹¹ and then later to the 2.1-m telescope at Kitt Peak.¹² While at Robo-AO, SAPHIRA produced its first science publication.¹³ The European Southern Observatory (ESO) has also been evaluating SAPHIRA arrays,¹⁴ and they are deployed to the Very Large Telescope's GRAVITY instrument for wavefront sensing at the individual telescopes and fringe tracking at the combined focus.¹⁵ Finally, First Light Imaging has developed a commercial system called C-RED One that combines the SAPHIRA detector, camera and associated cooling, and readout electronics.¹⁶

In this paper, we report on the operation modes and intricacies of behavior of SAPHIRAs. Unless stated otherwise, the data presented below were collected at a temperature of 85 K and at unity avalanche gain (i.e., a bias voltage of 2.5 V) with Mark 13 and 14 arrays on ME1000 ROICs.

2 Readout Modes

The readout modes available on SAPHIRA detectors differ according to when and the number of times that pixels are read

between resets. We have operated the SAPHIRA in three different readout schemes: (1) sampling up-the-ramp mode, which utilizes an initial reset of all pixels followed by multiple reads; (2) read–reset–read mode, which reads a row of pixels, resets it, and reads it again before clocking to the next row; and (3) read–reset mode, whereby a row is read and then reset before moving to the next row, and a reference image is subtracted to remove the pedestal voltage (sometimes called the fixed-pattern noise or bias). These readout modes are independent of subarray size. In practice, we use read–reset mode for AO-related observations because it provides regular time sampling (unlike sample up-the-ramp mode) and better noise performance and double the frame rate compared with read–reset–read mode. In the following sections, we discuss each readout mode in turn.

2.1 Sampling Up-the-Ramp Mode

In sampling-up-the-ramp (SUTR) mode, the entire array is reset, and then it is nondestructively read a user-specified number of times.¹⁷ These reads can be processed in one of two ways. First, the user can fit a line to the “ramp” (flux versus time) behavior for each pixel. This removes the pedestal voltage (fixed-pattern noise) and reduces the white read noise variance as $1/N$, where N is the number of times the array is read per ramp. A mathematical analysis of SUTR (and the conceptually similar Fowler sampling) is given by Garnett and Forrest.¹⁸ Additionally, if a pixel saturates partway through the ramp due to high flux or a cosmic ray strike, the reads in which the pixel is saturated can be excluded from the fit. When SUTR data are reduced by line-fitting, one “science” frame is produced for each ramp. This is well-suited for low-background observations in which read noise dominates and the exposures are long, but it is poorly suited for instances where fast analysis of the images is required (such as AO wavefront sensing).

The second way to reduce SUTR data is to subtract subsequent reads. If being used for AO wavefront sensing, this enables more frequent updates to the correction. By subtracting temporally adjacent frames, one removes the pedestal voltage and kTC (reset level) noise because it is present in both images, leaving only the flux that accumulated in the pixel between those two exposures and the read noise. This is known as a correlated double sample (CDS) image. However, this mode causes three challenges. First, it requires that the dynamic range be rationed across the ramp (if there are N reads per ramp, then no individual read can utilize more than $1/N$ of the dynamic range, or else the final reads of the pixel will contain no new flux). Second, there is irregular time sampling because it would not make sense to subtract the read immediately before a reset from the one immediately after the reset because this would result in negative flux and kTC noise. This is problematic for applications that require regular temporal sampling. Third, due to the intrinsic nonlinearity of infrared arrays, the responsiveness of the detector decreases as the pixels accumulate flux. This causes frame pairs toward the end of the ramp to incorrectly report less flux than an equally illuminated pair at the beginning of the ramp. However, one can compensate for this last effect with a standard linearity correction.

Often, the frame read immediately after a reset contains excess noise due to settling effects. In practice, we typically exclude the first frame following each reset from analysis. This exacerbates the irregularity of time sampling in SUTR mode and reduces its duty cycle, particularly when there are few reads between resets.

In both methods of reducing SUTR data, variations in the voltage of a pixel after one reset compared with its voltage after another reset has no effect because it is in all reads of a given ramp. This is called kTC noise because its variance is given by

$$\sigma_{kTC}^2 = k_B T C, \quad (1)$$

where k_B is the Boltzmann constant, T is the temperature, and C is the node capacitance. We typically operate the detector at $T = 85$ K and $C = 28$ fF for SAPHIRA, so $\sigma_{kTC} \approx 36e^-$. Unlike read noise, kTC noise is not affected by the pixel read rate. Because SUTR readouts are not affected by kTC noise, they enable the lowest noise of the readout modes discussed here. SUTR readouts are best-suited for low-background, long integrations.

2.2 Read–Reset–Read Mode

In read–reset–read mode, a row is read, reset, and then read again before clocking to the next one. The read immediately following the reset (which contains only the pedestal voltage) is then subtracted from the read preceding the next reset after scanning through the array (which contains the pedestal voltage plus flux). Therefore, there are two reads per science frame. This mode removes the effects of kTC noise.

In practice, read–reset–read mode presents two problems. First, and most importantly, settling effects immediately following the reset lead to spurious flux in the first column blocks read. This problem is shown in Fig. 1. The amplitude of this effect varies from one SAPHIRA array to another, and we theorize that it can be mitigated by inserting a delay between the reset and first read, though we have not tested this because it would slow the frame rate. Second, because there is read noise associated with both reads, these add in quadrature for the science frame. For the ~ 100 -kHz clocking rates used for

Fig. 1 On the left are consecutive unilluminated raw frames from SAPHIRA after (top) and before (bottom) the reset in read–reset–read mode. Approximately 2.7 ms has elapsed between the two reads. As is normal for images from infrared arrays, the pedestal noise dominates both images. However, when the before image is subtracted from the after image, not all the noise disappears. An entire row is reset at a time, the detector clocks from left to right, and the column blocks read immediately following the reset have spurious flux. Operating the SAPHIRA with a subarray that excludes these left columns does not solve the problem; the column blocks with spurious flux simply move to the edge of the subarray being read, since they are read immediately following the reset. All images have units of ADUs. We have noticed that this effect is reduced on other SAPHIRA arrays. Because of this postreset setting problem, we do not typically operate SAPHIRA arrays in read–reset–read mode.

most IR detectors, the read noise is much less than the kTC noise. However, the read noise approximately scales with the square root of the sampling frequency, and, for the 1- to 10-MHz sampling frequencies used for SAPHIRA wavefront sensing, these two noise sources become comparable. Because of the spurious flux and replacement of kTC noise with similar-magnitude read noise, we do not commonly use read–reset–read mode.

2.3 Read–Reset Mode

If the read noise of one frame is comparable to its kTC noise, then the correlated double sampling of read–reset–read mode (which produces an image free of kTC noise but containing read noise contributed by both frames) is no better than subtracting a reference dark frame that has been produced from averaging together many frames (and therefore contains negligible kTC and read noise) from a single read (which does contain these noise sources). The read noise depends on the sampling rate and specifics of the readout electronics, but, during typical usage of SAPHIRA, it is of the same order as the $36e^-$ kTC noise calculated from Eq. (1).

In read–reset mode, the pedestal voltage of the detector is removed by averaging together many unilluminated frames and then subtracting this reference frame from each individual frame. During science operation, these dark frames can be collected before or after the data frames. However, a new reference frame needs to be collected for each bias voltage (the bias voltage controls the avalanche gain), since the pedestal voltage pattern on the detector changes with bias voltage. Read–reset mode is our standard mode for AO observations because its noise is similar or lower than that of read–reset–read mode, it produces a “science” frame for every read, and the temporal sampling is regular. Additionally, read–reset mode does not exhibit the post-reset settling problem (the gradient of spurious flux) of read–reset–read mode.

2.4 Comparison of the Readout Modes

A comparison of the loop update rate and noise sources for each of the readout modes is provided in Table 1.

The SAPHIRA detector can be operated with global resets or line-by-line resets. The user can configure which regions of the detector are reset; one can save time by only resetting the pixels that are being read. We typically operate the detector with global

resets when SUTR and line-by-line resets in read–reset and read–reset–read mode. However, line-by-line resets can be used in all three modes. ME1000 and ME1001 ROICs support both global and line-by-line resets, whereas ME911 ROICs only support global resets.

Due to the fact that an entire row is reset at once, but then 32 pixels are read in each increment along the row, in read–reset mode, one side of the detector will see less flux. In other words, for pixels on one side of the array, there is a greater interval between the read and reset than for pixels on the opposite side of the array. This effect is worse for a short, fat subarray than for a tall, thin one. For a 320×256 pixel array, the difference in flux from one side to the other is $\sim 0.3\%$. For a 320×64 subarray, however, this effect is 1.3% . For most astronomical applications, a standard flat field correction effectively compensates for this effect. This effect is not relevant to SUTR readouts because one is calculating the slope of the flux between reads, and it does not affect read–reset–read images because the read of a given pixel is delayed relative to the reset an equal amount in both frames.

In addition to the postreset settling effects discussed in Secs. 2.1 and 2.2, we have also observed a smaller amplitude but longer duration thermal settling in the array. When an array begins clocking after sitting idle, the power dissipation in it changes significantly. This causes thermal variations in the array, and it can take hours for these to reach equilibrium. During this time, pixels exhibit slow drifts. We performed a test in which we read out an unilluminated array for several hours in read–reset mode. Every 10 min, we saved 1000 frames and averaged them to produce an image of the detector’s pedestal voltage with minimal read noise. We then subtracted subsequently collected average frames and computed the standard deviation to see how much the pedestal voltage drifted. The result is shown in Fig. 2. Again, an approximately 2-h settling timescale was observed. For this reason, it is important to run an “idle sequence” before observations in which no data are saved but the detector clocks as it would during science operations. The SAPHIRA detector should be operated like this for at least 2 h prior to observations where temporal stability is required (such as before collecting the reference frames for subsequent read–reset mode observations). It is worth noting that clocking-caused thermal drifts are not unique to SAPHIRA arrays; clocking the same way when idling and observing is key to obtaining good performance from CMOS detectors in general.

Table 1 A comparison of the different readout modes available for the SAPHIRA detector. σ_{RN}^2 is the variance in read noise, N_{reads} is the number of reads between resets in a ramp, σ_{kTC}^2 is the kTC noise, $\sigma_{\text{spurious}}^2$ is the noise associated with postreset settling in read–reset–read mode, $\sigma_{\text{settling}}^2$ is the noise due to thermal drifts and is negligible if the data and reference frames are collected near in time to each other or both are collected after at least 2 h of detector operation. σ_{RN}^2 is assumed to be white. Noise due to RFI (Sec. 4) is a component of σ_{RN}^2 .

	Maximum loop update rate	Noise variance $\sigma_{\text{total}}^2 = \sigma_{\text{photon}}^2 +$	Comments
Up the ramp (CDS pairs processed independently)	Every read, with interruption following reset	$2\sigma_{\text{RN}}^2$	Dynamic range must be split between the reads in the ramp
Up the ramp (flux determined by fitting to ramp)	Every ramp	$\sigma_{\text{RN}}^2/N_{\text{reads}}$	Cosmic rays and saturation can be mitigated by excluding those reads from fitting
Read–reset–read	Every two reads	$2\sigma_{\text{RN}}^2 + \sigma_{\text{spurious}}^2$	The first few column blocks may contain spurious flux
Read–reset	Every read	$\sigma_{\text{RN}}^2 + \sigma_{\text{settling}}^2 + \sigma_{kTC}^2$	Requires reference dark frame

Fig. 2 While the SAPHIRA continuously streamed images, frames were saved every 10 min for several hours. At each interval, 1000 frames were saved and averaged together to reduce read noise. Each average frame was then subtracted from the average frames collected 10 min earlier and later to see how the pedestal voltage of the detector changed. Approximately for the first 2 h of image streaming, the pedestal voltage exhibited slow drifts (as shown in the top left figure). However, after about 2 to 2.5 h, the detector had largely stabilized (top right). The bottom plot shows how the spatial standard deviation of these difference frames decreases with time. The two example frames shown at the top are identified with red points in the bottom plot. For this reason, it is important to clock the detector for at least 2 h before observations that require good stability.

This same 2-h timescale for SAPHIRA thermal drifts was also observed in dark current measurements performed by Atkinson et al.¹⁹ While operating SAPHIRA arrays using up-the-ramp sampling, they noted that it took ~ 2 h of continuous clocking before the dark current stabilized. They directly confirmed this thermal time constant during measurements to calibrate the detector temperature offset from the usual temperature monitoring point at the thermal control sensor/heater assembly. This was done using a temperature sensor permanently mounted to the SAPHIRA leadless chip carrier ceramic (this sensor cannot be powered up during science observations because the glow levels are prohibitive). For the temperature offset calibration, the system was allowed to stabilize with the SAPHIRA powered down (except for the temperature sensor) to establish the offset

between the two temperature sensors. The SAPHIRA was then powered on and clocked normally. After it stabilized, which took ~ 2 h, the on-chip sensor was used to measure the temperature offset.

3 Subarray Operation and Calculation of Frame Rate

There are 10 32-channel column blocks across the SAPHIRA, and it can be subarrayed to read integer numbers of column blocks (i.e., the frame size can be adjusted by 32-pixel increments in the x -direction). On the other hand, the detector will read any number of pixels in the y -direction. Multiple subarrays can be supported at once. The 32 outputs on the SAPHIRA read out 32 adjacent pixels in a row at a time.

Table 2 A summary of the various times for a full-frame and 128×128 subarray readout of the SAPHIRA. These times can be used in Eqs. (2)–(5) to calculate the expected frame rate. Except for the time to reset the detector, the timings are independent of window size. Except for the postreset dead time, the timings are independent of reset mode.

Parameter	Definition	Value	
		Frame size	
		128×128 px	320×256 px
t_{fd}	Duration of frame demand clock pulse	310 ns	310 ns
t_{fc}	Duration of frame completion clock pulse	410 ns	410 ns
n_r	Number of rows being read	128	256
t_{rd}	Time that an output spends on each pixel	1 μ s	1 μ s
n_{cb}	Number of 32-pixel column blocks in the image	4	10
t_{rst}	Time to reset a row	100 ns	1 μ s
t_{prst} (global resets)	Postreset delay	400 ns	400 ns
t_{prst} (line-by-line resets)	Postreset delay	120 ns	120 ns
t_{rc}	Duration of row completion clock pulse	210 ns	210 ns
t_{ra}	Duration of row advance clock pulse	210 ns	210 ns

Greatly increased frame rates during subarray readouts are possible because the interleaved nature of the outputs enables all of them to still be used. To illustrate this, consider the case of a 128×128 pixel subarray. A HAWAII-2RG has 2048 pixels per side and 32 outputs (therefore, 64 columns per output), or a HAWAII-4RG has 4096 pixels per side and 64 outputs (so again, 64 columns per output). Either HAWAII array would utilize two outputs to read the 128×128 subarray, whereas the SAPHIRA could read it using all 32 outputs. At identical pixel clocking rates, the SAPHIRA could read it 16 times faster.

The SAPHIRA frame rate scales roughly inversely with the number of pixels in the window. This rule is approximate because the array must be reset, and there are extra clock cycles at the ends of rows and frames; the full equation for duration of one frame when operating in read–reset mode is given as

$$t_{img,rr} = t_{fd} + t_{fc} + n_r(t_{rd}n_{cb} + t_{rst} + t_{prst} + t_{rc} + t_{ra}), \quad (2)$$

where $t_{img,rr}$ is the time to collect one image in read–reset mode, and the other parameters and their typical values are summarized in Table 2. These clock pulses are all necessary for the detector to operate reasonably and cannot be set to 0; the acceptable ranges are specified in the SAPHIRA manual. Read–reset–read mode is similar, but each pixel is read twice, and there is an extra t_{rc} per row

$$t_{img,rrr} = t_{fd} + t_{fc} + n_r(2t_{rd}n_{cb} + t_{rst} + t_{prst} + 2t_{rc} + t_{ra}). \quad (3)$$

For sample up-the-ramp mode, the time to collect a frame (not including the time to reset, since that typically is not done every frame) is

$$t_{img,sutr} = t_{fd} + t_{fc} + n_r(t_{rd}n_{cb} + t_{rc}). \quad (4)$$

Naturally, the frame rate f is

$$f = t_{img}^{-1}. \quad (5)$$

The timing values for our current 1-MHz pixel rate are summarized in Table 2. These are subject to tuning but can be used as typical values. In line-by-line reset modes, these enable a 1.69-kHz frame rate for a 128×128 pixel subarray or 339-Hz frame rate for 320×256 pixel full frames.

It should be noted that, in all readout modes described above, the detector is continuously clocking and reading pixels. This is known as a rolling shutter; different parts of the frame are read at different times. This can be problematic for observations for which the flux is rapidly changing and synchronization is important. For example, in present pyramid wavefront sensors deployed on-sky, the beam is modulated in a circle around the tip of the pyramid to increase the linearity of the wavefront sensor's response. This causes each pupil image to be illuminated at a different time. In SCEXAO, the beam completes one full modulation per EMCCD image, and the EMCCD has a fast frame transfer that functions as a global shutter, so the detector does not notice the pupils being illuminated at different times. On the other hand, a detector with a rolling shutter might see gradients across the pupil images caused by the modulation, interpret it as an optical aberration, and drive the deformable mirror to a suboptimal shape. The Keck SAPHIRA solves this challenge by synchronizing the start of each frame's readout to the modulation, so the part of the frame being read does not include the pupil, which is changing in illumination. The time to modulate is slightly greater than the time to read one image, but the dead time between reads is minimized for the reason described in the next paragraph.

Fig. 3 Shown at top is a sample single frame (with the pedestal noise subtracted) which exhibits the RFI of the SCExAO environment. In the middle is a power spectrum calculated from the corresponding full image cube. Apart from a major spike at 60 Hz (likely corresponding to ground contamination from other electronics), the noise is largely white. Plotted over the power spectrum is a Welch-filtered version. The Welch filter²³ reduces noise in a power spectrum, but this comes at the expense of spectral resolution. As a result, the lowest frequency bins of the Welch-filtered data are much broader. At the bottom is the cumulative sum of the unfiltered power spectrum.

The integration times of an image given by Eqs. (2)–(4) depend entirely on the durations of the various clock pulses and number of pixels being read. A user could insert an additional delay to increase the detector's integration time. However, unless storage space is limited, this is suboptimal. Instead of inserting a delay to make each exposure longer, the user does better to clock the detector at its maximum rate, increase the avalanche gain to use the entire well depth with each integration (assuming read–reset or read–reset–read mode), and then average images together. This reduces the read noise and, therefore, improves the signal-to-noise ratio. This contrasts with a normal (nonavalanche multiplying) detector, wherein reading and resetting the array when very little target flux has accumulated and then coadding the exposures results in the read noises adding in quadrature and, therefore, a reduced signal-to-noise ratio. (In both normal and avalanche-multiplying detectors, nondestructively reading up the ramp at maximum readout rate until the entire well depth has been used results in optimal signal to noise, assuming that ROIC glow is negligible. For SAPHIRA, one can trade-off the ramp length and avalanche gain; the optimal signal-to-noise ratio can be obtained by selecting the avalanche gain that produces the minimum dark current and then using a ramp length that fully utilizes the detector's dynamic range.)

4 Radio Frequency Noise and Cryogenic Preamplifiers

Radio frequency interference (RFI) noise is the dominant noise source for SAPHIRA's deployments to Subaru Telescope with the SCEXAO instrument, but it is negligible in other environments. In the laboratory, the measured SAPHIRA CDS read noise ($9e^-$ RMS⁸ at unity avalanche gain and a Generation III Leach Controller²⁰ at a 265-kHz clocking rate) is very comparable to the read noise of conventional HAWAII series arrays.^{21,22} The read noise was similar in telescope environments, such as the NASA IRTF and Kitt Peak. However, when mounted to the SCEXAO instrument at the Subaru telescope, the noise is increased by a factor of 4.5, and we have been unable to reduce this through the normal procedures. The SAPHIRA is ground isolated within the cryostat with the analog and digital grounds brought out separately to the controller and then to a “star” ground point. The internal radiation shield and outer vacuum vessel are tied separately to this same star ground to provide concentric Faraday cages. The entire camera system is isolated from the SCEXAO bench except through a tie from the star ground to the primary SCEXAO ground. SCEXAO is characterized by an extremely challenging RFI environment and nearly all the instruments experience higher levels of noise than they did during laboratory testing elsewhere.

RFI shows up as 32-pixel-wide blocks that oscillate in value (Fig. 3). A power spectrum can reveal some of the sources of RFI. In our current SCEXAO deployment, the noise is white except for a 60-Hz spike. In previous deployments, it was dominated by a few discrete frequencies (and therefore could be mostly filtered out from the data). The RFI changes with time, grounding setups, and the arrangement of wires on the outside of the camera. We found that RFI was decreased by increasing the shielding of cables, using twisted-pair cables, eliminating ground loops, and using a “clean” (i.e., with minimal other electronics connected) ground point. RFI behaves as read noise in the noise comparisons of Sec. 2.4.

We are currently testing cryogenic preamplifiers that were developed at Australia National University. These sit near the detector inside the camera and amplify its electrical signals prior to the wires to the readout electronics. We anticipate that these will greatly reduce the RFI noise. Additionally, our maximum pixel rate until now has been limited to about 1 MHz per output due the ROIC output drivers causing settling issues with our ~1-m cable lengths. We expect that the preamplifiers will enable faster readout rates.

5 Avalanche Gain, Charge Gain, and the Excess Noise Factor

In a conventional source-follower HgCdTe array, such as the Teledyne HxRG family⁵ or the Raytheon Virgo,⁶ there are two gains of interest: the voltage gain (which has units of $\mu\text{V}/\text{ADU}$) and the charge gain (which has units of e^-/ADU). The charge gain is related to the voltage gain through $Q = CV$, where C is the integrating node capacitance (the sum of the photodiode and ROIC capacitances) and has a typical value of 20 to 40 fF . The voltage gain is a property only of the ROIC. In the above-mentioned infrared detector arrays, the bias voltage is low enough that there is minimal change in the photodiode capacitance as the pixel accumulates photons, and this effect can be compensated for with a linearity correction. However, the situation is somewhat more complicated for SAPHIRA. The user adjusts the bias voltage across the detector over a 20-V range to set the avalanche gain. At bias voltages too low to produce any avalanche gain ($\lesssim 1$ V), operation is similar to the conventional source-follower arrays. However, as the bias voltage of SAPHIRA increases, avalanche multiplications begin to occur and the diode capacitance, and therefore the charge gain, decreases. Both the avalanche multiplication and the decreasing charge gain should be taken into account when converting between collected photons, electrons, and ADUs for data from SAPHIRA detectors.

During SAPHIRA operation, the input node of each pixel is reset to the selected bias voltage relative to the detector common voltage. After the reset is lifted, photoelectrons discharge the integrating node, and the pixel saturates when the node voltage equals common. The node voltage is read out by an ROIC with a typical gain of 0.8 to 0.9 (for the single source-follower design of the ME1000/ME1001) and is further amplified by the controller preamplifier chain before conversion to ADU. The voltage gain of the system (in units of $\mu\text{V}/\text{ADU}$) can be calibrated by varying the reset voltage in known increments and measuring changes in the corresponding digitized signal. The charge gain is normally determined through the signal versus variance method, although Finger et al.²⁴ directly measured a V versus Q relation for an H2RG detector by connecting it to a calibrated large external capacitor and measuring voltage changes in each while pixels were exposed and then reset.

The photodiode capacitance of SAPHIRA increases as the node is discharged, and this causes a flux versus time plot to depart from linearity as the pixel accumulates flux. Figure 4 shows the photodiode capacitance of a Mark 13 SAPHIRA from 1 V (where the avalanche gain is negligible but the capacitance varies significantly with bias voltage) to 9 V (where the diode is fully depleted so the capacitance is insensitive to bias, but the avalanche gain is ~ 5). Avalanche gain becomes appreciable at a bias voltage of 2 to 3 V, and, by biasing the detector up to 20 V, avalanche gains of ~ 600 can be obtained. In all cases, the unit cell source follower still limits dynamic range

Fig. 4 The measured integrating node capacitance versus bias voltage for a Mk. 13 SAPHIRA detector. The data are from Ref. 25. 19.9 fF of the integrating node capacitance is due to the silicon readout, and the remainder is the diode. It is important to take into account the increasing capacitance at low bias voltages when measuring avalanche gains.

to a few hundred mV. The avalanche gain is usually measured by varying the bias voltage under constant detector illumination and defining that the flat portion of the response curve (which occurs at 1- to 2.5-V bias) corresponds to an avalanche gain of 1. From Fig. 4, it is evident that for avalanche bias voltages > 10 V, the change in capacitance with bias must be included in this calculation.

Read noise and RFI noise are independent of avalanche gain, so one can improve the signal-to-noise ratio of observations by amplifying the signal of photons with a higher avalanche gain. However, as avalanche gain increases, the dynamic range (in photons) of the detector is reduced because the well depth in electrons remains constant, but more electrons are produced per photon. At sufficiently high bias voltages, tunneling current (a type of dark current) increases dramatically,¹⁹ can dominate over the read noise, and thereby can degrade the signal-to-noise ratio of observations³.

The excess noise factor F characterizes noise that occurs during the multiplication process of avalanche photodiodes. F is the ratio of the signal to noise of the photon-induced charge to that of the avalanche-multiplied charge;²⁶ noise-free amplification produces $F = 1$. The advantage of HgCdTe over other semiconductors is that only electrons participate in the avalanche process, and this has the potential for noise-free amplification. Finger et al.³ measured an excess noise factor of $F = 1$ at 60 K at all gains and a maximum of $F = 1.3$ at an avalanche gain of 421 at 90 K.

6 Conclusions

For most AO-related observations, we operate SAPHIRA in read–reset mode, wherein each row is read before being reset, and the pedestal noise of the detector is removed by subtracting a reference dark frame. This mode enables the maximum usable frame rate and regular time sampling. When a minimum of noise is desired and long periods between detector resets are acceptable, the SAPHIRA can be operated in up-the-ramp mode, wherein there are many reads between resets. The user then fits a line to the flux of each pixel over time, which reduces read noise. If a pixel saturates, one can estimate the correct flux by fitting only to the reads that occurred before the

saturation. Read–reset–read mode, wherein a row was read, reset, and read again, and this repeats for the next line, failed to deliver the hoped-for low noise due to a postreset settling effect that caused spurious flux in the columns read in the first few microseconds after resets. For this reason, we do not commonly use read–reset–read mode.

Because the pixel rate is determined by the readout electronics (and therefore is relatively constant) and the outputs are designed to read contiguous 32-pixel blocks, higher frame rates can be achieved by reading $32x \times y$ pixel subarrays, where $1 \leq x \leq 10$ and $1 \leq y \leq 256$. For example, at a 1-MHz clocking rate in read–reset mode, one can read the full 320×256 pixel frame at a rate of about 339 Hz or a 128×128 subarray at about 1690 Hz. The detector itself is capable of 10-MHz pixel rates, and ESO operates their SAPHIRA arrays at typically 5-MHz rates.

SAPHIRA detectors enable the potential for observing modes due to their high frame rate, low noise, and infrared sensitivity. We hope that this paper can assist users in utilizing the detectors to their maximum potential.

Acknowledgments

The authors acknowledge support from NSF award AST 1106391, NASA Roses APRA award NNX 13AC14G, and the JSPS (Grant-in-Aid for Research Nos. 23340051 and 26220704). Sean Goebel acknowledges funding support from Subaru Telescope and the Japanese Astrobiology Center.

References

1. P. Wizinowich et al., “Near-infrared wavefront sensing,” *Proc. SPIE* **9909**, 990915 (2016).
2. I. Baker et al., “Leonardo (formerly Selex ES) infrared sensors for astronomy: present and future,” *Proc. SPIE* **9915**, 991505 (2016).
3. G. Finger et al., “Sub-electron read noise and millisecond full-frame readout with the near infrared eAPD array SAPHIRA,” *Proc. SPIE* **9909**, 990912 (2016).
4. D. N. B. Hall, I. Baker, and G. Finger, “Towards the next generation of L-APD MOVPE HgCdTe arrays: beyond the SAPHIRA 320×256 ,” *Proc. SPIE* **9915**, 991500 (2016).
5. J. W. Beletic et al., “Teledyne imaging sensors: infrared imaging technologies for astronomy and civil space,” *Proc. SPIE* **7021**, 70210H (2008).
6. B. Starr et al., “RVS large format arrays for astronomy,” *Proc. SPIE* **9915**, 99152X (2016).
7. S. B. Goebel et al., “Evolutionary timescales of AO-produced speckles at NIR wavelengths,” *Proc. SPIE* **9909**, 990918 (2016).
8. D. Atkinson et al., “Observatory deployment and characterization of SAPHIRA HgCdTe APD arrays,” *Proc. SPIE* **9154**, 915419 (2014).
9. D. L. Fried, “Probability of getting a lucky short-exposure image through turbulence,” *J. Opt. Soc. Am.* (1917–1983) **68**, 1651 (1978).
10. D. Mawet et al., “Keck planet imager and characterizer: concept and phased implementation,” *Proc. SPIE* **9909**, 99090D (2016).
11. C. Baranec et al., “High-speed imaging and wavefront sensing with an infrared avalanche photodiode array,” *Instrum. Methods Astrophys.* **809**, 70 (2015).
12. R. Jensen-Clem et al., “The performance of the Robo-AO laser guide star adaptive optics system at the Kitt Peak 2.1 m telescope,” *Astron. J.* **155**(1), 32 (2018).
13. E. Han et al., “Magnetic inflation and stellar mass I: revised parameters for the component stars of the kepler low-mass eclipsing binary T-Cyg1-12664,” *Sol. Stellar Astrophys.* **154**, 100 (2017).
14. G. Finger et al., “Evaluation and optimization of NIR HgCdTe avalanche photodiode arrays for adaptive optics and interferometry,” *Proc. SPIE* **8453**, 84530T (2012).

15. L. H. Mehrgan et al., “GRAVITY detector systems,” *Proc. SPIE* **9907**, 99072F (2016).
16. T. Greffe et al., “C-RED One: the infrared camera using the Saphira e-APD detector,” *Proc. SPIE* **9907**, 99072E (2016).
17. R. Glendinning et al., “Implementation of a charge integration system in a low-background application,” *Proc. SPIE* **1235**, 34–42 (1990).
18. J. D. Garnett and W. J. Forrest, “Multiply sampled read-limited and background-limited noise performance,” *Proc. SPIE* **1946**, 395–404 (1993).
19. D. E. Atkinson et al., “Dark current in the SAPHIRA series of APD arrays,” *Astron. J.* **154**, 265 (2017).
20. R. W. Leach and F. J. Low, “CCD and IR array controllers,” *Proc. SPIE* **4008**, 337–343 (2000).
21. D. N. B. Hall et al., “Performance of the first science grade $\lambda C = 2.5\ \mu\text{m}$ HAWAII 4RG-15 array in the laboratory and at the telescope,” *Proc. SPIE* **9915**, 99150W (2016).
22. O. D. Fox et al., “Performance and calibration of H2RG detectors and SIDECAR ASICs for the RATIR camera,” *Proc. SPIE* **8453**, 84531O (2012).
23. P. D. Welch, “The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms,” *IEEE Trans. Audio Electroacoust.* **15**, 70–73 (1967).
24. G. Finger et al., “Conversion gain and interpixel capacitance of CMOS hybrid focal plane arrays,” *Exp. Astron.* **19**, 135–147 (2005).
25. I. Baker, Leonardo MW Ltd., Hampshire, UK (2017).
26. R. J. McIntyre, “Multiplication noise in uniform avalanche diodes,” *IEEE Trans. Electron Device* **13**, 164–168 (1966).

Sean B. Goebel is a PhD student at the University of Hawaii Institute for Astronomy. He works with Hall and Guyon to evaluate SAPHIRA arrays, deploy them to the Subaru SCExAO instrument, and use them for wavefront measurements. SAPHIRAs are high frame rate, low noise HgCdTe L-APD arrays, and SCExAO is an extreme adaptive optics (AO) instrument intended for direct observations of high-contrast objects, such as exoplanets and debris disks.

Donald N. B. Hall received his PhD from Harvard in 1970. He introduced InSb for astronomy. His research currently focuses primarily on the development of HgCdTe arrays for astronomy. He was principal investigator for the development of the NASA HAWAII-2RG and SIDECAR ASIC and the NSF HAWAII H4RG-15, all teamed with Teledyne. He is also PI on optimization of SAPHIRA APD arrays and the development of a $1\text{ k} \times 1\text{ k}$ APD array for low-background astronomy—both teamed with Leonardo.

Olivier Guyon is an expert in exoplanet imaging techniques: coronagraphy and AO. He leads the Subaru Coronagraphic Extreme-AO instrument and designs coronagraph systems for future space-based exoplanet imaging missions. He received the US Presidential Early Career Award for Scientists and Engineers and the MacArthur fellowship for his innovative contributions to astronomical optics. He is an avid amateur astronomer and citizen science enthusiast and founded the citizen science project Panoptic Astronomical Networked Observatories for a Public Transiting Exoplanets Survey.

Eric Warmbier has been an electrical engineer at the NASA IRTF since 2007. Much of his work includes building and designing the electronics for infrared instruments that use the H2RG and Aladdin detectors. Other work includes telescope and facility hardware, such as piezo-driven choppers and drives for the shutter, telescope, and dome. He is the hardware designer for the Pizza Box infrared controller, which supports H2RG and SAPHIRA detectors and is currently under development.

Shane M. Jacobson is an instrumentation engineer for the University of Hawaii Institute for Astronomy. His specialization revolves around HAWAII HxRG series arrays and SAPHIRA detectors. He provides support to the Hilo IfA Instrumentation Division’s launch of these new detector technologies within astronomical systems. Notable support of past instruments includes UH88” ULBCam, Subaru Telescope’s HiCIAO, Cerro Armazones Observatory’s IRIS, Subaru Telescope’s IRD, and numerous deployments of GLS SCC cryostats containing SAPHIRAs to various observatory settings.