

CCD speed-noise optimization at 1 MHz

Roger Smith and Stephen Kaye

California Institute of Technology, 1600 E. California Blvd, Pasadena, CA 91125, USA

ABSTRACT

The Zwicky Transient Facility (ZTF)[1] is a CCD mosaic requiring 64 differential output channels to be transmitted to electronics located over 2 m from the CCDs and digitized with less than 10 e⁻ read noise at 1 MHz pixel rate. To minimize pixel overhead, the Reset Gate pulse is generated inside the dewar by a pin driver controlled by a very short pulse using an LVDS interface. Overlapping serial clocks spanning the entire pixel are made entirely triangular with slopes tuned to cancel substrate return current and minimize high frequency content to improve common rejection by the fully differential signal path. We document the trade between settling time after charge dump and linearity and illustrate the desirability of generating both Summing Well and Reset Gate edge close to the CCD. The parallel clocking overhead is exacerbated in ZTF by ganging multiple CCDs but is hidden by overlapping the parallel shift with pixel readout. To suppress fixed pattern due to the concurrent parallel clocks, slow overlapped triangular waveforms panning the entire line time are employed to null the substrate current, in the same manner as the serials. Both noise and speed requirements are exceeded on all 64 channels, with margin. At all pixel rates the median noise is as good as can be expected for differential transmission being $\sqrt{2}$ times the single sided noise published in the data sheet for the CCD231-C6 CCDs. Linearity is preserved even at 840 ns pixel time, and crosstalk is less than 10 ppm.

Keywords: CCD, speed, noise, differential, feedthrough, correlated double sampling

1. INTRODUCTION

This paper discusses the optimization of the speed-noise trade. It is in fact applicable to any pixel rate but was motivated by the additional challenges faced at high pixel rates where activities other than noise averaging tend to dominate pixel time.

CCD controller designers typically must minimize time between exposures to maximize time collecting photons. Data transfer time can be hidden by buffering data within the CCD controller [1], or by transmitting data to the host concurrently with readout. Though not yet standard practice, parallel transfer time can be hidden by parallel clocking during pixel readout [2]. Although pre-scan pixels can be shifted faster than pixel reads, this usually comes with a penalty of a strong fixed pattern at line start as serial clocks recover after the high loading during fast pre-scan. The number of overscan columns can be reduced to about 20, since overscan only needs to be wide enough to support accurate fitting to the drift in offset during a frame, after discarding (~5) columns corrupted by serial charge deferral.

For any given pixel rate, readout time can be reduced by adding readout amplifiers to the CCD, but the cost of additional electronics is commensurately greater. While possible in theory, the cost of developing a CCD with more output amplifiers is prohibitive for small projects. Amplifier count is limited ultimately by total power dissipated in the output amplifiers (5-20 mW/ch) and the density of traces on the flex circuit connecting to the CCD. While several hundred amplifiers per side might be possible, only a few designs exist with more than 8 amplifiers per side. Thus there is always pressure to minimize the pixel time.

The pixel is comprised of CCD reset pulse and the settling time for the transient it produces, the “dwell time” when post-reset level is measured, settling time for the low going edge on Summing well which dumps charge onto the sense node, and the post-charge-dump measurement. The difference between these two measurements is the signal. We have succeeded in spreading the serial clock transitions over the entire pixel, without impacting noise performance, so that the transport of charge from serial register to SW does not represent an additional overhead.

The post-reset level must be re-measured after each reset, since thermal noise in the reset switch resistance produces a voltage fluctuation that is trapped on the sense node capacitance when the reset switch opens (called kTC noise). The subtraction of the post-reset level also removes drift due to 1/f noise in the CCD output MOSFET, and noise capacitively coupling to the sense node from bias and clock voltages at frequencies lower than the pixel rate. Our goal is to minimize

the “overheads” (Reset, serial shift time, settling times) so that more time can be allocated to signal averaging to reduce noise in the frequency range between the pixel rate and the bandwidth-limit imposed by the antialiasing filter of the Analog to Digital Converter.

The CCD controller used for the tests reported here is a 16 channel Archon made by Semiconductor Technology Associates, sampling at 100 MHz and averaging ~30 samples during each of the dwell times at 1.2 MHz pixel rate. The clock driver incorporates an independent state machine, which updates a 100 MHz DAC to implement programmable slew rate (in reality a staircase) on each clock edge. The sixteen e2v CCD231-C6 CCDs employed in the ZTF focal plane are each read out through 4 *differential* outputs per CCD. After AC coupling, preamps amplify the signal by 1.3 and transmit 64 CCD output signals differentially through 3 m long twinax cables, 100 Ω series termination at the driver. The signal path is fully differential from CCD to 16 bit 100 MHz ADC. The Archon gain is set for a 4V full scale input, which combined with 7 μV/e- CCD sensitivity produces 5.8 e-/ADU conversion gain, so that the 350 ke- well capacity maps to 60,300 ADU. A programmable voltage is supplied to the black-level clamp at the Archon video board input to allow the post-reset level to be offset to make the full ADC range available for the signal swing.

2. SERIAL CLOCKS

The pixel can be divided into the following sequential periods:

1. Reset Gate pulse (RG = RØ is high)
2. Reset settling time
3. Reference level sampling
4. Charge dump and settling
5. Signal sampling

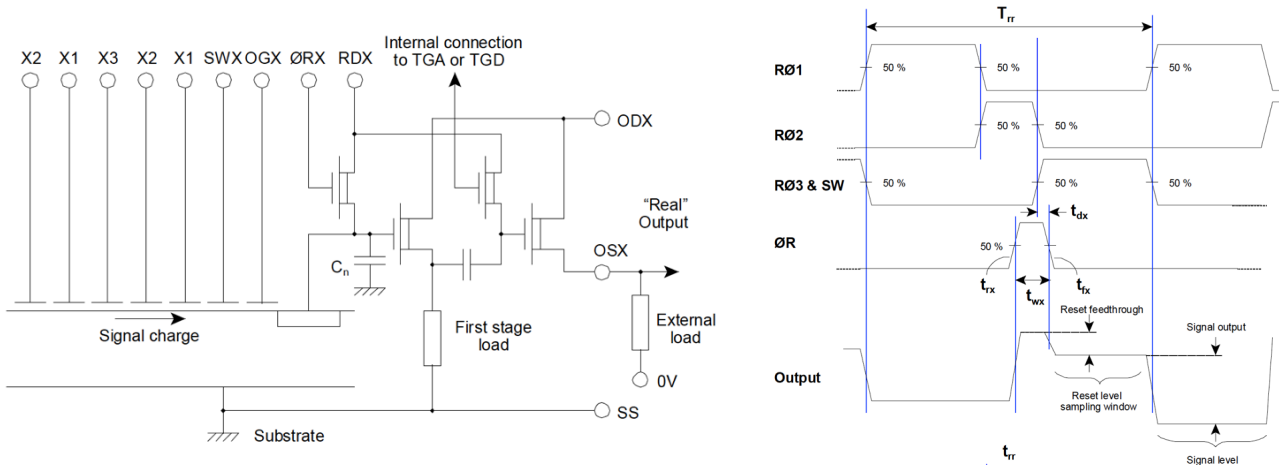


Figure 1: e2v CCD231-C6 output amplifier (left) and e2v's recommended readout timing (right)

In the timing waveform recommended by e2v (Figure 1), signal sampling may occur without corruption from transients capacitively coupling from the serial clocks, only if the sampling of the reference occurs while serial clock phase 3 (RØ3) is high and sampling of the signal occurs while phase 1 (RØ1) is high. However the time RØ2 is high is then unavailable for sampling.

Instead, we spread the clock edges evenly throughout the pixel then set the slew rates to the lowest possible value that still allows the clock to reach its high or low level before reversing. Slowing the clocks edges is advantageous when using differential CCD outputs, since the common mode rejection is always better at low frequencies. Feedthrough to the video is caused not only by capacitive coupling from clocks to sense node but by the voltage induced by return currents as they flow through finite resistance in the substrate, traces, and return wiring to the clock (Figure 3).

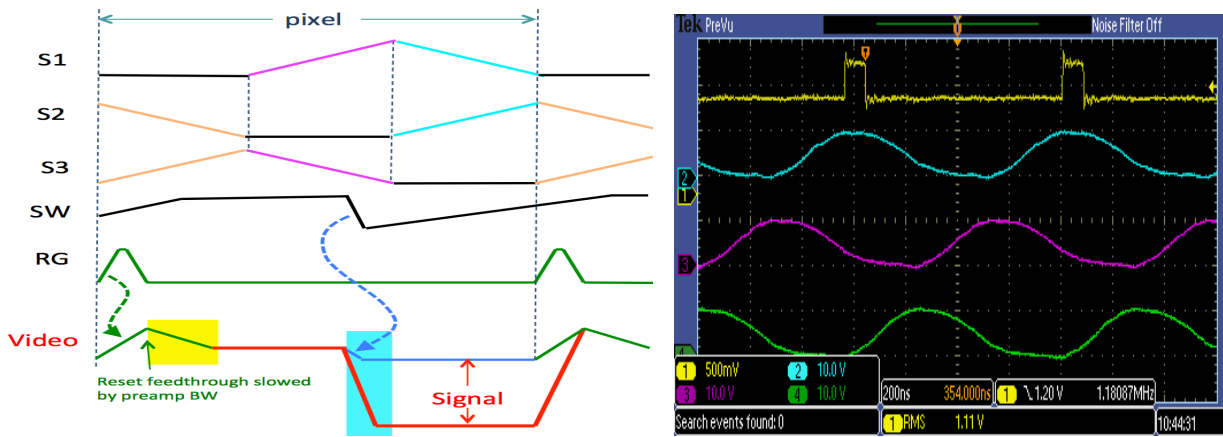


Figure 2: Left: Pixel timing diagram using equispaced triangular serial clocks to minimize slew rates. Right: measured waveforms showing RG (LVDS, yellow), S2 (cyan), S1 (magenta) and S3 (green). Waveforms appear more sinusoidal than triangular since serial clock capacitance attenuates higher frequency components.

For coincident clock edges, substrate currents have opposite sign can be arranged to cancel exactly.

$$i_1 + i_2 = 0$$

when

$$C_{P1} \frac{\partial V_1}{\partial t} + C_{P2} \frac{\partial V_2}{\partial t} = 0$$

So for linear slopes,

$$C_{P1} \frac{\Delta V_1}{\Delta t_1} + C_{P2} \frac{\Delta V_2}{\Delta t_2} = 0.$$

We require the edge transition times to be matched ($\Delta t_1 = \Delta t_2$). The cancellation condition is then

$$C_{P1} \Delta V_1 + C_{P2} \Delta V_2 = 0.$$

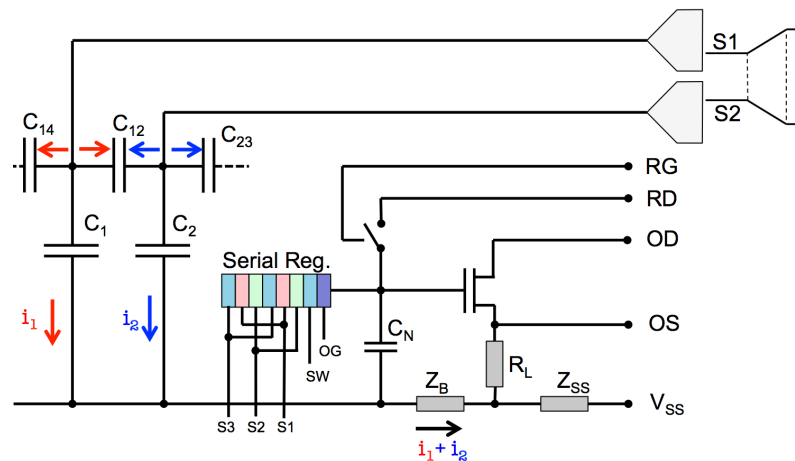


Figure 3: Currents flowing from S1 and S2 drivers through serial clock capacitances and returning via Vss. Currents i_1 and i_2 have opposite sign and thus tend to cancel when S1 and S2 edges have opposite slopes.

In practice one tunes the voltage swing while simultaneously adjusting slew rate so slew times remain the same, searching for the case where serial clock feedthrough is minimized (Figure 4).

This approach eliminates the serial clock overhead while minimizing both loading on the clock driver and the feedthrough to the video signal. Even when substrate currents are not perfectly nulled, the slow edges are still beneficial in reducing ringing caused by inductance in the long cables back to the CCD controller, which are 3m long in ZTF.

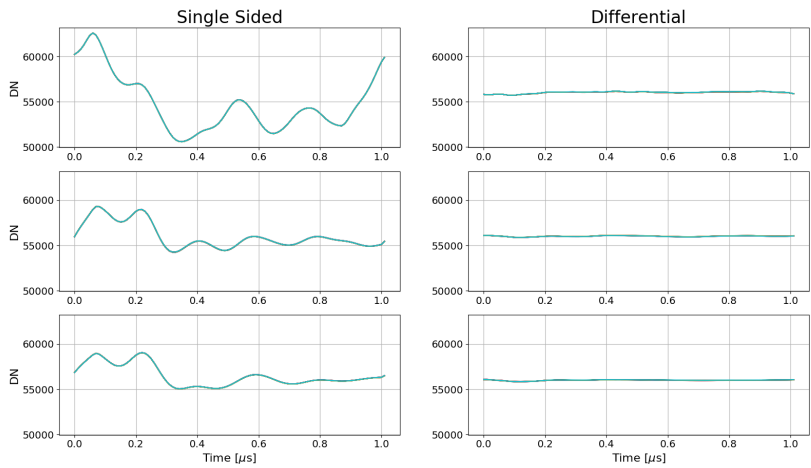


Figure 4: Video waveform, acquired using the Archon’s oscilloscope mode, showing serial clock feedthrough pulses for conventional fast edges (top), triangular clocking (middle) and triangular clocking with amplitudes and slopes adjusted to null the substrate current (bottom). Single-sided video at left; differential video at right. The SW pulse has been held high in this example to highlight serial clock feedthroughs.

3. SUMMING WELL

3.1 Video Settling

Operational amplifiers generally exhibit lower bandwidth for large signals. Consequently video settling (after charge dump) to a given percentage of step size takes longer for larger signals. This is a source of non-linearity that becomes significant if signal averaging begins too soon after charge dump as illustrated in Figure 5.

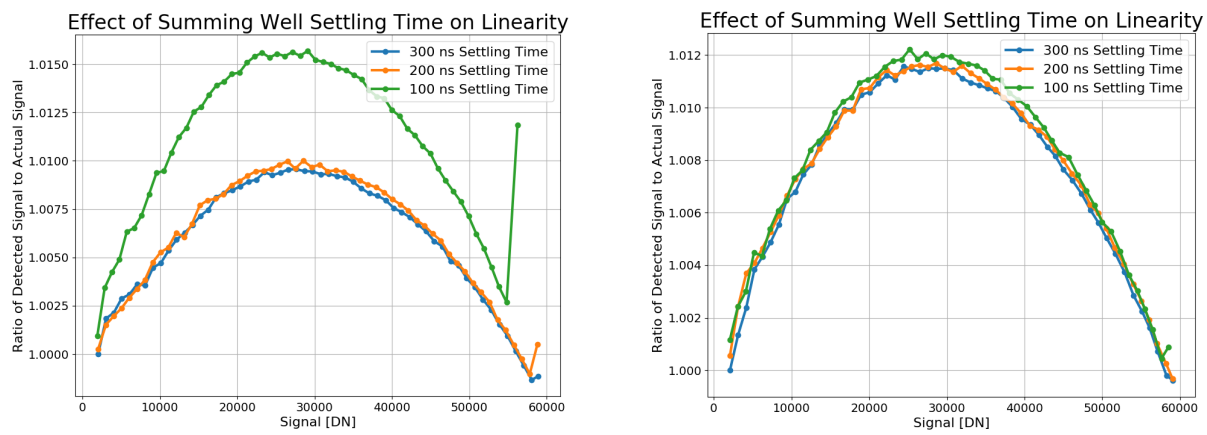


Figure 5: Non-linearity curves. At left: measured signal divided by expected as a function of measured signal, for 100 ns, 200 ns and 300 ns settling time between charge dump and start of a 200 ns sample averaging period. Right: The same but with 1μs sample averaging to illustrate how the degradation of the non-linearity is diluted at longer pixel times where signal settling represents a smaller fraction of averaging time.

3.2 SW Feedthrough

Even in the absence of charge, a step occurs in the video signal due to capacitive coupling of the SW clock to the sense node. The charge dump should be fast enough for the video bandwidth to limit the settling rather than the clock slew rate. Unfortunately the CCD controllers for ZTF had to be located outside the telescope tube so that the clock driver is 2m from the CCD. Inductance in the long cable tended to produce some ringing (figure 4). Reducing slew rate did help but this method could not be employed to entirely eliminate the ripple without impacting settling time. Fortunately, the ringing is sufficiently reproducible, and the timing jitter for sampling is low enough that noise is not seriously impacted provided that 160 ns is allowed for settling before start of sampling. We will see below that large signal settling requires a longer delay. In future, for systems running at or near 1 MHz, we plan to minimize ringing on SW feedthrough by placing the SW driver in the dewar. We will also select a video preamp with higher bandwidth – closer to the Nyquist frequency – so that large signal settling after charge dump can be faster and thus have less effect on linearity.

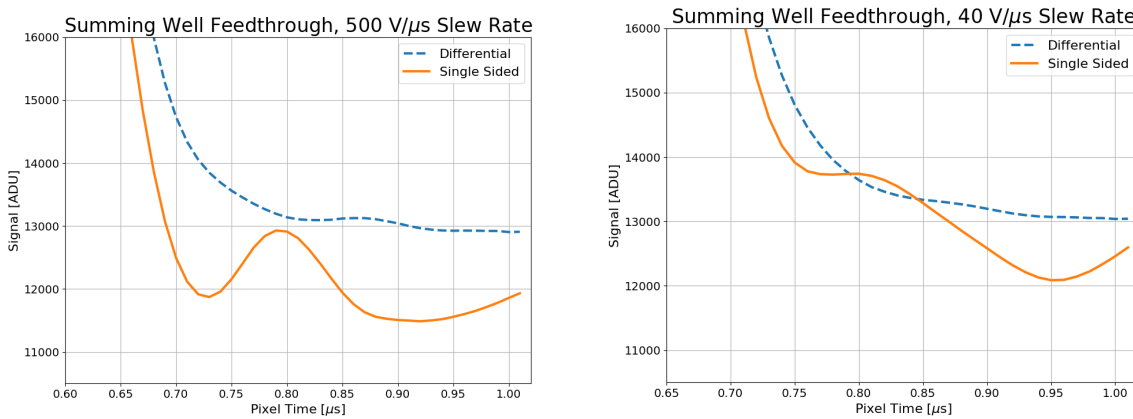


Figure 6: Summing Well feedthrough pulse for two falling edge slew rates, 500 V/μs and 40 V/μs. Single sided (solid) and differential video (dotted). Ripple after slower edges is better suppressed by the differential receiver.

3.3 Asymmetric Slew Rates

Only the falling edge of SW needs to be fast. We slow the rising edge (so that it generates less ripple) by starting the rising edge as soon as it reaches its most negative value, and then extend the rise to the beginning of falling edge on S2 (R02) at start of the next pixel (Figure 2).

4. RESET GATE

4.1 Reset Pulse Width

Figure 7 shows feedthrough, in single sided and differential modes, for two different RG pulse widths. If RG pulse is short compared to the time constant of the video signal path, the feedthrough pulse is reduced in amplitude. The time for the reset to occur is negligible compared to the typical slew rate on RG. Even if the FET on-resistance never falls below 10 kΩ, the sense node capacitance is so low (~25 fF) that the RC time constant would still be only 200 picoseconds.

Reset Gate (RG) capacitively couples to the sense-node, mostly via the Reset FET Gate-source capacitance, causing a level shift in the video only while the reset switch is open. As a result, RG feedthrough is proportional to the difference between the low level on RG and the switching threshold of the Reset MOSFET and is independent of the RG high level. Raising RG_low to within 2 V of the switching threshold will limit the feedthrough amplitude, but moving closer to the threshold raises the risk of the reset switch not turning off completely. One must keep in mind that this switching threshold is relative to Reset Drain (RD).

Multiple methods are available for determining whether the reset pulse is too short. The linear scaling of feedthrough amplitude versus RG_low , and insensitivity to RG_high , allows the switching threshold to be estimated correctly, while also confirming that switching is indeed occurring. Insufficient RG pulse duration also causes a shift in baseline in response to signal and consequent loss of gain after CDS processing. Using the incremental binning method [3] one can measure linearity using just two flat field exposures making this a practical method for detecting this gain loss quickly.

In ZTF we found that 80 ns was sufficient provided that RG_high was increased to boost slew rate enough to reach the switching threshold. This shortened RG pulse reduced feedthrough as shown in Figure 7.

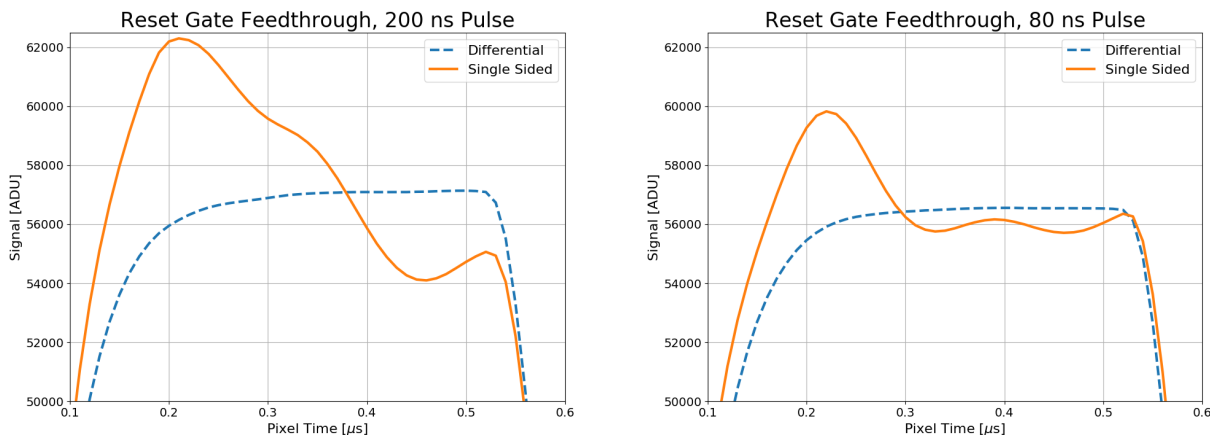


Figure 7: Reset Gate feedthrough for 200 ns and 80 ns RG pulses, in differential and single sided mode.

4.2 Local Reset Gate Driver

RG couples strongly to the sense node and thus must have low noise, fast rise time and minimal overshoot and ripple. This would have been difficult to achieve in ZTF [4] given that the CCD controller had to be located 2 meters away (outside the telescope tube). Instead, a fully balanced Low Voltage Differential Signal carries the timing information, while two low noise biases supply the high and low levels, with switching provided by a high speed pin driver located close to the CCD connector (Figure 8) inside the dewar. Bypass capacitors close to the pin driver then supply the transient current during the RG edges to minimize the length and inductance of the wiring. The high speed component of the current now only circulates between CCD and local bypass capacitors instead of through the 2m cable. The traces on flex circuit are quite thin and thus have enough inductance to setup ringing in conjunction with these bypass capacitors so $100\ \Omega$ is added in series to damp the oscillation. $100\ pF$ has been added to limit slew rate to only that required. In future we will use a similar driver for SW to allow faster charge dump with minimal ringing.




Figure 8: Reset Gate driver located close to CCD connector.

4.3 Signal Settling after Reset

Thus far we have only discussed settling of the capacitively coupled RG control signal but one must also consider settling of *signals* after reset. When the signal settling is insufficient, the pixel following a bright pixel will be offset negative by an amount that is roughly proportional to signal. This pixel-lag effect tends to be worse at high signals due to slew rate limiting. The effect is diluted for longer signal averaging times and is thus not commonly seen. We aim to keep pixel lag below 0.05% so that it is well below effects such as deferred charge (~1%), non linearity (1%), and far below the level at which Point Spread Function would be significantly affected.

To test this we require a signal with a contrast level that is difficult to achieve by imaging since the overscan edge is contaminated by deferred charge. Instead we synthesize high contrast edges using the clock pattern as follows. One can generate a linearly increasing signal by incrementally binning lines after flat illumination. Isolated bright pixels can easily be created by suppressing serial clocking for all except one in every N pixels read out, without any change in cadence throughout the line.

```
Illuminate to ~1% of well capacity.
For i = 1 to  $\sqrt{2 \cdot \text{NumRows}}$ 
  Shift i lines into serial register;
  For j = 1 to X_size/N {
    Read normal pixel;
    Read N-1 pixels pulsing SW but not serial clocks }
  Flush (Xsize - X_size/N) pixels }
```

5. NOISE VERSUS SPEED

The measures adopted above have minimized the pixel overheads as follows:

- RG pulse width = 80 ns
- RG feedthrough settling time 80 ns for 0.05% pixel-to-pixel memory.
- Charge dump settling = 160 ns for less than 0.05% increase in non-linearity.

The remainder of the pixel time is then divided equally between integrating on the post-reset and signal levels. Figure 11 shows read noise versus pixel time, for all 64 channels in ZTF. The read noise histograms are shown in Figure 10.

The incremental binning method[3] is used for conversion gain measurement (in addition to linearity), because it is very fast, and because it is immune to the brighter fatter effect[4]. We found that the variation in conversion gain from channel to channel was significantly lower than the 0.25 e-/ADU on each channel's conversion gain measurement. This is illustrated in Figure 9 where for each channel a point is plotted for which the X axis is the average of two conversion gain measurements and the Y axis is the average of another two measurements for that same channel. The spread along the positive diagonal is due to the combination of random errors and systematic variations between channels, while the spread in the orthogonal direction is due to purely random noise. The similarity in the width of the distribution in the orthogonal directions shows that random errors strongly dominate systematics, so we chose to average across all 64 channels and all four measurements of each and use this common conversion gain estimate, 6.02 ± 0.02 e-/ADU for all channels in Figure 10 through Figure 12. In reality, systematic gain variations could be as large as ± 0.1 e-/ADU.

5.1 Comparing differential to single sided video

A CCD with differential output has a dummy serial register and output stage that exactly replicate the normal output. These are connected to the same clocks and biases so that clock feedthrough and noise are replicated as closely as possible. The subtraction process serves to reject crosstalk as well as clock induced transients in the video waveform. The differential receiver also rejects noise on clock or bias feedthrough and radiated interference. Noise induced by sample timing jitter in the presence of ripple on the clock feedthrough is reduced to the extent that the differential receiver is able to exclude this ripple. The penalty is the addition of the noise on the dummy output buffer or downstream preamplifiers which is uncorrelated and thus adds in quadrature, so that the differential noise is $\sqrt{2}$ times greater than ideal single sided noise.

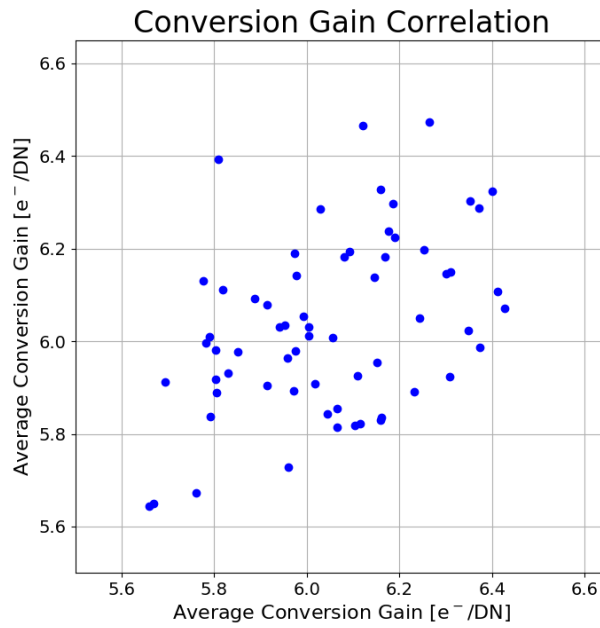


Figure 9: Correlation plot for repeated measurement of conversation gain. For each channel, the average of two measurements is plotted against the average of another two. The weakness of the correlation indicates that the channel-to-channel differences are much smaller than the 0.25 e⁻/ADU rms variation in the entire sample set.

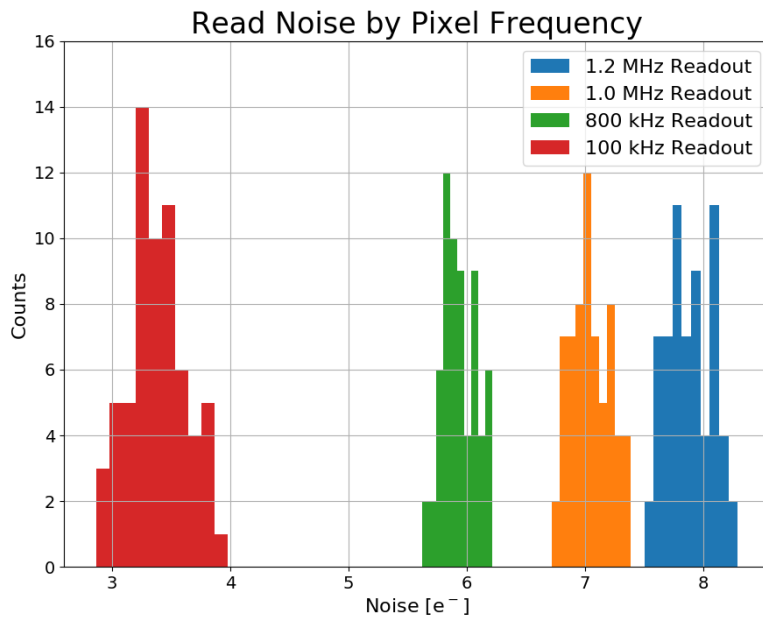


Figure 10: Read noise (in electrons) for the 64 channels in ZTF, at 100 kHz , 800 KHz, 1Mhz , 1.2 MHz for differential CCD outputs and signal path.

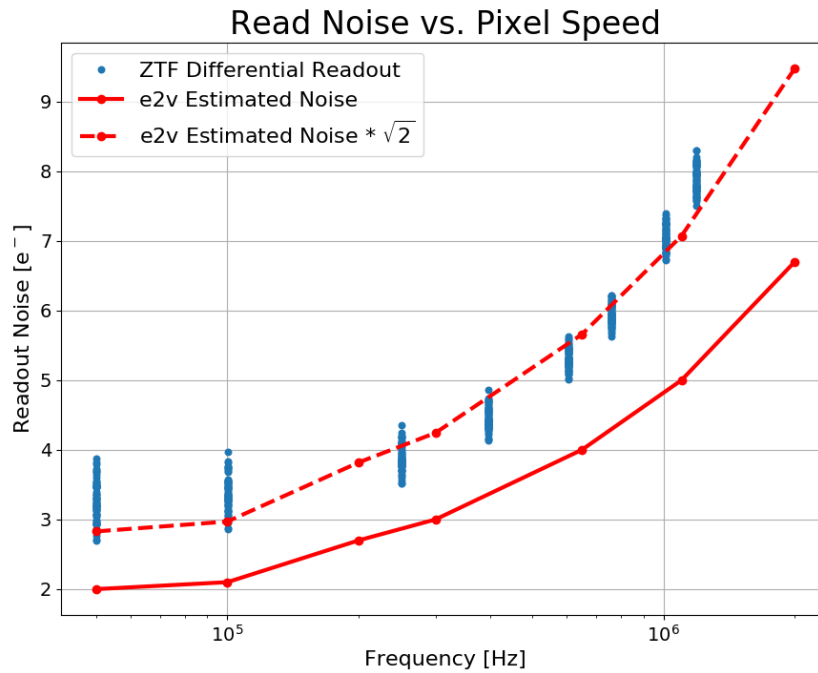


Figure 11: Read noise versus pixel time. Blue dots = 64 channels in differential mode. The red dotted line is the ideal noise for differential mode, which equals $\sqrt{2}$ times e2v's prediction for single sided noise (solid red line.)

Susceptibility to different noise sources varies with the passband shape of the CDS process. At long pixel times the CDS passband is narrower and centered on lower frequencies and thus total noise tends to be dominated by $1/f$ noise in the CCD output FETs, and single sided video is preferred. At high frequencies the additional noise sources listed above may dominate the $\sqrt{2}$ penalty for random noise so that differential signal path produces lower noise than single sided. Figure 12 compares the speed noise curves for differential and single sided cases.

In ZTF, use of a true differential signal path was motivated both by the need to minimize crosstalk and to reject common mode clock transients and noise. There is no evidence of crosstalk even from full blooming along entire columns with well depth of 350 ke⁻. This indicates that cross talk is *well below* the 25 e⁻ shot noise in the sky, and thus lower than ~20 ppm.

While the differential signal path did reduce noise, we were surprised to find that even at 1.2 MHz pixel rate single sided noise was slightly under the ZTF 10 e⁻ requirement. It is plausible that if preamplifier bandwidth was increased to reduce settling time, and cable length was reduced then single sided noise would remain lower than differential to even higher frequencies, but reverting to single sided signal transmission would probably increase crosstalk by several orders of magnitude.

Figure 12 compares noise curves for single sided video at three different settling times (after reset and again after charge dump). While increased settling time reduces sensitivity to clock feedthrough, it consumes time that would otherwise be used for averaging noise. The fact that the shortest settling time (160 ns, orange curve) produces the lowest noise at all frequencies indicates that the extra time made available for noise averaging is of greater benefit than the suppression of noise due to variability in clock feedthrough amplitude or noise due to timing jitter for samples on a slope. This timing jitter is of more relevance at high pixel rates since a larger fraction of the samples occur during transient conditions. However it is only above 500 kHz that the noise induced by clock transients exceeds the $\sqrt{2}$ noise penalty for differential signal processing. This is rather remarkable given that the clocks travel to the dewar over a 2m cable and the video returns over a 3m cable. The only measures taken are mild pre-amplification by a factor of only 1.3 in the dewar, with the same op-amp driving a twinax cable with grounded shield and reference side though 100 Ω series resistance at the

driver to match the characteristic impedance of the line. The input to the differential receiver has high impedance producing a reflected wave that is absorbed at the driver.

The sharp turn up in noise at high pixel rates (red and green curves) in Figure 12 is caused by lack of sample averaging where too large a fraction of the pixel time has been allocated to signal settling.

Another surprising feature in Figure 12 is the fact that even though, at mid to high frequencies, the differential noise curve is very close to $\sqrt{2}$ times the e2v prediction, the single sided noise does not fall on the e2v prediction. This suggests that the single sided noise is contaminated by a common mode component such as would occur if a bias or clock voltage was noisy and coupling equally to both signal and reference sides. Since the differential path suppresses this common mode noise source, to safely exceed the system requirements we did not pursue this further, though in principal it should have been possible to do even better at low frequencies.

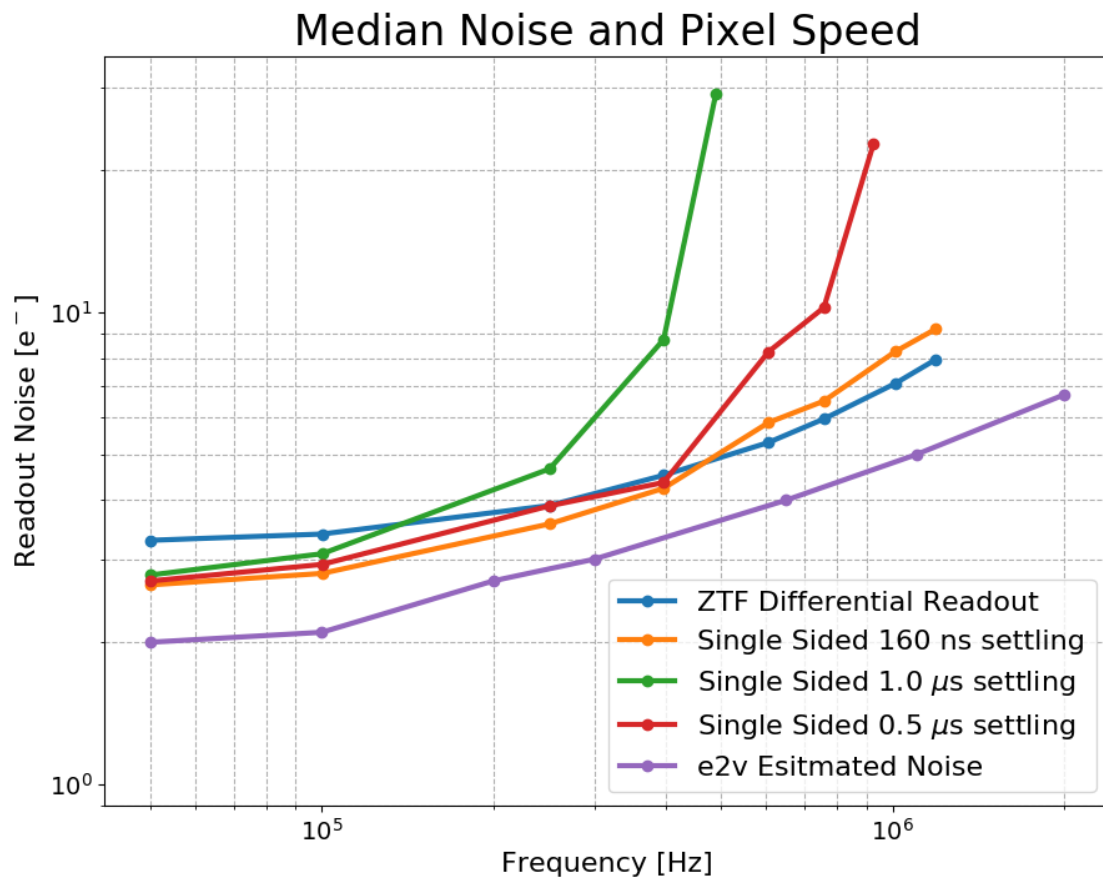


Figure 12: Median speed-noise curves for all 64 channels in ZTF for differential (blue) and single sided video (orange). RG and SW settle times are 160 ns for differential, but must be extended to 1 μ s for single sided to achieve the noise claimed in the data sheet (green curve) at low pixel rates. Differential video achieves superior noise performance above 300 kHz pixel rate.

6. CONCLUSIONS

Coincident triangular serial clocking is useful to minimize clock currents and feedthrough to the video, allowing serial clocking to be spread across the entire pixel, thereby eliminating this source of pixel overhead. Locating the RG driver

close to the CCD proves to be beneficial since it allows very short RG pulses, saving time and reducing the energy in the reset feedthrough pulse. Differing slopes on the SW pulse work well. Generation of SW close to the CCD would be helpful for reducing ripple.

Our implementation could have benefited from a preamplifier with greater large-signal bandwidth to speed up settling on both edges and thus allow more time for signal averaging to beat down noise, without risk of non-linearity. Nonetheless we have succeeded in exceeding both the readout speed and noise requirements for ZTF, achieving 8 e⁻ read noise at 1.2 MHz pixel rate, while also achieving 1% non-linearity (peak to peak gain variation) [3]. (All values are medians among 64 channels.) Differential CCD outputs, with preamplifiers in the dewar driving twinax transmission lines, impedance matched using series terminations were essential to achieve this performance given the 3m video cables from CCDs to controllers. We find that digital correlated double sampling, with differential signal transmission, provides lower noise than single sided video for pixel rates greater than 500 kHz. At low frequencies (50 kHz) the same signal chain with reference side clamped to ground delivers only slightly higher noise (2.5 e⁻) than that advertised in the data sheet for single sided video (2 e⁻), in spite of the 2m clock cables and 3m video cables.

REFERENCES

- [1] Smith, R., Dekany, R., et al. "The Zwicky Transient Facility Observing System", Proc. SPIE 9154 (2014)
- [2] Bilgi, P. and Smith, R., "Concurrent Clocking" Proc. Scientific Detector Workshop, Baltimore (2017)
- [3] Kaye, S., Smith, R., Mao, P., Cunningham, G. and Greffe, T. "CCD linearity measurement by incremental binning." Proc. SPIE 10709 (2018)
- [4] Bredthauer, G. "Archon: A modern controller for high performance astronomical CCDs," Proc. SPIE 9147, 91475B (2014)
- [5] Astier, P., Antilogus, P., Doherty, P., Guyonnet, A., and Regnault, N., "The brighter-fatter effect and pixel correlations in CCD sensors," JINST, 9C03048 (2014).