How to Drive the LTC2387 part II

A Driver for Imaging (short version)

The circuit shown below works great. You should try it.

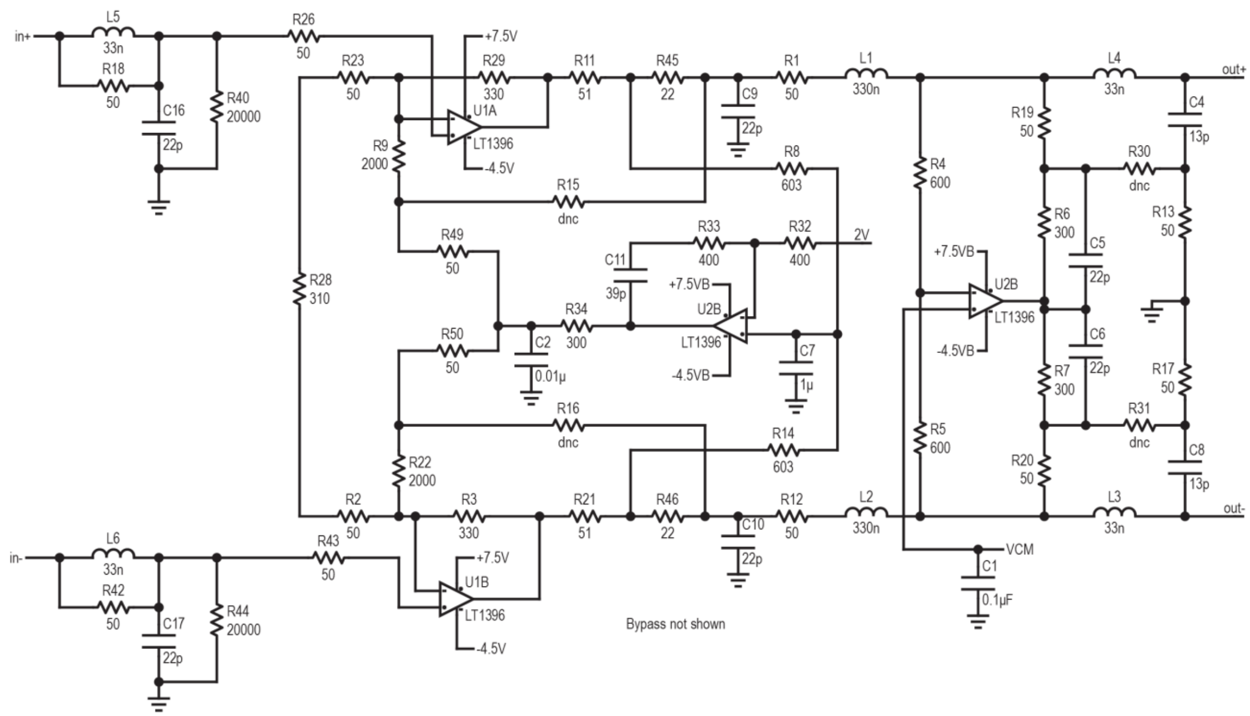

This is intended for nominally 0-3V single ended or differential signals, from CCDs, CMOS image sensors, or other similar signal sources. As shown, SNR is 88 dB. With the AD8008, in place of U1, and a few other minor changes, it produces 92 dB SNR, settles to 0.02% in one pixel at 15 Msps. See the long version of this article for more detail. Stable bias may be required at in- to optimize use of the input range.

How to drive the LTC2387 part II

A Driver for Imaging (long version)

Let me begin by stating that this driver is not necessarily just for imaging applications. By imaging, we actually mean: any signals that involve synchronous steps to DC levels. Or more precisely, abrupt steps followed by a relatively static dwell at an arbitrary voltage level, preferably with some significant portion of a sample period to settle before sampling. These levels may be pixels. This can mean, for example, CCD field or linear array sensors, optical, Xray or infrared image sensors, or in fact multiplexed signals. In the event that this driver is preceded by multiplexers, spectral power distribution of the signals selected by the multiplexer should be preponderantly below 1 MHz, at least in the period some 10-15 nsec prior to sampling.

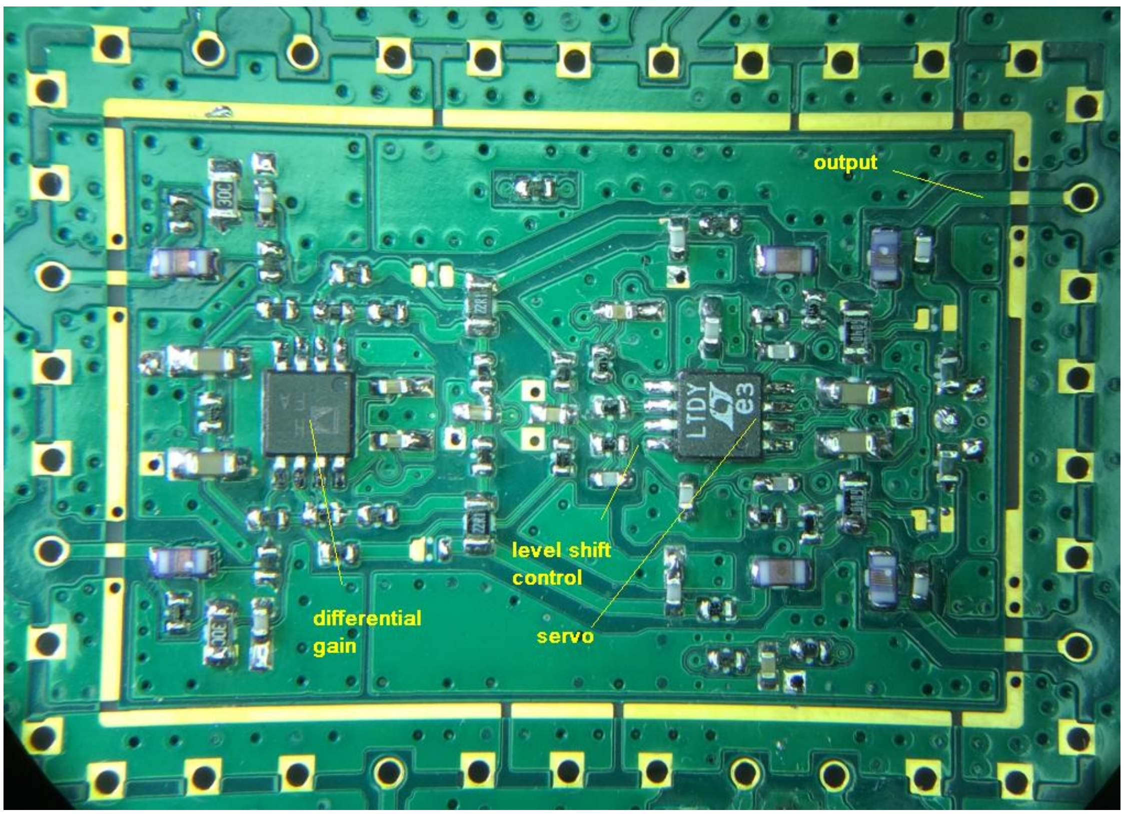

Figure 1 Prototype with two dual current feedback amplifiers from different vendors

The original intent was to use the LT1396, for both U1 and U2, but lower noise can be achieved with other amplifiers. The benefit of this higher SNR depends on the characteristics of the sensor, and to some extent on the nature of information required from the sensor. This circuit was tested with the AD8008 several months prior to the time of writing.

At the time, this other population was not intended to be published. It was done for investigative purposes only. Other current feedback amplifiers (CFAs) from various vendors have also been tried, as they have been in part I, and part III. But only one vendor's parts produced meaningful improvement over the LT1396. Now, it seems appropriate to release it.

This topology is related to the July 2013 article "Near noiseless ADC drivers for imaging". That version was published for the 20 Msps LTC2270, with its 2.1V p-p input range. I would suggest reading that article as well if better understanding is desired. It may be appropriate to interrupt this and read that article if this becomes confusing.

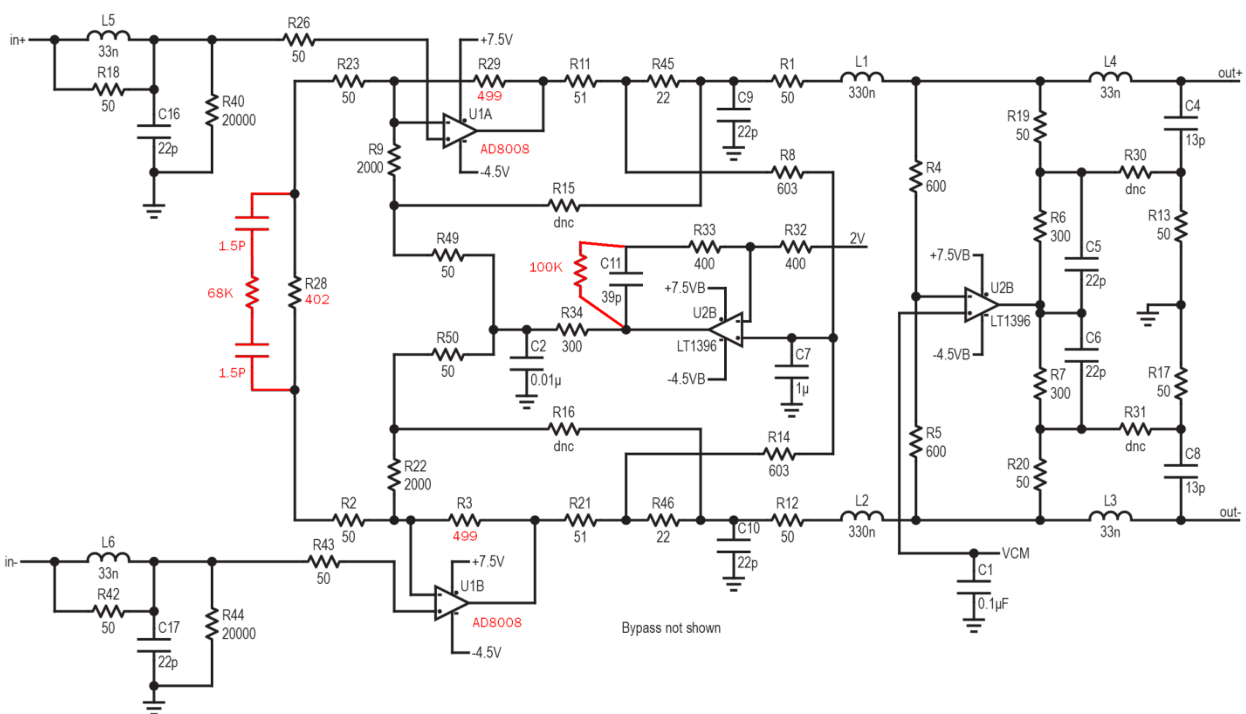

Figure 2 Values in red reflect changes from the LT1396 version to the 650 MHz CFA from vendor A

Components in red are suggested additions not present on the layout

The difference between this present circuit and the LTC2270 version is that this version must provide gain, in the input stage, to achieve adequate signal swing for the 8V peak-peak input range of the LTC2387. This assumes that the "video" signal is in the commonly encountered 3-5V peak to peak range. The signal level at the outputs of the first stage must be greater than the input range of the ADC as it is attenuated in a common mode servo stage prior to the ADC. The attenuation in this version is only 3.8dB, less than in other cases where I have used this approach. This is because the undistorted signal swing at the output of the driven input amplifier is limited to about 8V p-p on a 12V supply. The driven amplifier, assuming single ended drive, has to deliver almost twice the excursion of the

complement amplifier. If driven with a nominally balanced differential signal, and this version differs from the original article in that it is differential, the input stage gain could be higher, and SNR would improve. This is partly as noise current would be less significant, and partly due to greater attenuation in the common mode servo. The suppression of the common mode translates single ended drive, or in this case, rather poorly balanced differential drive, to relatively good differential drive. The original article used the LT1395 family as a unity gain input buffer, taking the power output, the actual power used to drive the ADC, directly from the inverting input, as if it were an emitter follower. This could possibly be done in precisely this fashion with this board, if the video signal were 10 Vp-p differential. It would require somewhat different population. The previous topology translated the 4-5V p-p single ended input to 2.1V pk-pk differential. The LT1396 was performing essentially as a pair of complementary emitter followers within a feedback loop. The output of the amplifier was only used to close the loop, via the minimum 402 ohm FB resistor.

The LT1395 used in that fashion, is quite low noise, and very fast settling. In producing gain however, its 4.5 nV/√Hz input noise voltage, and the inverting input noise current of 25 pA/√Hz are the dominant noise sources. Unfortunately, lower noise amplifiers, and rail-rail output amplifiers, often promoted as suitable for imaging drivers, are often too slow in settling.

Pattern noise and settling

Slow settling will translate any timing variation, any periodic clock feed through, any variable charge injection, and almost any coupling artifacts between pixels, into what is commonly referred to as "pattern noise". This is often in the form of diagonal lines, sometimes with other more interesting, but nevertheless unwanted effects. The patterns often give clues as to the origin of that variation. A deterministic variation in clock feed-through, or charge injection in the CCD will produce similar effects to timing variation if settling is not complete. The same situation would exist in multiplexed signals with variation in timing, or with variable charge injection. The timing variation may originate in processors, in timing engines or in FPGAs, rather than the sensor itself. It can originate in crosstalk mechanisms in poor PCB layout design, or of course in power supply design as well.

There have been cases where it was asserted by imaging customers that slow settling should not be an issue, because it was believed that there were no disturbances between the pixels and no timing variation. In retrospect these seemed often to be those cases where pattern noise was later found to be a problem. In some cases however, the pattern noise was undoubtedly the result of ringing of a poorly chosen differential amplifier unsuited to low gain, high slew applications. That can also translate timing variation into pattern noise as it is another form of memory.

An amplifier that does not settle adequately will also appear to have poor PSRR, as the power supply will have some effect on slewing and settling. This can be misinterpreted as being simply a power supply issue. Some amplifiers have asymmetrical slewing, which, combined with a high dv/dt signal source, and slow settling, or even filtering, may translate clock feed-through variation, charge injection, and other inter-pixel disturbances into discernable artifacts in a flat field. In some cases, it is advisable to have some band limiting before such an amplifier, even though one would expect it to exacerbate slow

settling. Consistent settling time, without approaching slew rate limiting would be easier to correct in software, assuming a model of the simple settling behavior can be computed.

At the root of some of these pattern noise problems is the fact that the settling time that is available to the amplifier is often less than anticipated, and often less than 50% of a pixel interval. At 15 Msps, this means only 33 nsec. As an example, the LT6237, sometimes unwisely promoted for imaging, will need about 1 usec to settle to 18 bits, in a 4V p-p scenario. This is arguably too slow by a factor of about 30. Of course, the appropriate amplifier must also be low noise.

At 15 Msps, and for approaching 18 bits, this appears to leave CFA's as the only apparent choice. The LT1396 settles to 0.1% in 25 nsec, and is not characterized in the data sheet beyond this. In any case, this is not quite good enough in terms of settling, nor in terms of noise. The settling time of the AD8008, to 0.1%, is 18 nsec (2Vp-p), and the AD8002, is 16 nsec. Note that I am only using settling to 0.1% for comparison purposes, and yet settling to even 0.01% is arguably not adequate.

The circuit as shown above settles to within 0.02% in one pixel interval, after a full scale step, at 15 Msps, and within 0.01% at 10 Msps.

SNR

The settling of the amplifier is not however, the only factor in choosing an amplifier. Inverting input noise current is dominant in low gain (<7dB) applications of all these CFAs.

The LT1395(6) at 4.5 nV/√Hz, and 400 MHz, has often been my workhorse. In the circuit below, with LT1396 as U1, we see SNR of 87.5 dB. This is 4 dB better than the 83.5 dB of the previous LTC2270 version as published. It is not however better in terms of density at the input of the ADC; as it is the 12 dB increase in input range that is responsible for the 4dB increase in SNR.

However, using the AD8008, we see 90.7 dB SNR; and with the AD8002, SNR of 91.5 dB. There is some trade-off between settling and SNR. Versions with somewhat better settling are worse in terms of SNR by about 1 dB.

Thermal tails

Another factor that appears to be an issue is that of thermal effects in the gain setting resistors, and in the CFAs themselves. We had observed differences in settling performance that appear unrelated to the settling performance of the amplifiers, and determined largely by resistors. Thermal tails in the amplifiers do eventually manifest themselves at the levels mentioned, once the resistors are good enough. Low TC resistors would seemingly not be necessary based on a simple calculation; however, high current density hot spots at the end of laser trim lines have been a problem in the past. Resistors with trimmed slots 50% of the way across the element, or more, have been found to cause distortion. The most sensitive resistors are the gain setting resistors R3, R28 and R29. R6 and R7 are the next most likely to participate in producing thermal tails. However, R11, R45, and R1 and could be additive if they were the same device, and had similar effects. They do not all have the same effect, FB resistors R29, R3,

and the shunt elements in the common mode servo will increase the gain with positive TC, and the series elements, and R28, the reverse. If in prototyping, they all happened to cancel, that would be unfortunate, and concluding that it performs well would be a classic mistake.

All of the three CFAs that were tested, in the end, appear to exhibit a similar thermal tail, likely associated with the second transistors. This similarity initially, to me, suggested that it was not the amplifiers. After much experimentation, and discussion, it was decided that they are seemingly very similar. There is a first order compensation network, *that I did not make provision for on the PCB*, that consists of 0.75 pF, in series with 68K ohms, placed across R28. This was tested in the form of 1.5pF-68K-1.5pF.

This has a tau of 51 nsec. Although the thermal tail is a more complex function, associated with several thermal time constants, and a thermal lag associated with heating the neighboring input transistors, at which point there is some degree of cancelation, the simple network does a very good of hastening settling, in comfortably less than 1 clock period.

The testing of the linearity of this circuit, similar to the original article, involved a high level square wave at fs/2, or fs/4, combined with a lower amplitude sinusoid, the square-wave at Nyquist emulating pixels. If the timing of the sampling, relative to the square wave, is towards the end of the "pixel", as such, where you would strive to sample a pixel, the distortion should be at a minimum. Sampling earlier is not advisable as this would provide less settling time, and as well, the common mode servo will not have settled, and hence, will not have effectively suppressed common mode, as it does towards the end of a pixel. However, the linearity of the LT1395 version may not in fact be a concern for many imaging applications, and the 87 dB SNR may be acceptable, as it is likely to be dominated by the sensor. In this case, the lower current consumption of the LT1396 may be attractive. The AD8002 has somewhat more distortion than the AD8008 version, but lower noise, so it may be preferable for many customers. However many imaging devices will dominate even the -87dBfs noise floor of the LT1396 version by 20 dB, or thereabouts. With short aperture times, low illumination, low dosage, or small area imagers, the SNR of the driver will not matter, much. In fact, the LT6411, a device derived from the LT1395 family, still on a 12V process, has considerably faster settling time, due to internal feedback resistors, and yet, should produce about the same noise floor as the LT1396 version, and may be used in a new iteration of the above board. There are other amplifiers under consideration.

The present circuit differs from the original 2013 imaging topology for the LTC2270 in that the input stage also contributes to the translation from single ended signaling to differential. As it has a differential gain of approximately 3, the output from the first stage has a common mode component that is reduced relative to the differential component, by about 14 dB. As such, the output common mode servo could remain much the same as the original design even though the input range of the ADC is 4X higher. This higher gain, higher input range version does however, have an additional common mode loop that controls level shift, and reduces the output servo's role in level shift. This level shift however may need some involvement on the part of a microcontroller and a DAC. The prototype had a high gain in the servo, and in the case of dramatic changes in illumination, the servo would have to respond. Reduced gain in that servo may be necessary. (C11 paralleled with a resistor). It is a question

of the video level, as 0-4V or 0-5V do not allow as much headroom. The level shift may be simply a fixed bias if input common mode is relatively stable.

The common mode output servo must operate into a pair of nodes, that in common mode, must not appear as capacitive reactance. The output filters isolate the common mode servo from both the output impedance of the first stage, the level shift servo, and the input capacitance of the ADC, as well as various transmission lines. The filter also reduces aliased noise bandwidth, as well as the frequency content of the transients which the common mode servo must suppress.

The common mode servo in this LTC2387 version, U2A operates into the output filter, via C5 and C6, and part of the network intended to appear as absorptive to the ADC. Without this feature, following the servo with a diplexer stage may compromise common mode settling. The circuit board has provision to move the last absorber from R13 to R30, for example, or split between the two. However, returning the entire common mode transient power to the output of the common mode servo is likely unwise.

L1, L2, L3 and L4, are an integral part of the common mode loop, isolating the controlled node(s) from capacitive loads C9, and C4.

Much as explained in part III of this series, loose tolerance of components in the output servo can compromise SNR and distortion. Rather than repeating details explained in the original article, I would ask that you read those if this article appears a bit short on details.

The use of CFAs with their relatively high input current could potentially be a source of problems for some sensors. The maximum in+ current for the LT1395 is 30 uA. However, many lower noise VFB amplifiers also have quite high input current. The LT6200, for example, has a max bias current spec of -40 uA, and the LT6237 has a 12 uA input current spec, and yet both are too slow. The AD8008, and the AD8002, have 8 uA, and 10 uA max input current. The use of the LTC6268 FET amplifier (5 nV/√Hz) as the input stage could be entertained if sample rates are less than 10 Msps. However, the LTC6268 is limited to a 5V supply, so a differential video signal may be required to be practical.

Testing and Validation

As there is no budget to acquire, or time to devise an 18-bit-linear CCD emulator, nor likely any CCD or light source, for that matter, that is likely to provide us any way of proving that this solution is linear to the extent that we would like to believe, we have no choice but to clarify the case that was made in the original article, and refine it. We did receive comments that imaging customers did not know how to interpret the case we had presented.

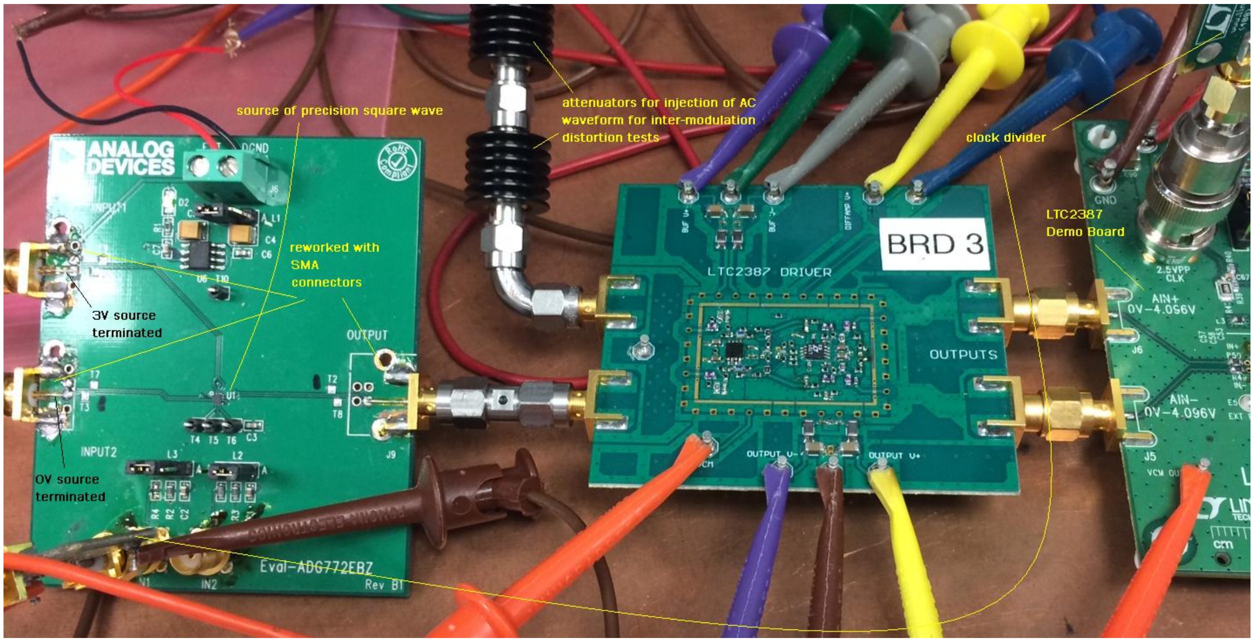

Figure 3 experimental set up with precision square wave source

The test case in the original article was a square wave at Nyquist, combined with a 70 KHz tone. In the presence of the square wave, there were some IM products produced, but no change in the apparent power of the fundamental, at least not out to 3 decimal points. The level of IM distortion in the sinusoid was quite low, in the 15-16 bit level, with the AD8008.

In the present design, using the LT1396, the spur at Nyquist-fin, was at about -65dBfs, or about -45 dB relative to the -20 dB sinusoid. With the AD8008, it was at -98dBfs. That is the same as the power in the quantization error of a 16 bit ADC. I am not saying it originates in the quantization error; it is just comparable in terms of power.

In this new case, the source of our single ended square wave is an ADG772 analog switch switching between two reference voltages, each with source termination for the transmission lines. They are not end terminated at the amplifier other than at high frequency, in R18. This switch is differential, and the unused lines are also source terminated.

In the LTC2387 version of the inter-modulation tests, we have increased the AC frequency to 300 KHz, for reasons related to generators, filters, and DC blocks, but dropped its power level to -20dBfs, such that the square wave drive has near full amplitude. IM products are typically maximal when power is equal in the two tones, but the higher amplitude transient is closer to actual imaging signals. This scenario is intended to emulate near full scale steps in an imaging signal, but using the IM distortion of the weaker sinusoid as a measure of the non-linearity. If there were significant non-linearity, either in the form of compression, or in the form of differential phase error, the 3rd order inter-modulation products would manifest themselves offset by 300 KHz from the square wave. If there were compression, the amplitude of the fundamental would also change.

The symmetrical, bimodal nature of the signal resulting from a sinusoid superimposed on two distinct logic states, for want of a better term, would perhaps not produce significant 3rd order IM if there were the same degree of compression (reduction of gain) in those two states. It would produce products at $2f_{sine}$ below Nyquist, resulting from alternation of the phase of the 2nd order distortion between those two states. So we also tested the combination of 300 KHz sine wave with the 1/2fs square wave toggling between bipolar zero and near either maximum of the transfer function. Compression would produce AM, which would appear as a tone at Nyquist minus f, seemingly the lower sideband, but in fact both sidebands combined as the upper sideband aliases on top of the other. This scenario did not in fact produce more severe 3rd order inter-modulation distortion, so linearity is presumed to be acceptable.

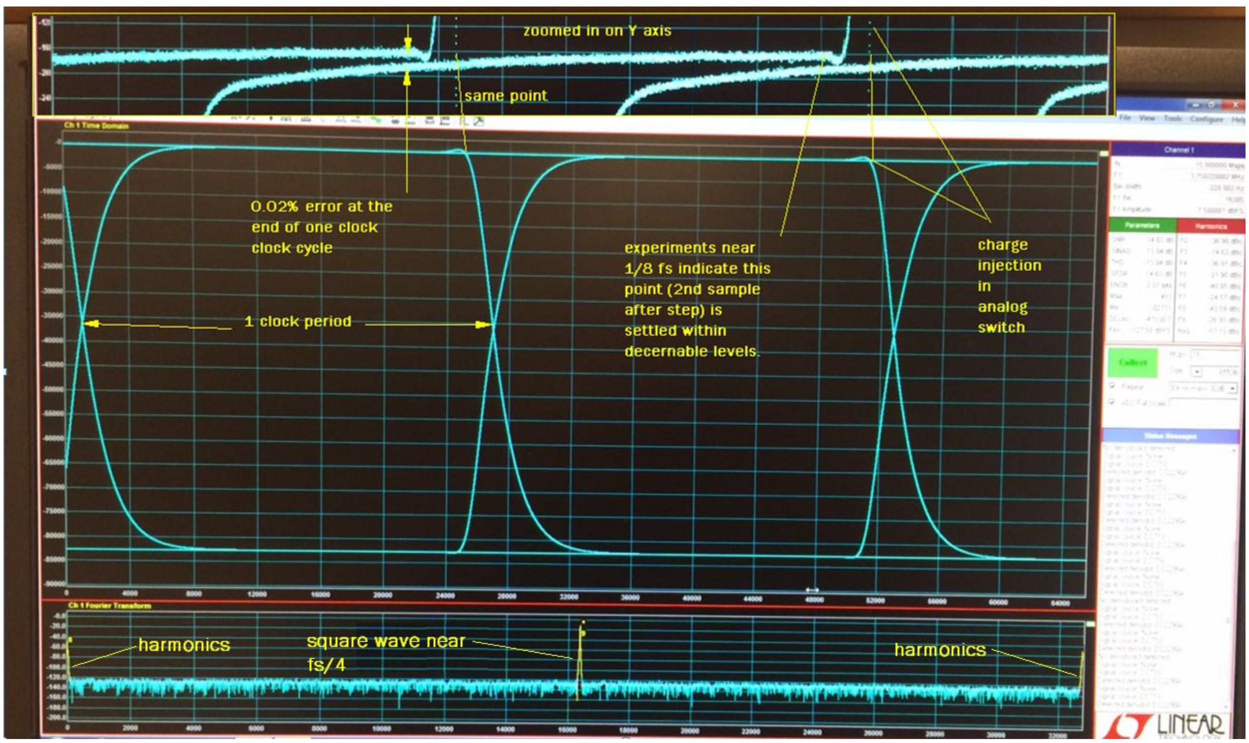

Figure 4 Square wave test near 1/2 Nyquist (3.75 MHz), showing thermal tail in zoomed Y axis (top inset)

Criticism and Defense of complexity

The most vehement criticism of this circuit, and other circuits in this series, is that of complexity. There are several basic truths that impose this level of complexity. High linearity requires high GBWP of the amplifiers, which necessarily means a great deal of noise BW. The input BW of the ADC is 200 MHz, or 27 Nyquist zones at 15 Msps. If sampling were subject to this entire noise BW, assuming it is flat, the noise from the amplifier would be 14.3 dB* than if it were limited to 1 Nyquist zone. *(10log(200/7.5)=14.3 dB) As the amplifier dominates, that would be almost a 14 dB penalty.

This means the output must be filtered to get the benefit of the high SNR. However, you cannot limit the BW to 1 Nyquist zone and get fast settling. Some 3-4x the Nyquist BW is required. There is still an 8.3 dB benefit in limiting the BW from the full 200 MHz BW of the ADC, to 30 MHz.

There are two aspects to settling of a filter, the response to the input waveform (S21 for RF engineers), and the reflection in response to the transients received from the ADC (S22). The time available for settling out in response to the disturbances from the LTC2387 is 1/fs-39nsec, or 27 nsec at 15 Msps. A blighted reasoning would suggest that the 25 nsec settling of the LT1396 is adequate for this scenario. However, the amplifier must settle out interacting with the delayed response through the band-limiting filter, and driving the complex impedance of the filter. In addition, inter-dependent amplifiers result in extended settling time relative to a single amplifier. The differential input stage, composed of two CFAs, is however much less inter-dependent than if it were composed of voltage feedback amplifiers. To visualize this effect, for single ended input signals for example, the signal at the inverting input of the driven amplifier is directly buffered by the input stage, much as a follower, and injected into the opposing amplifier, behaving as an inverting amplifier. This aspect dramatically improves the initial settling in response to fast input steps, but is perhaps less significant in response to reflections into the output stages.

On a happy note however, the filter mitigates the disturbance reflected into the amplifier. Fast settling of the output network is required for both high linearity from a direct sampling ADC, to dissipate any non-linear charge from the sampling process, and for imaging applications, to avoid pattern noise. This typically means a Gaussian absorptive filter. Unfortunately, an absorptive filter is typically 2-3 times more complex than a reflective filter. Optimal damping in a driver such as this would require a filter with an impulse response that settles completely by the required settling time, perhaps taking into consideration the slew rate of the amplifier. If the slew rates are asymmetrical, achieving optimal settling for both excursions is unlikely. Assuming symmetrical slewing, the optimal impulse response of the filter itself would be that with slewing of the amplifier involved rather than an ideal step. However, this would mean a different solution for different excursions if the amplifier is slew rate limited above some amplitude. Hence, the need for some slew rate limiting prior to the amplifier, as determined by C16 and C17. If the slew rate is limited prior to the amplifier, keeping the amplifier out of slew rate limiting, the filter should settle well for any amplitude excursion. Keeping the amplifier out of slew limiting also prevents a large error voltage from developing across the inputs, keeping it out of non-linear operation, and avoiding an extended recovery process.

The common mode servo must operate into a pair of nodes that appear to be a real impedance, rather than reactive, and as a result, the output filter is again slightly more complicated than a simple reflective filter. These two nodes under control by a single amplifier must match, in terms of complex impedance, or it will translate the common mode errors, and noise of this amplifier into differential components, and would be seen by the ADC.

The control of reflections on transmission lines, and other elements involved in isolating the high speed path from a lower speed common mode loop also add complexity. This prototype board was designed to allow us to evaluate the LT1396 vs the AD8008, and other devices, in the MS8 package. There are other devices with different pin-outs that are faster settling, lower noise, and perhaps better suited to producing a compact layout. The use of resistors arrays has been considered, which may produce a more compact layout, and perhaps better mitigation of thermal effects, but those would be the subject of a later article.

If there is interest in this circuit, help is available, and the existing board can be made available in blank or partially populated form. The suitability of this driver must be evaluated for each sensor and application. However, the dynamic range, and fast settling of this topology may eliminate the need to customize for each type of sensor. That statement assumes that gain is low enough to accommodate the highest signal levels expected.

This is a work in progress. There are some other topologies under consideration that should improve settling, SNR. But as you may expect, they are somewhat more complicated.