

CS448 Isolated Channel Oscilloscope



Bart Schroder, of Cleverscope, recounts the saga of the CS448 isolated channel oscilloscope project.

Way back in Sep 2011 Dave Jones of EEVblog fame interviewed me at ElectroneX 2011 in Melbourne and asked what's next? I glibly answered we planned to make an isolated channel oscilloscope. After a long road taking us more than 5 years we first showed our CS448 at PCIM in Nuremberg, Germany in July 2017. This year I'm back at ElectroneX in Sydney, and maybe I can have another conversation with Dave, only this time my hair will be greyer, and I'll be talking about the CS448 instead of our CS328A.



I thought I'd share some of this journey with you. After all, this magazine is about electronics projects, and the CS448 has been quite a project.

Why and what?

This all started when I was designing Variable Speed Drives (VSDs) for electronic motor speed control. A VSD uses three half bridges to generate a three phase signal which rotates a three phase motor. You can also use just two half bridges to control a stepper motor, or permanent magnet DC motor. This is called a full bridge. Each half bridge uses two FET or IGBT switches to switch a load between a negative and positive bus. Fig 1 shows a full bridge which switches Z_{Load} alternately between V_{bus+} and V_{bus-} and then V_{bus-} and V_{bus+} . If Z_{load} is inductive - as most motors are - current increases or decreases slowly when a voltage is applied. By varying the pulse width of the switching, we can vary the current into the load in a sinusoidal manner, and make the motor rotate. This is known as Pulse Width Modulation (PWM).




Fig 1. A Full bridge using FETs and showing probes.

When you design something you need to make sure it is right, otherwise Murphy will get you. Commonly you will put it through its paces doing whatever real world thing it has to do - this is functional testing. But before that point you want to know the internal parts are doing what they need to do, and you have to make allowances for temperature, component variation and drift. At this point you will probably use an oscilloscope to measure the signals at various circuit nodes and make sure they are correct, according to your design calculations.

Non ground referenced signals

But, and it's a big but, for the Half Bridge, some of those signals are not ground referenced, and pretty much every scope is ground referenced. Look at the purple and green probes which are measuring the gate drive V_{g1} and V_{g2} . These are referred to the high side FET source, and that source is switching rapidly between V_{bus-} and V_{bus+} . You could 'float the scope', but that only gives you one channel, and besides, it's dangerous, and there might be quite high capacitance or inductance to ground through the power supply. FETs (including GAN and SiC varieties) can switch in 10 - 100ns. Take the worst case 10ns, with 1pF capacitance, 680V bus, and you'll get $i = CV/dt = 1p \cdot 680V / 10n = 68 \text{ mA per pF}$. With 100 pF you get 6.8A! That current will flow into the scope probe and disturb the measurement, so you need low capacitance.

A second way is to use a differential probe - but even good ones have poor Common Mode Rejection Ratio (CMRR) at high frequencies. As an example the Tektronix P5200A has a CMRR of 30 dB at 3.2 MHz. Now 3.2 MHz is a rise time of $t_r = 1/\pi f = 100ns$. Lots of modern switches are faster than that. If you had a 680V bus (typical for three phase = $2 \times \sqrt{2} \times 240V$), the probe will generate a spurious response equal to -30dB of 680V = $0.032 \times 680 = 21.8V$. That obliterates pretty much anything interesting. Why

measure the gate drive you say? Well after blowing up quite a few IGBT modules, I can tell you that the gate drive has to be right.

Adequate signal resolution

Next you might want to make sure that there was not too much power loss in the transistors. A reddish look followed by a loud bang is not the best way to test this! A better way is to measure the voltage across the transistor while measuring the current through it, and multiply to get power. Referring to Fig 1, you want to measure V_{DS} across the transistors, and you want to measure the current by looking at the voltage across R_{IL} and R_{IU} , which are typically low value (1-10mOhm) coaxial resistors. V_{DS} will transition

between the saturation voltage, say 0.2 - 3V, and the off state voltage, say 680V. So you really need a resolution of 0.1V, or 1 part in 6800. Your average scope is 8 bit, or 1 part in 256. Assuming the display fits to 680V, ADC resolution is 2.6V, and noise will mean minimum real resolution is likely to be > 5V. Not very useful.

So we need an isolated scope with good CMRR at high frequencies, and high enough resolution to do 1 part in 6800. Onward ho!

Key Specifications

The key specifications we came up with were:

- 1 part in ± 8000 resolution. This requires a 14 bit ADC, and very low noise. Resolution will be 0.1V in 800V. We decided on less than than 2 lsb (least significant bits) RMS noise, which gives a real resolution of 0.2V in 800V - just enough.
- Less than 1% error when measuring current in a current sense resistor with a 1x connection while slewing 680V. This is the same as less than 10% error when measuring a gate drive with a 10x probe.
- Channel capacitance ≤ 10 pF to limit common mode capacitive current.

Design Thinking

1. Noise and ADC

The first question was where to put the ADC - on the isolated side, or the non isolated side. The noise floor sets the dynamic range and is proportional to the number of components in between the input and the ADC. Fewest components means it needs to be on the isolated side. Starting with the ADC we wanted 14 bit, low power (because we have to get the power across

the isolation gap), and a method of sending the digitized signal to the non-isolated side. The only realistic transfer method is serial, and the only method that allows synchronization is JESD204B, a standard for ADC and DAC data transfer. After a search we settled on the Intersil ISLA214S50, a 500 MSPS 14 bit ADC which could transfer all the samples over two serial lanes, at 4.375 Gbps per lane using compression.

2. Front end ranges and parts

The best parts we could find for the front end was the Analog Devices ADA4817, a 1 GHz BW FET input op-amp with only 4 nV/ $\sqrt{\text{Hz}}$ voltage noise, and low well defined distortion and slew rate (unlike a discrete solution). We matched this with the ADA4937 differential ADC driver, only 5.8 nV/ $\sqrt{\text{Hz}}$ output noise, 1.9 GHz BW, and -102 dB THD. We talked to AD, and discovered that we could mux the ADA4817, so the plan was to have two ranges, and mux between them to keep everything as simple as possible. We'd make the ranges ± 800 mV and ± 8 V. With these two ranges we could use a 10x probe to get ± 8 V or ± 80 V with full bandwidth, and the ± 800 mV range would cover the low level current sense application with 100uV resolution. A 100x probe would give us a ± 800 V range, and we'd be covering pretty much everything.

3. Fibre Isolator

We did a market search looking for the way to transmit the two serial data streams to the FPGA we would control everything with. Eventually we found an English company, Advanced Fibreoptic Engineering to make us TOSA/ROSA pairs with a holder and fibre link. We paired these with TI ONET4291VA and PA VCSEL TOSA driver and Limiting Amp ROSA receiver. Sounded simple!




Fig 2. Initial Digitizer channel PCB

4. Clock Generator

We wanted all 4 channels to use the same clock source- but that would have meant another fibre channel, and besides the jitter on a fibre channel is way too high, and our 14 bit ADC would become an 8 bit one. So we settled on using a programmable clock oscillator - the Silabs Si598 as the low jitter clock source. The idea was that we could measure the channel frequency from the serial data coming back, and adjust the clocks to make them all the same.

5. Power Supply

As we have seen we need low capacitance between the channel common and the real system ground. The capacitance is set by the power transformer capacitance, and the capacitance between the channel and the chassis, and the scope probe and the surrounding environment. We could control the power transformer, and the channel placement. We settled on a MAX13256 H bridge driver for isolated supplies, and the companion Halo TGMR-501V6LF low capacitance (10 pF) transformer as specified in their EVM. The transformer was UL/EN60950 approved, which we needed.

6. Control

We decided to use a cheap as chips STM8 processor for all the control. It would communicate using an opto isolated serial link with the system FPGA.

7. 50 ohm termination

Most scopes offer 1M and 50 ohm terminations. So we put in a relay to switch in the 50 ohms. You need a relay, because the contact impedance and capacitance is small.

8. Shielding

We knew we'd need shielding to stop noise from all those high frequency FPGA signals from getting into the sensitive analog front end. So we planned a shield that would be U shaped with fingers which could be pushed down through slots in the board, to make a shield right round the board, and mate with a ground plane on the main board the digitizer would be plugged into.

Prototype test results

We found a lot wrong! I guess only by building something do you learn the harsh lessons.

These were the worst:

- The ADA4817 has bugs in it - you could not mux as specified, and when the device was disabled, it dragged the input connections to -5V instead of them being high impedance. I finally confirmed this with the designer. Our two range design was unusable.
- The MAX13256 H bridge driver + transformer generated large common mode transients on the isolated ground which added to any signal being measured. Our power supply design was unusable.
- The 50 Ohm relay and 50 ohm load resistor could not be turned off fast enough when the 1kV maximum input voltage was applied to the inputs, with the resistor and relay disappearing in a puff of smoke. We had to abandon a 50 Ohm load option.
- The Si598 drifted at the rate of about 15 Hz/sec, which meant that long duration captures would be out of synch after 60 msec or so. It also meant our intent to do Frequency Response Analysis (FRA) would fail. We needed a better clocking system.
- The ISLA214S50 ADC lost gain/offset alignment between the two internal ADC's used to achieve 500 MSPS, and became horribly non-linear if the signal exceeded the digitizer range by even 1mV. We could not use the ADA4937 Differential Amp, because overloading is very common when a user is looking at a portion of a signal. We needed limiting.
- The shield design was fine at stopping noise but useless at achieving a good CMRR. Because the shield was referred to the system earth, any capacitance between components on the board and the shield injected current into the front end circuit polluting the measured signal. We needed a better shield design.
- We had different RC time constants between the AC and DC paths in our two ranges. These generated slowly rising or falling pulse responses when fed using 10x probes.

Of these the power supply turned out to be the most problematic. Making an isolated supply that injected only microvolts into the system being measured became one of the most difficult challenges.


Final Digitizer Design

Our final digitizer design (and we have done three major versions, with two tweaks to the last version) is based on lessons learnt!

The difficult bits


Isolated power supply

This part was a year in the making. Lessons we learned: **everything symmetrical!** You needed a very symmetrical power switch, controlled equal slew rates on the power switch edges, and a symmetrical power transformer. The transformer needed to be balanced and center tapped with minimal inter-winding capacitance. Our final design has two very separated windings with very low capacitance between the windings. The windings are wound bifilar so that each half of the winding is symmetrical to the other.



Clock System


The only way to have all the clocks be the same was to have a common clock. The only way to do that was to put it with the FPGA as a master, and distribute it to all the channels. This approach means we can also synchronize more than one unit together. The only way to have very low uncertainty about the clock phase was to use another fibre, so we needed one coming back. We used a fixed 100 kHz clock from the FPGA board, and multiplied it up to our 500 MHz ADC clock using a Silabs 5344, a truly magical device. This device reduces jitter to 0.1ps - good enough for 85 dB SNR at 100 MHz and does not compromise the ADC SNR, as shown on this TI graph:



The Si5334 generates an output precisely in phase with the 100 kHz master clock, meaning all 4 channels (and any downstream units) are synchronized together.

Fibre Isolator


Our two way fibre isolator was no longer good enough. Murata in Japan came to the rescue with Fibre Optic Transceivers (FOTs) and interconnect fibre. These dual channel bidirectional 10 Gbps units have only 60ps edge uncertainty unit-unit. This meant we could do a good job of synchronizing our 2ns clock periods (our final system achieves ± 160 ps phase variation between channels). Here are the FOTs and the interconnect fibre, showing the 1kV isolation gap and the isolation power transformer:



With bidirectional fibre we could also have bidirectional communications with the digitizer, and can locate the digitizer remote from the main unit.

Shielding

The shielding is absolutely crucial to getting good CMRR. Key is that common mode current must flow along the outside of the shield to the common point, which is the centre tap of the isolation transformer. From there it flows through the transformer interwinding capacitance to the case.



The shield goes right round the PCB, and solders to the BNC. It has a heatsink for the high power ADC and clock generator. The plastic cover is for 1 kV isolation.


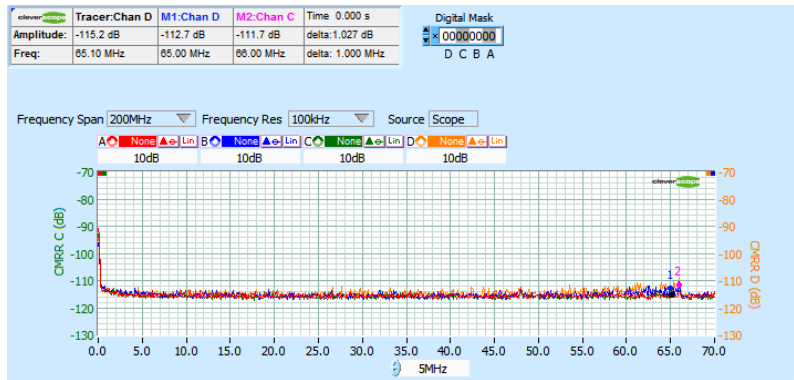


Fig 3. Final Digitizer channel PCB

The final design is shown in Fig 3. We changed the front end to have two ranges taking into account the ADA4817 deficiencies using RF photomos switches (the two white packages). We used a clamping LMH6553 differential ADC driver to avoid saturating the ADC, and we got rid of the 50 Ohm option.

Some Results



It's always good to end with some results. Here is the CMRR for Channel D of Serial Number EQ10019. It's below -110 dB right up to 65 MHz!

-110dB on 680V is 2mV. With a 10x probe that's 20mV. Assuming a current probe, 2mV as 1% means 200mV full scale. Sounds pretty good to me.

And here is the direct measurement of two high side gate drives, where the switch common is attached to the switched signal slewing 500V in 8ns, as in Fig 1. We can clearly see the Miller plateau (which is where the gate voltage stops rising as the gate charges up) on Gate 1, and the results of droop caused by capacitive voltage divider of Cgd and Cgs acting on V_{DS} as the switch goes high (red). Similarly, on Gate 2, we see a pulse caused by the divider as the Switch (in blue) goes low.

I doubt that anyone has seen graphs like these for an actual working bridge before!

Come see us at stand A13 at ElectroneX, 5-6 Sep, 2018

