

TMS Electronics

Tom LeCompte (SLAC)



U.S. DEPARTMENT OF
ENERGY

Office of
Science

The SLAC logo is the acronym "SLAC" in white, bold, sans-serif font, centered within a solid red square.

SLAC


Outline

- Description of what we put in the CDR (following the signal)
 - Quad Counter Mounting Boards
 - Digitizers
 - Data Concentrators & Timing
- Problems and Complaints
- Some Options (not following the signal)

Following the Signal - QCMB

- Quad-counter Mounting Board: the analog SiPM signal gets placed on an HDMI connector
 - Also some LED flasher controls
- HDMI cable runs up to 1 meter to digitizer box
- Why HDMI?
 - It's commercial, widely available, with lots of downward price pressure
 - It has 4 pairs intended for data – we have four channels per quad-counter
 - Locking versions are available

Base cost: \$540,701
Contingency: 34.6%



The mu2e equivalent

Essentially, the mu2e CMB board re-layed-out to fit the DUNE TMS footprint.




Quad-counter footprint



Following the Signal – Digitizer Boards

Base cost: \$1,154,522
Contingency: 58.6%

- Accepts signals from QCMBs and digitizes them using the TI AF5807 ultrasound chip
 - Several variants of this chip exist: mu2e CRV uses one.
 - Yes, I know this chip is single-ended. I don't like it either.
- Signals are sent to an Artrix-7 FPGA which zero-suppresses them and sends the data out on Display Port cables
 - Channel/energy/time + packet overhead
- Why Display Port?
 - It's commercial, widely available, with lots of downward price pressure
 - It has 4 pairs intended for data – 3 out and 1 in fits naturally to 48 counters
 - Timing could come in on the 4th channel
 - I don't like using the same connector twice if there is any risk of miscabling



Following the Signal – Data Concentrator/Timing Boards

- These each support 5 digitizer boards
- Take the signals from the Display Port Cables
- Use a Zynq FPGA/SoC to suppress noise, format the TCP/IP packet the DAQ expects and send it to a switch
 - This is a commercial switch, most likely RJ-45 + SPF+ uplink
 - We hope to fit on a 10 GbE optical link. We could write out more SiPM noise, but why?
- Right now this is on a custom motherboard in a 19" enclosure
- This board provides power and timing to the digitizer boards

Base cost: \$211,523

Contingency: 75%

Contingency is 50% from the estimate + 25% in case we can only support 4 panels instead of 5.

We know the sum of the cost of the last two boards better than we know the individual costs.

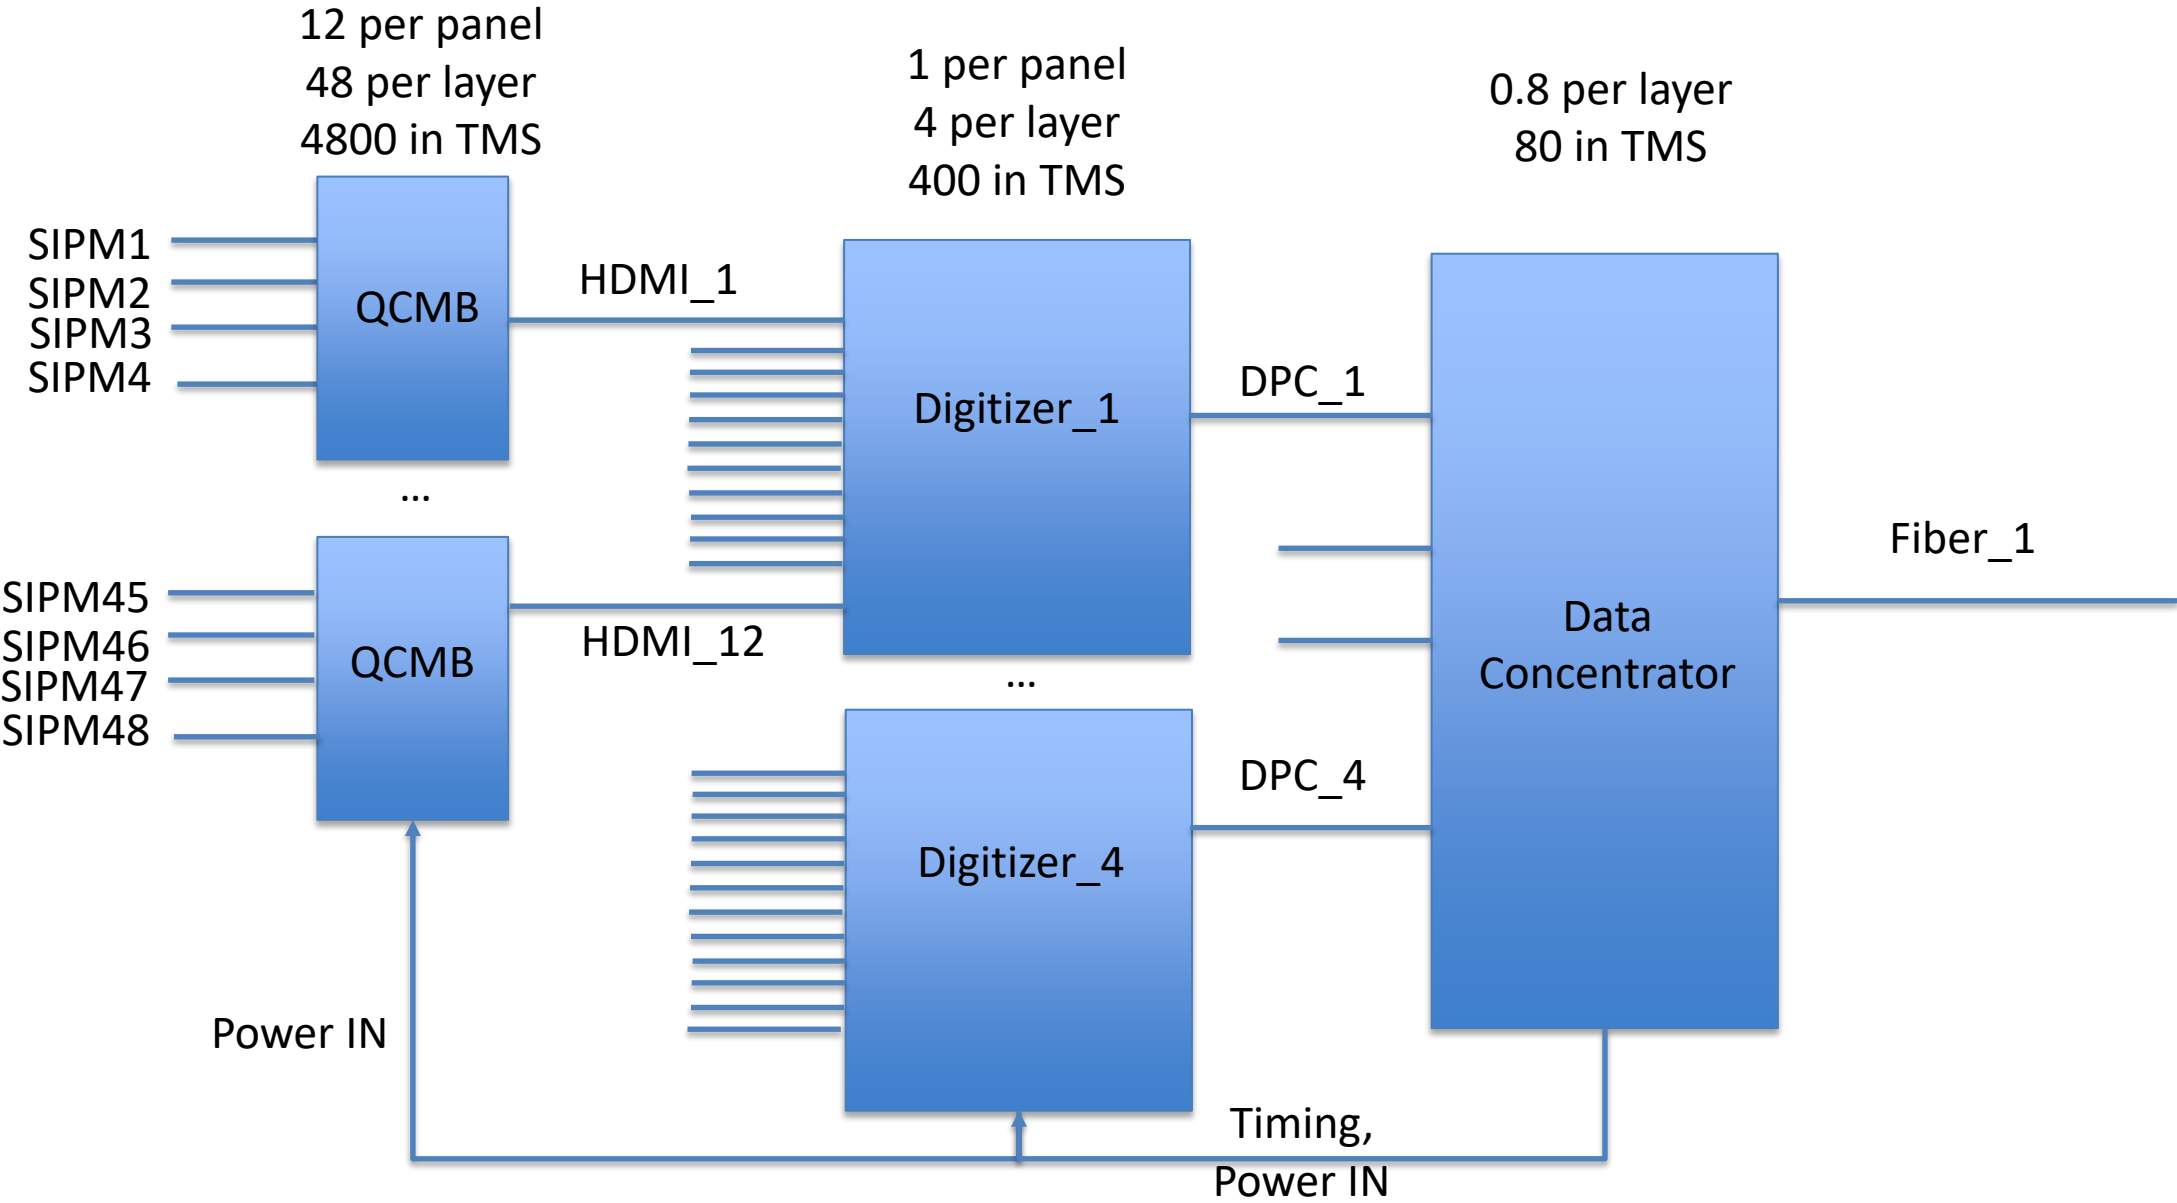
Electronics Power



- This will work, and can be costed. As the electronics design has evolved, some legacy features look kind of goofy. We have time to change things.
- 120VC AC powers ethernet switches (w/Power Over Ethernet)
- The Data Concentrators – in the same racks as the switches - are powered by PoE and send power to the digitizers on D-sub connectors
 - One could also send timing this way (or through the DisplayPort cable)
- The Digitizer Boards send power to the QCMBs where the SiPMs are located.
- Aside: all float with respect to the building ground – for us, the PRISM rails.



Block Diagram



Pros and Cons

- Pros
 - We know this will work – and reasonably well
 - Mu2e has already done most of the R&D for us
- Cons
 - Much of the electronics is inaccessible - what if it breaks?
 - There is no provision to gain match the SiPMs by adjusting their voltage
 - Setting up the minimal test stand supports a lot of channels (240)
 - Data volume is dominated by SiPM noise
 - Cosmic rays (especially east-bound) look different than CC muons
- .

Analog Front End Alternatives

- “Type 1A”
 - The TI ultrasound chip you just heard about
- “Type 1B”
 - An updated version of this chip
 - One exists that is faster (2X the Main Injector RF). 16 channels and originally dismissed as too expensive (2.1x the price of 1A). It’s now at 1.5v, “uses less support circuitry”, whatever that means
 - TI might be trying to move people to this chip
- “Type 2”
 - The KlauS ASIC from Germany

I lead off with this because it influences other decisions .

Because it was designed for calorimeters, the more we are interested in energy as well as position, the better it will suit us.



Comparison

	AFE5807	KlauS
Cost	\$170,000 (base)	Potentially in-kind from Germany
Channel count	8	32 (need 48)
ADC	12-bit	10/12-bit
Sample Time	12.5 ns	Sample and hold "0.2 ns"
Deadtime	Deadtimeless	Average 3-4% (but adjustable)
System Power	1 kW	< 100 W
Input	Single-ended	Single-ended
Support Circuitry Impact	Unknown	Unknown

- Green is better, gray is good enough, red is potentially troublesome

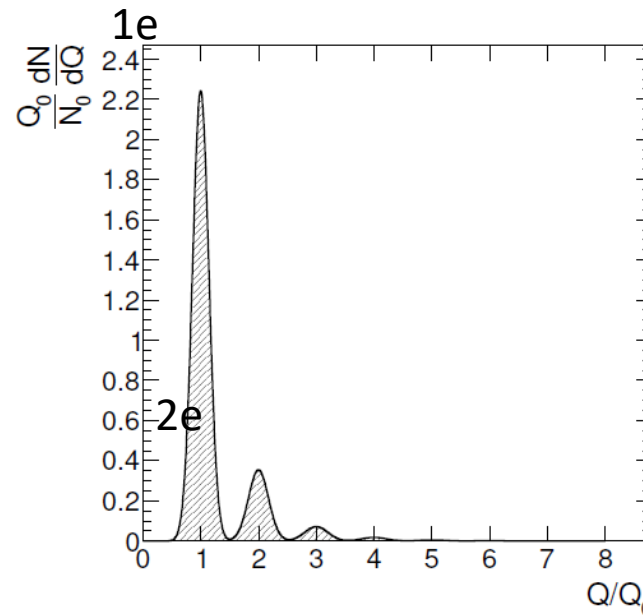
Deadtime and the Operating Point

- Originally I had not considered deadtime
 - Commercial front-end chips can handle our rates with no deadtime
 - However, noise is still noise – signal pulses can be distorted
- I would rather have a 99% efficiency than a 1% deadtime
 - An efficiency is just that – deadtime is correlated with beam intensity, particle location, etc.
- We would probably operate KlauS at a higher threshold than the TI AFE5807
 - Lowers deadtime (at a small cost in efficiency) → that's the tradeoff
 - Looks like ~4 photoelectrons (~¼ mip) would be where we end up – a fraction of a percent of each (not the 3-4% at the lower threshold)




How Noisy are the SiPMs?

- All the SiPMs we have looked at have noise (dark count rate) in the 0.5-1.0 MHz ballpark
 - For a 2 bucket-wide pulse, that's an "occupancy" of 2-3%.
 - "Occupancy" is for at least 1 photoelectron
 - At 2 pe, it's more like 0.4%
- We don't know the noise spectrum for any of the candidate SiPMs.
 - Much less coherent noise
 - Reminder: single-ended front-end chips are more susceptible to coherent noise



arXiv:1710.11410



Deadtime And Cheating



- Can we cheat and reduce/eliminate the KlauS deadtime?
 - A panel box contains 16 extra channels – can we use them somehow?
 - Seems non-optimal to have deadtime and unused channels at the same time if it can be avoided.
 - Can we do analog shaping of the signal?
 - We need an output time-stamped every 19 ns, but we can wait microseconds or longer for it
 - Double-up the SiPMs
 - Add an analog coincidence circuit (they exist – and go back to the days of vacuum tubes)
 - Doubling up the SiPMs is on the edge of what is financially possible; doubling the channel count is probably a whole ‘nother kettle of fish
 - Other ideas?

QCMB Alternatives

- Obviously we need something to do this - -our discussions have revolved around where to put them.
- Commercial and custom “light cables” are in the \$2M range
 - See Andy’s previous talk
 - The odds of getting this \$2M from DOE are slim – we would need another funding source.
- SiPMs are not phototubes – experience (or at least lore) is that they are much less likely to die.

TMS and “Smart Sensors”

- TMS Panels are only a little more accessible as the South Pole. Could the 400 TMS Panels be self-configuring and self-calibrating, with only a single RJ-45 connector? (Power over gigabit ethernet)

Almost

- Placing a small computer (like an ARM Cortex – 100’s of VAX equivalents) on the digitizer board is a few dollars. Some other elements of the system might need to be reworked, but this should be possible. **The unresolved issue is timing.**

- Timing:

- Needs to drift less than ~half an RF bucket over a second: 10^{-8} or so (this is an oversimplification: we should be using Allan variances)
 - Ethernet cannot be used to send timing signals this precise
 - Smallest packet is 19 RF buckets long
 - Architecture (CSMA-CD) doesn’t support precision timing
 - Therefore this depends on a good local clock
 - With 400 clocks talking to each other (via an NTP derivative like PTP or White Rabbit) we need individual clocks good to a few $\times 10^{-7}$. This is in the range of TCXO’s.



Smart Sensors II


- Smart Sensors were a DOE initiative a few years back
 - I presume that pot of money is spent
 - This is nonetheless a reasonable Research program
- Doesn't this move in the wrong direction – bury even more electronics?
 - Maybe – it does that, but it also reduces connector counts (a less reliable component), has shorter signal runs, etc.
 - I don't think it's immediately obvious whether this is a net winner or net loser
- It certainly makes checkout and QA/QC smoother
 - One click and you can run your tests.
 - Any laptop can be a test stand!
 - Probably means more reliable components go into TMS – as said, it's not obvious if this is a net winner or net loser

Replaceable Digitizers



- What if we put the digitizers on sockets – maybe we could replace one with some custom tool?
 - Uh....
 - This sounds hard. Real hard.
 - It's also not clear that the socketing wouldn't add to the unreliability
- What might work is to pre-install spares – 5 panels of readout per 4 panels
 - Boot configuration could select which units to use
 - Cost would be about \$250K, plus additional mechanical costs and cables
 - I'd like to run these on spare cables, so as to protect against cable/connector issues.

PCIe Data Concentrators



- There is no longer a good reason to make data concentrators as custom PCBs in their own enclosure.
 - When we started this, very few CPUs had enough PCIe lanes to run this – no longer the case
|
- The job looks like a firewall – take data on one side, look at it, maybe tweak it, and send it to the other.
 - There exist commercially many-lane PCIe enclosure to do this for cloud computing – about \$8K/enclosure with a “DPU” – data processing unit
- Moving to a set of PCIe cards probably won’t save any money (dominated by engineering) but it makes it easier to get small units in the hands of the Consortium”
 - Minimum is probably one panel, not 5/



Summary

- Now is the time to get going on the electronics
- I believe both the TI AFE5807 (and variants/successors) and KlauS can be made to work as front-end chips, although neither is perfect.
 - The biggest issue with the AFE5807 is speed
 - The biggest issue with KlauS is deadtime
 - Neither seems to be a show-stopper, but both have physics impact
 - There are some other smaller annoyances with each option
 - But we need to settle this soon – it drives other decisions
- There are other decisions to make and they are coupled – doing X here can preclude doing Y there.

Backup


Why Does Sample Time Matter?

- Signal formation time is ~30-40 ns.



- With 106 MHz, you can do so-called “optimal filtering” – take the four or five measurements and quickly convert them to ΣQ , t_0 , and χ^2 .
 - With 80 MHz (AFE5807), you can do this, but “even” events look different than “odd” events – offset by $\frac{1}{2}$ an RF bucket. Probably more trouble than its worth.
- How to do this with Klaus’ sample and hold remains to be seen.


Noise & Efficiency



- TMS measures momentum by range
 - Missing the last hit causes us to underestimate the momentum → keep the threshold low
 - Adding an extra hit at the end causes us to overestimate the momentum → keep the threshold high
 - Both can be corrected for statistically, but small corrections are better than large ones
- An important decision will be the operating point
 - Likely to be chip dependent.

If the TMS is a pure downstream tracker, this is not too complicated – we'll end up at around $\frac{1}{2}$ mip. (Zero is too low...one is too high) If we start looking at dE/dx , this gets more interesting.

Cables




Cables to the Digitizer:

Base cost: \$22,038
Contingency: 20%

Cables to Digitizer – HDMI, four pairs, one per SiPM

Cable cross section:



Cables from the Digitizer:


Base cost: \$31,005
Contingency: 15%

Cables from Digitizer – DisplayPort, three pairs used, one per 16 SiPMs

Electronics Labor

- What we do and do not have

- The estimate described on the last slide is based on functional blocks. We have a BOM for each block.
- Which blocks go on which boards and how the signals go from one to the other is a physicist-level (me)
- We have an engineering estimate on what engineering needs to be done
 - We don't have a final design, which means we don't have a layout, which means we don't have a vendor estimate on what it costs to make and stuff a board.



- Programming

- FPGA programming is costed as Lab EE effort.
 - There have been bad experiences with assigning this to universities
- “Regular” programming is uncoded scientific labor
 - The majority of this is interfacing with the DAQ. This will be largely dealing with exceptions and edge cases.
 - To remind people, the DAQ does not do data acquisition. We do. The DAQ does event building and writing to storage. The Data Concentrators are essentially our DAQ.
 - The risk mitigation will be “offload to Labs if necessary”, but this probably isn't the right answer. If the collaboration can't find someone to write the DAQ software, is that a project problem or a collaboration problem?