CALICE ECAL Readout Electronics: CDR Report

October 16, 2002

A summary of the issued raised and the discussion of them follows.

1 System-wide issues

1. Transmission of bussed signals on the backplane. It was thought LVDS signals on the bus might cause problems with jitter and slew. "Bus LVDS"[1] was mentioned as a possibility; this features similar voltage swings, but provides increased current drive to handle the double terminations required in multi-point applications.

   The more critical signal on the backplane from the timing perspective is the trigger as this sets the sample-and-hold timing. Even here, as the sample-and-holds for each cable are adjustable separately, then slew (as opposed to jitter) can be corrected cable-by-cable. In fact, the trigger could be considered as the "system clock" to a large extent.

   The other signal is the 12.5 MHz system clock; this is less important as the individual readout boards do not have a requirement to run exactly synchronised. In fact, the whole system would work with each readout board having its own separate clock and this option has been considered.

   It was considered "good housekeeping" to keep the trigger (and clock to some extent) solid. At least the trigger should be transmitted point-to-point from the trigger board to the readout boards. There are 64 spare pins for standard VME on J2; point-to-point LVDS to the 19 readout board slots in a crate would take 38 pins of the trigger board connector. Clearly, also having the clock point-to-point would take 76 pins, more than available so only the trigger or the clock could be sent this way.

   Having point-to-point connections clearly forces one slot only for the trigger; it was suggested that this should be central in the crate to keep the trigger paths as equal as possible. This may have some influence on servicing of interrupts if both the readout and trigger boards generate interrupts on a trigger.

2. Event buffering. Almost all beam lines have a bunch structure where the beam is delivered for a time, usually of order 1 sec, and there is then a dead period of around a factor of 10 longer. Hence, for the required sustained rate of 100 Hz, a peak rate of 1 kHz must be achievable. This clearly means buffering events during the bunch and then reading during the dead period would ease the rate requirement enormously.

   However, the proposed system has no event buffering and so needs to read every event before the next trigger. This is because buffering events would require significant extra complexity in the trigger distribution system as event numbers to tag the data for each event would be needed, or else a missed trigger in a board would cause all subsequent events to be misconstructed. While this would be straightforward to do within the crate

[1] See `http://www.national.com/an/AN/AN-971.pdf`

of ECAL readout electronics, it is not clear if it would be so for the other (undefined) systems which need to be included in the DAQ, i.e. the HCAL, trigger and beam monitor. In addition, these systems may not allow multiple buffering in any case. Certainly, it is unlikely the pre-installed systems in any beam line would allow multiple events before readout.

These problems were considered significant. However, one suggestion was that the readout boards could still store multiple events during the bunch time but have the other systems read out event-by-event. This would be useful as the ECAL is likely to be 20 of the 25 kBytes per event and so this would save 80% of the data from being read. To ensure no trigger has been missed, a small amount of data (such as a trigger number from an internal board counter) could be read from each board. This ensures no trigger has been lost and hence that the events can be reconstructed correctly.

There are other possibilities for speeding up the readout. One would be to split the system between two crates and have separate PCI-VME interfaces for the two. Ideally, these would go into the same PC to simplify the DAQ. However, there is no experience with multiple PCI interfaces in one PC so it was not clear if this would double (or even increase at all) the total VME data rate. This depends on the details of the PCI interface within the PC. The proposed PCI-VME PC-end cards have RAM buffering, so the data rate may improves as the VME-Interface card DMA transfer is in real-time, while the Interface card-PC can be done later. This would need testing.

3. Trigger latency. The trigger has to be generated from scintillators in the beam, put through discriminators and some NIM (or similar) logic, sent to the trigger board, distributed on the backplane of the crate, be received (and possibly delayed) by the readout board, transmitted as the sample-and-hold down the cables to the VFE-PCB's and then propagated across the VFE-PCB to the VFE chip, where it is used. The total path has to take 180 ns in order than the shaped signal is sampled at the peak.

   Concern was shown as to whether this would be possible, particularly if the scintillators where some distance upstream of the calorimeter. No detailed estimate of the time expected has been done. There are no obvious sources of a significant delay in the trigger and readout boards. However, because the beam line to be used is not known and the the trigger logic has not been designed, it is unfortunately not known what the total latency will be. Almost by definition, the trigger must be built close enough to the calorimeter to ensure the latency from signal propagation is not too long.

4. Test board. Although not discussed in detail during the review, the test board was mentioned briefly as an "inverted readout board" where, roughly speaking, all ADC's are replaced by DAC's and vice versa. This would allow production testing of all signals on the cables from the readout boards. The current proposal has all six cables of the test board populated so that a complete readout board can be tested at one time, with all possible combinations of cross-talk being measurable.

   This was considered to be a "Rolls Royce" solution. Although only costing around £3k (out of £100k total), it was thought to need around 6 months of engineering design effort in addition to 1 month of layout technical effort, even when leaning heavily on the readout board design. The cost could be reduced, e.g. by only having one cable rather than six testable at one time. However, the main limitation would be the effort needed and time in the drawing office. It would be most efficient for the test board to reuse as much of the layout of the readout board design as possible. Alternatively, one (or six, if there is space) cable output could be built into the trigger board and so reuse the VME interface firmware design as well as only requiring two boards to be laid out in total.

5. Overlap between current system and long term linear collider electronics. There was concern that the long term direction for the electronics was not being considered while designing the current system.

   The electronics for any final calorimeter would be completely different from what is needed for the beam test. Here, the aim is to test only the calorimeter itself; the VFE chip is nothing like the chip that would be needed eventually and so the downstream electronics similarly will be very different. Any long term developments should be considered as a separate project.

6. Altera vs Xilinx. The University groups all have Xilinx design tools. RAL have used both and have no strong preference about which they prefer, although Xilinx is possible a little ahead at present. Europractise includes design tools for both.

   Altera allows smaller components to be used in the footprint of larger devices, if it is found that only the smaller component is needed after the board is fabricated.

7. Crate power 3.3V or 5V. The 3.3V is natural for the Spartan-II FPGA's and bulky regulators may be needed if 5V is used. However, only 5V is standard for VME. There may also be complications with the ramp of the Spartan-II at power-up under 5V.

8. VME64 and VME64x. VME64 works through having all lines bidirectional. It sends an A32 address and then uses both the A32 and D32 sets of lines to send 64-bit data words. Hence, the rate is doubled compared with standard VME and around 50 MByte/s should be achievable. This would make the 1 kHz rate requirement much easier to match. VME64x uses further extensions but would not be needed for this system.

   RAL have some experience and are built their own crates. A VME64 backplane costs around £1k and the total crate around £5k. An off the shelf crate was thought to be less than £10k. The PCI-VME interface card for VME64 would also be more expensive than assumed. VME64 crates come with geographical addressing and both 3.3V and 5V as standard, which solves the issue with which to use. There is no standard bus for user signals so the crate would still need customisation; however, the central 1mm-pitch connector area is often used bussed, so these would be the obvious pins to use.

   CompaqPCI, which also has much higher rates than standard VME, has been used at RAL but was not strongly recommended. It is harder to use if starting with no experience. Also, despite general opinion, it was thought unlikely to become the future standard as there are already competing protocols with much higher rates.

   There seems to be a very strong incentive to move to using VME64.

9. Schedule and effort. The 5 months effort allocated to layout and fabrication was considered too low; 8 months was thought a more reasonable figure if the boards were sufficiently similar that parts of the layout could be reused. In addition, the scheduled 2 months (including the design reviews before the layout) for this task was also thought too small. The layout should start towards the end of the design period and take around 3 months in total.

   Ideally, the layout should be done iteratively, with a preliminary layout done early to check for board space. This is not always possible, given the way the RAL drawing office works. In addition, the drawing office is known to be heavily booked for the first half of 2003 already. It might be worth putting the trigger board in first, if there is sufficient overlap to mean this gets a rough layout of the readout board also.

2 Trigger board

1. Logic device choice. The UCL engineers have experience of using a Mach5, 256-pin PLD device and have a working VME interface for it. This component is quite old but seems sufficient for the trigger board. However, would a Spartan-II be preferable?

   It was thought that an FPGA would be much more powerful and allow greater robustness, given the uncertainties in the trigger board requirements. Using the Spartan-II to drive LVDS directly was considered to be useful as it gives a lot of flexibility for timing and jitter issues and otherwise the board will have a lot of buffer components. However, driving LVDS directly clearly doubles the number of IO pins required, which might be a problem. Note, the Spartan-II should not be used to drive VME directly and VME buffers were considered essential.

   As for voltage requirements, the Mach5 runs at 5V while the Spartan-II (but not the Spartan-IIE) is designed for 3.3V but is "5V tolerant"; it requires a DC-DC converter, 5V to either 3.3V or 1.8V. On the trigger board, there should be space for this component.

2. Reuse of readout board VME interface. The layout time would be significantly shortened if the boards could use the same VME interface. There are issues of the readout board being substantially more complicated than the trigger board, both in VME interface and PCB number of layers, etc. This may result in the trigger board FPGA being much bigger than needed and also the PCB being more complex and expensive. This may be offset to some extent if the readout board layout connected with the master FPGA can be restricted to a subset of layers, so these would be the only ones required for the trigger board. In any case, as only four trigger boards are needed, the extra complexity might not be a significant expense compared with the effort saving.

3 Readout board

1. Merge slaves into master? The issue of whether the slaves needed to be separate FPGA's from the master was raised. In principle, all the functions of both could be done in one FPGA, and so remove the need for the configuration bus and fast data path. However, the master FPGA is already very heavily I/O limited at around 220 pins and merging would certainly increase this count unless a large amount of functionality (or flexibility) was dropped. In addition, the designs are being done at separate institutes and so would make the parallel design work harder to coordinate.

2. Updates of slave FPGA firmware. It was thought to be useful to be able to update the slave firmware through VME, as well as by attaching a header. The boards will have a connection from the master FPGA to the slaves JTAG ports for the loading of new firmware into these slaves via VME and hence from the host (DAQ) CPU. It was thought that complications of wrapping the firmware code to allow this transport should be straightforward.

3. VFE-PCB interface. There was unanimous agreement that using non-standard "LVDS" shifted down by -5V was extremely unusual and could cause all sorts of complications throughout the design. It would be highly unlikely that such an offset could be accomodated in the other potential uses of the board for the HCAL option readouts.

   The reasoning behind this choice, being due to a PMOS input stage in the VFE chip, was understood but it was thought highly unlikely this would be the limiting factor in the noise performance of the system as a whole. There was very strong support for having the VFE chip run between +5V and 0V (rather than the proposed 0V and −5V) and the alternative should only be considered as a very last resort.

Given that the VFE-PCB is operating at negative voltages, it was not understood how the analogue differential voltage levels of $+2 \pm 2$V would be possible. This needs to be clarified. "Proper" signal termination, rather than open termination, for this line was thought a good idea, even with short (2m) cables. This might prevent problems with beam pickup, etc, in the noise beam line environment.

The grounding for the VFE-PCB's and the readout electronics needs to be thought through and a grounding scheme document should be produced. Even assuming standard LVDS, the operating range for common mode around the nominal 1.2V is limited so the grounds need to be kept quite close. The two systems cannot be allowed to float but must have the grounds connected somewhere. This should not be through the cable itself (in case of ground loops) but preferably through a connection between the VFE-PCB power supplies and the readout VME crates. Ideally, this means the power supplies and crates should be close to each other; the distance betwen them is not yet known.

The proposed SHL connector set was not considered ideal for the reasons outlined in the document. The 4.5mm clearance at the VFE-PCB end was seen as a major limitation, but further investigation of alternatives would be a good idea. It might also be a good idea for the VFE chips to be on both sides of the wafers to minimise the track length of the signals. This would complicate the cabling to the readout boards and might increase the assumed 2m length, but would not be impossible. It is not clear if the tungsten support structure would allow the VFE-PCB to come out at both sides.

It is clear there issues would have a major impact on layout and so it was thought that they must be settled as soon as possible and treated as a matter of urgency.

4. Synchronisation issues. There was a long discussion on whether the slaves need to run synchronously. Note, they all run on the board clock (so the system is "isosynchronous"), but the sequence timing edges could be on different ticks of this clock.

   It was agreed that each sample-and-hold should be separately adjustable because of the tight timing requirement and the possibility of variation in the shaping times of the VFE chips. The trigger board uses digital delay units and it may make sense to use them here too. These provide ns-accurate delays; it would require an extra six (not large) components per readout board and would add 10ns to the latency. For example, a DS1020[2] would take three extra IO pins per slave FPGA.

   However, the timing of the rest of the readout sequence, which drives the VFE multiplex shift register, is much less critical and so could in principle be run synchronously. One issue is the variation in the ADC busy times, as they run off an internal clock and so can finish at different times. The current design assumes the next multiplex clock is generated when the ADC busy line (actually the OR of all six ADC busy lines) goes off. Having them all synchronous would require instead that the multiplex clock is set for a fixed time after the maximum possible ADC busy. This would clearly mean the multiplex sequence would be slower but by how much was not known.

   The main advantage of being synchronous would be on the fast data path to the master FPGA dual-port RAM. In the current design, this would have to have six independent logic blocks for receiving the data from the six slaves. If they were synchronised, all the duplicated control logic could be removed. In fact, synchronising the data transfer but leaving the slaves unsynchronised would in principle solve most of the issues. However,

[2]See `http://www.maxim-ic.com/quick_view2.cfm/qv_pk/2606`

the potential use of the readout boards for the digital HCAL option causes a complication. This option has to have zero-suppression, so the amount of data going through each of the slaves to the master will be different. Hence, if it is intended to keep the ability to read the digital HCAL, then at least variable amounts of data per slave must be accomodated in the slave-master interface.

Another potential advantage of synchronising all the slaves is that this would mean all signals on the board are clocking at the same time during every sample. If the analogue lines or ADCs are picking up noise from the board, then synchronising all parts means the same noise is picked up for each sample.

However, synchronising the slaves would require a major change to the design. As the sequence timing is currently assumed to be configurable then the only way to force synchronised slaves is for there to only be one source of configuration data. There are several ways to achieve this; obvious choices are by forcing all slaves to receive the same data by only sending it broadcast on the configuration data bus or by having a single slave control all six cables. In all cases, the cable signals must be inidividually controlled and not be driven from a single source through a signal splitter. This is again because of the digital HCAL option, where separate data may need to be downladed to the HCAL boards via the LVDS lines.

5. ADC choice. The ADC is a major fraction of the readout board price, being 1/3 of the cost and £20k in total. It is therefore important to be sure of the component and cost soon. Clearly, given the tight space available on the proposed 6U boards, a dual ADC component would be very useful. There were not thought to be fast (500kHz), 16-bit dual packages available, although a draft data sheet for a new component was circulated at the review; a Texas Instruments ADS8361[3]. However, it is not known if this will be readily available and cheap enough on the timescale of this project. Otherwise, the proposed AD7664 was thought to be a reasonable choice.

6. Cost estimate. The fabrication and assembly costs for the readout board of £850 were thought too low. For the production run of 20 boards, then assembly alone would be around £800, with fabrication being around another £200. For the prototype run of only 2 boards, assembly could be £1000.

[3]See `http://focus.ti.com/docs/prod/folders/print/ads8361.html`